# Configuring VMware Aria Operations

FEBRUARY 2024

VMware Aria Operations SaaS

vmware by Broadcom

You can find the most up-to-date technical documentation on the VMware by Broadcom website at:

https://docs.vmware.com/

**VMware by Broadcom**
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to https://www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. Copyright and trademark information.

# Contents

Configuring VMware Aria Operations 18

1 Accessibility Compliance 19

2 Collecting Data with Cloud Proxies 22

Configuring Cloud Proxies 23

Activating Data Persistence in Cloud Proxy 25

Configuring Cloud Proxies in AWS and Azure 26

Monitoring the Health of Cloud Proxies 27

Deleting Cloud Proxies 31

Upgrading Cloud Proxy 31

Using APIs with VMware Aria Operations Cloud Proxy 34

Cloud Proxy FAQ 35

Cloud Proxy Troubleshooting 37

Configuring Collector Groups 41

Adding a Collector Group 45

Editing Collector Groups 46

3 Configuring Policies 48

Policies 48

Policy Decisions and Objectives 50

Policies Library 51

Operational Policies 53

Types of Policies 54

Custom Policies 54

Default Policy in VMware Aria Operations 56

Policies Provided with VMware Aria Operations 56

Using the Policy Workspace to Create and Modify Operational Policies 57

Policy Workspace in VMware Aria Operations 59

Assigning Policies 83

4 Integrating Data Sources with VMware Aria Operations 86

The Integrations Page 93

The Repository Tab 94

The Accounts Tab 95

Adding Accounts 98

Exporting and Importing Accounts 98

Configuring Credentials in Integrations 99

vSphere 102
Configure a vCenter Server Cloud Account in VMware Aria Operations 103
Configure User Access for Actions 110
Cloud Account Information - vSphere Account Options 111
VMware Cloud on AWS 113
Configuring VMware Cloud on AWS in VMware Aria Operations 113
Configuring VMware Cloud on AWS Government Cloud Endpoint in VMware Aria Operations 119
VMware Cloud on AWS Outposts 124
Configuring VMware Cloud on AWS Outposts in VMware Aria Operations 125
Azure VMware Solution 131
Configuring an Azure VMware Solution Instance in VMware Aria Operations 131
Known Limitations 135
Oracle Cloud VMware Solution 135
Configuring an Oracle Cloud VMware Solution Instance in VMware Aria Operations 136
Known Limitations 139
Google Cloud VMware Engine 140
Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations 140
Known Limitations 144
VMware Cloud Foundation 144
Configuring VMware Cloud Foundation Cloud Account in VMware Aria Operations 145
VMware Infrastructure Health 147
Configuring the VMware Infrastructure Health 147
Monitoring the Health of VCF Deployments 148
Monitoring the health of vCenter Server 149
Monitoring Health for NSX 150
VMware Cloud on Dell EMC 152
Configuring a VMware Cloud on Dell EMC Instance in VMware Aria Operations 152
Known Limitations 156
Generating CSP API Token for VMware Cloud on Dell EMC 156
Verify that the NSX Adapter Instance is Connected and Collecting Data 157
Amazon Web Services 158
Supported AWS Services 158
Charges for AWS Metrics 163
Configuring AWS 164
View AWS Objects 175
Microsoft Azure 176
Supported Azure Services 176
Charges for Azure Metrics 180
Configuring Microsoft Azure 181
View Objects for Microsoft Azure 186
OS and Application Monitoring 187

Supported Application Services 189
Supported Platforms 191
Prerequisites 192
Steps to Monitor Applications 198
Troubleshooting 288
Monitoring Application Services and Operating Systems using Open Source Telegraf 296
Monitoring Application Using Open Source Telegraf 297
Troubleshooting (Open Source Telegraf) 346
Monitoring Physical Servers 346
Service and Application Discovery 346
Supported Platforms and Products for Service Discovery 347
Supported Services 348
Configure Service and Application Discovery 349
Manage Services 353
Create Application Definition 355
Manage Applications 356
View Applications 357
Discovered Services 358
Discovered Applications 359
Discovered Rule-Based Applications 360
Alert for Service Unavailability 361
Service Discovery Metrics 361
Configuring Business Applications 363
The Business Application Page 363
Add Business Application 365
Configure, Activate, or Deactivate Synthetic Monitoring 368
Application Integration 369
Application Discovery 370
Integrating Applications 370
Configuring Ping Adapter Instances 371
VMware Aria Automation 372
Advanced Workload Placement for Allocation Model 373
Integrate VMware Aria Automation SaaS Service with VMware Aria Operations SaaS Service 374
Importing Accounts from VMware Aria Automation 375
Supported VMware Aria Automation Versions 376
Object Types 376
Workload Placement 377
Pricing for VMware Aria Automation Components in VMware Aria Operations 378
Managing Public Cloud Endpoints with VMware Aria Automation Integration 379
Cloud Zones in VMware Aria Operations 380
vSAN 381

Configure a vSAN Adapter Instance 381
Verify that the Adapter Instance is Connected and Collecting Data 383
vSAN Log Analytics Enhancements 385
VMware Aria Operations for Networks 386
NSX 386
Configuring the NSX Adapter 387
Support for Principal Identities Authentication for the NSX Management Pack 390
Configuring Cloud Federation Adapter 390
Google Cloud Platform 391
Supported GCP Services 392
Charges for GCP Metrics 393
Configuring Google Cloud Platform 394
View Objects for Google Cloud Platform 398
Integration of VMware Aria Operations and VMware Aria Hub 399

5 Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria Operations 401
Activating VMware Chargeback Capabilities 402
Register VMware Aria Operations in a VMware Cloud Director Instance 403
Migration of Chargeback Data to VMware Aria Operations in Case of an Upgrade 403
Managing Chargeback Administration Settings 405
405
406
406
407
Configuring VCD Pricing Details 407
Creating Notification Rules for Tenant 420
Configuring Tenant Email 421
Managing Chargeback Reports 422
Generating Bills 422
Scheduling Bill Generation Using Automation Central 424
Viewing Chargeback Summary 426

6 Optimizing Capacity and Improving Performance 427
Capacity Optimization Concepts 428
How Does Capacity Optimization Work 428
Calculating and Forecasting Capacity 429
Allocation and Demand Model in Workload Optimization 434
Policy Settings for Capacity 435
How to View and Assess Capacity 435
Assessing Capacity in the Capacity Page 435
Viewing Object Capacity in the Capacity Tab 436

Predefined Dashboards for Capacity 443
How to Optimize Capacity and Improve Performance 443
Using the Capacity Page to Asses and Optimize Capacity 443
Using Reclaim to Free Up Resources 447
Using Rightsize to Adjust Resource Allocation 453
Using Workload Optimization to Improve Performance 456
How to Plan for Capacity Changes 478
What-If Analysis: Modelling Workload, Capacity, or Migration Planning 478
What-If Analysis - Workload Planning: Traditional 481
What-If Analysis - Workload Planning: Hyperconverged and VMC on AWS 486
What-If Analysis - Infrastructure Planning: Traditional 490
What-If-Analysis - Infrastructure Planning: Hyperconverged 493
What-If-Analysis - Migration Planning: VMware Cloud 495
What-If-Analysis - Migration Planning: Public Cloud 498
What-If Analysis - Data Center Comparison 500
Retain Historical Data of VMs Migrated Using VMware Hybrid Cloud Extension 502

7 Configuring Cost 504
Cost Overview 504
VMware Cloud on AWS and VMware Cloud on Dell EMC Cost Management in VMware Aria Operations 507
Google Cloud VMware Engine Cost Management 508
Reference Based Costing for Azure VMware Solution/Oracle Cloud VMware Solution 510
Cost Analysis 510
Analysis of Cost Metrics 511
Analysis of Price Metrics 512
Analysis of VMware Cloud Bills 513
Cost Analysis - Metrics and Metric Keys 513
New or Edit Cost Analysis 516
Saved Analyses 518
Cost Settings for Financial Accounting Model 519
Overview of Cost Drivers 521
Import or Export Cost Drivers 525
Cloud Providers Overview 526
Add or Edit Cloud Provider 526
Billing Enhancements for Horizon Management Pack and Virtual Hosts 527
Editing Cost Drivers 527
Editing Server Hardware : Traditional 528
Editing Server Hardware: Hyper-Converged 529
Edit Monthly Cost of Storage 530
Edit Monthly Cost of License 531
Edit Monthly Cost of Maintenance 533

Edit Monthly Cost of Labor 534
Edit Monthly Cost of the Network 535
Edit Monthly Cost of Facilities 536
Editing Additional Costs 537
Edit Application Cost 537
Cluster Cost Overview 538
Cluster Cost Computation with Allocation Model 539
Editing Cluster Cost Calculation Methods 540
Publish Daily Cost Metrics for Virtual Machines 541
Publish Tag Based Cost as Individual Metrics 542
Pricing Overview 542
Add New Pricing Card 543
Migration of vCenter Pricing Cards to vCenter Pricing Policies 545
Cost Calculation Status Overview 546
Costing Enhancements 546
Reclaimable Hosts Cost Metric 547
Realized Cost Savings Using Reclamation Suggestion 548
Costing for Oversized VM and Undersized VM 548

8 Viewing and Configuring Compliance 550
Measuring Compliance of Objects 550
VMware SDDC Benchmark Details 553
Regulatory Benchmark Details 557
VMware Cloud Foundation Benchmarks based on VMware Cloud Foundation Compliance Kits 562
Compliance Score Cards 563
Compliance Alerts 566
How To Configure Compliance Benchmarks 566
Activate VMware SDDC Benchmarks 566
Activate VCF Benchmarks 567
Create a New Custom Benchmark 568
Import or Export a Custom Benchmark 568
Install a Regulatory Benchmark 569

9 Configuring Green Score to Track Sustainability 571
Configuring Green Score Sustainability Data at the Organization Level 571
Configuring Green Score Sustainability Data for Physical Data Centers 577

10 Configuring Automation Jobs 585
Automation Central 585
Create Job from Automation Central 587
Create Job from Reclaim or Rightsizing 592

11 Log Analysis with VMware Aria Operations for Logs 594

12 Configuring Alerts and Actions 598

Understanding Alerts 599

Accessing Alerts 600

Viewing Alert Information 605

Intelligent Alerts 607

Alert Definitions 609

Creating Alert Definitions 612

Symptom Definitions 621

Defining Symptoms for Alerts 621

Understanding Negative Symptoms for Alerts 638

Recommendations 639

Defining Recommendations for Alerts 641

Notifications 642

Creating Notification Rules for Alerts 644

User Scenario: Create an Email Alert Notification 650

Notifications - User Scenario: Create a Webhook Alert Notification 651

Outbound Settings 653

Adding Outbound Notification Plug-Ins 655

Exporting and Importing Outbound Settings 665

Payload Templates 666

Creating Payload Templates for Outbound Plugins 667

Managing Alert Groups 670

Grouping Alerts in VMware Aria Operations 671

Deactivating Alerts in VMware Aria Operations 671

Alert Definition Best Practices 672

Create a Simple Alert Definition 673

Create a New Alert Definition 674

Create an Alert Definition for Department Objects 675

Add Description and Base Object to Alert Definition 676

Add a Virtual Machine CPU Usage Symptom to the Alert Definition 677

Add a Host Memory Usage Symptom to the Alert Definition 679

Add Recommendations to the Alert Definition 680

Create a Custom Accounting Department Group 682

Create a Policy for the Accounting Alert 683

Configure Notifications for the Department Alert 684

Create a Dashboard to Monitor Department Objects 686

Actions 687

Viewing Actions List 690

Actions Supported for Automation 691

Integration of Actions with VMware Aria Automation 693
Working with Actions That Use Power Off Allowed 693

13 Configuring Objects 697
Object Discovery 697
About Objects 699
Managing Objects in Your Environment 701
Managing Custom Object Groups 715
Managing Application Groups 725

14 Configuring Dashboards 729
Accessing Predefined Dashboards 731
Types of Dashboards 731
Custom Dashboards 731
Create and Configure Dashboards 731
Dashboard Name 732
Widget or View List Details 732
Widget and View Interactions Details 734
Dashboard Navigation Details 734
Manage Dashboards 735
Manage Summary Dashboards 738
Auto-Rotate Dashboards 738
Manage Dashboard Folders 739
Share Dashboards with Users 740
Dashboards Actions and Options 741
Options for Sharing Dashboards 741
Manage Widgets in Dashboards 742

15 Predefined Dashboards 743
Application Monitoring Dashboards 746
Linux OS discovered by Telegraf Dashboard 746
Windows OS discovered by Telegraf Dashboard 747
Availability Dashboards 748
VM Availability Dashboard 748
vSphere Availability Dashboard 750
Ping Overview Dashboard 751
Capacity Dashboards 752
Cluster Capacity Dashboard 753
Datastore Capacity Dashboard 755
ESXi Capacity Dashboard 756
VM Capacity Dashboard 757

Reclamation Dashboard 758
vSAN Capacity Dashboard 760
vSAN Stretched Clusters 760
vSAN ESA Capacity Dashboard 761
vSAN OSA Capacity Dashboard 762
Configuration Dashboards 762
Cluster Configuration Dashboard 764
ESXi Configuration Dashboard 766
Network Configuration Dashboard 768
VM Configuration Dashboard 769
vSAN Configuration Dashboard 770
Workload Management Configuration Dashboard 771
Consumer \ Correct it? Dashboard 772
Consumer \ Optimize it? Dashboard 773
Consumer \ Simplify it? 774
Consumer \ Update it? Dashboard 776
Provider \ Correct it? Dashboard 777
Provider \ Optimize it? Dashboard 778
Provider \ Simplify it? Dashboard 780
Provider \ Update it? Dashboard 781
Cost Dashboards 782
Consumer Layer 782
Provider Layer 784
Cost Optimization 787
Performance Dashboards 789
Guest OS Performance Profiling Dashboard 796
Network Top Talkers Dashboard 798
Storage Heavy Hitters Dashboard 799
VM Contention Dashboard 800
VM Performance Dashboard 802
VM Utilization Dashboard 804
Troubleshoot an Application Dashboard 805
Cluster Contention Dashboard 805
Cluster Performance Dashboard 809
Cluster Utilization Dashboard 810
VM Rightsizing Dashboard 812
Datastore Performance Dashboard 813
ESXi Contention Dashboard 814
ESXi Utilization Dashboard 815
Network Performance Dashboard 817
vSAN Contention Dashboard 818

vSAN File Services 820
vSAN Performance Dashboard 821
vSAN Utilization Dashboard 822
vSAN ESA Performance Dashboard 823
vSAN OSA Performance Dashboard 824
Sustainability Dashboards 825
Carbon Efficiency with Virtualization Dashboard 825
Carbon Transparency Dashboard 827
Environmental Impact of Idle VMs Dashboard 828
Green Supply Dashboard 829
Dashboard Library 830
Executive Summary Dashboards 830
Network Operation Center 832
Software Defined Wide Area Network Dashboard 837
Troubleshoot SD-WAN Dashboard 838
Troubleshoot SD-WAN Gateway Dashboard 838
Troubleshoot SD-WAN Orchestrator Dashboard 838
VMware Aria Automation Dashboards 839
Cloud Automation Environment Overview 839
Cloud Automation SDDC Project Price Overview 840
Cloud Automation SDDC Resource Consumption Overview 840
Cloud Automation Deployment Overview 841
Cloud Automation Top-N Dashboard 842
VMware Aria Operations Dashboards 842
VMware Aria Operations Cloud Billing 842
VMware Aria Operations Cloud Universal Billing 843
Service Discovery Dashboards 843
Service Distribution Dashboard 844
Service Relationships Dashboard 844
Service Visibility Dashboard 844
Virtual Machine Relationships Dashboard 845
Inventory Dashboards 845
vSphere Compute Inventory Dashboard 846
vSphere Network Inventory Dashboard 846
vSphere Storage Inventory Dashboard 846
Workload Management Inventory Dashboard 847
Microsoft Azure Dashboards 847
Azure Inventory Dashboard 848
Azure Virtual Machine Dashboard 849
Azure Kubernetes Service Overview Dashboard 850
Azure Virtual Network Overview Dashboard 851

Azure Storage Account Overview Dashboard 851
AWS Dashboards 852
AWS Instance Utilization Dashboard 855
AWS Auto Scaling Group Dashboard 855
AWS Troubleshooting Dashboard 855
AWS Instance Heatmap Dashboard 856
AWS Volume Performance Dashboard 856
AWS Disk Space Dashboard 856
Alerts 856
AWS Inventory Dashboard 856
Compute EC2 Dashboard 857
AWS Network Services Overview Dashboard 858
AWS Storage Services Overview Dashboard 859
Dashboards in VMware Cloud on AWS 860
VMC Capacity Dashboard 860
VMC Cost Overview Dashboard 861
VMC Inventory Dashboard 861
VMC Management VM Monitoring Dashboard 862
VMC Utilization and Performance Dashboard 862
VMC Configuration Maximums Dashboard 863
Google Cloud VMware Engine Dashboards 864
GCVE Configuration Maximums Dashboard 864
GCVE Cost Overview Dashboard 864
GCVE Inventory Dashboard 865
Dashboards in NSX Management Pack 865
NSX Configmax Metrics 866
NSX Inventory Dashboard 867
NSX System Dashboard 868
NSX Edge Dashboard 868
NSX Switch Dashboard 868
NSX Transport Node Performance Dashboard 869
NSX Edge Performance Dashboard 869
NSX Switch Performance Dashboard 869
NSX Load Balancer Performance Dashboard 870
Cloud Federation Adapter Dashboards 870
Getting Started - Cloud Federation Dashboard 871
Infrastructure Capacity 871
Infrastructure Configuration 873
SDDC Health and Configuration 875
Dashboards in VMware Infrastructure Health 876
Dashboards in Google Cloud Platform 877

GCP Inventory Dashboard 878
GCP Compute: CE Dashboard 879
GCP Kubernetes Service Overview Dashboard 880
GCP Network Overview Dashboard 881
GCP Storage Bucket Overview Dashboard 882

16 Configuring Widgets 884
Widget Interactions 884
Configuration Files 885
Configuration Files for Widget Metric Configuration 885
Configuration Files for Text Widget Content 887
Management Packs Configuration 889
Configuration Files for the Topology Widget 891
Widget Definitions List 893
Alert List Widget 895
Alert Volume Widget 902
Anomalies Widget 905
Anomaly Breakdown Widget 906
Container Details Widget 909
Capacity Remaining Widget 911
Container Overview Widget 912
Current Policy Widget 916
Data Collection Results Widget 918
DRS Cluster Settings Widget 920
Efficiency Widget 923
Environment Widget 925
Environment Overview Widget 927
Environment Status Widget 930
Faults Widget 934
Forensics Widget 936
Geo Widget 937
Heatmap Widget 940
Health Widget 945
Health Chart Widget 948
Log Analysis Widget 953
Mashup Chart Widget 955
Metric Chart Widget 957
Metric Picker Widget 965
Object List Widget 967
Object Relationship Widget 973
Object Relationship (Advanced) Widget 976

Property List Widget 979
Recommended Actions Widget 984
Risk Widget 987
Rolling View Chart Widget 990
Scoreboard Widget 995
Scoreboard Health Widget 1002
Sparkline Chart Widget 1005
Tag Picker Widget 1012
Text Display Widget 1014
Time Remaining Widget 1016
Top Alerts Widget 1018
Top-N Widget 1020
Topology Graph Widget 1026
View Widget 1031
Weather Map Widget 1034
Workload Widget 1037
Workload Pattern Widget 1039

17 Configuring Reports 1042
Create a Report Template 1042
Name and Description Details 1043
Views and Dashboards Details 1043
Formats Details 1044
Layout Options Details 1045
Manage Report Templates 1045
Accessing Report Templates 1046
Schedule a Report 1047
Generate and Regenerate a Report 1049
Accessing Generated Reports 1050
Download a Report 1051
Upload a Default Cover Page Image for Reports 1052

18 Configuring Views 1053
Accessing Predefined Views 1055
Views Overview 1055
Views and Reports Ownership 1057
Create and Configure a View 1057
List View 1058
Summary View 1064
Trend View 1070
Distribution View 1074

Text View 1079
Image View 1084
Editing, Cloning, and Deleting a View 1085
Share a View 1086
Including Deleted VMs in List View 1087
User Scenario: Create, Run, Export, and Import a VMware Aria Operations View for Tracking Virtual Machines 1088
Create a VMware Aria Operations View to view VM Memory Overhead 1089
Run a View 1089
Export a View 1090
Import a View 1090

19 Configuring Super Metrics 1092
Create a Super Metric 1093
Enhancing Your Super Metrics 1096
Exporting and Importing a Super Metric 1098
Super Metrics Tab 1098
Manage Super Metric Workspace 1100
Super Metric Functions and Operators 1100

20 Configuring Administration Settings 1105
Maintenance Schedules 1105
Manage Maintenance Schedules 1106
Managing Users and Access Control 1107
Roles and Privileges 1107
User Scenario: Manage User Access Control 1108
Access Control 1109
Audit Users and the Environment 1120
Certificates 1125
Importing CA Certificates 1126
Removing an Adapter Certificate 1127
Adding Physical Data Centers in VMware Aria Operations 1127
Modifying Global Settings 1128
Access Global Settings 1129
List of Global Settings 1129
Managing Content 1136
Creating a Backup 1137
Importing Content 1138
User Scenario: Migrating Content from VMware Aria Operations On-prem to VMware Aria Operations SaaS 1139
Best Practices for Migrating Content 1140
Managing Orphaned and Unassigned Content 1140

Dynamic Thresholds 1141
Customizing Icons 1142
Customize an Object Type Icon 1142
Customize an Adapter Type Icon 1144

# Configuring VMware Aria Operations

The *Configuring VMware Aria Operations* guide describes how to configure and monitor your environment. It shows you how to connect VMware Aria Operations to external data sources and analyze the data collected from them, ensure that users and their supporting infrastructure are in place, configure resources to determine the behavior of your objects, and format the content that appears in VMware Aria Operations.

## Intended Audience

This information is intended for VMware Aria Operations administrators, virtual infrastructure administrators, and operations engineers who install, configure, monitor, manage, and maintain the objects in your environment.

# Accessibility Compliance

1

VMware Aria Operations accessibility compliance provides several interactive elements that can be operated using a keyboard and screen reader.

## Keyboard Support

Table 1-1. Tooltips, Grid Sorting, Drag and Drop, and Combo-box with the X icon

| Component | Description | Examples |
| --- | --- | --- |
| Open and Close Tooltips | Navigate through elements using the **TAB** key.<br>Open tooltips using the **Ctrl + i** keys.<br>Close tooltips using the **ESC** key. | Navigate through elements in the workbench page and open and close tooltips.<br>1 In the menu, click **Home** > **Troubleshoot** and then click **Workbench.**<br>2 Click on a card.<br>If no card is available, search for a resource and click on it.<br>3 Use the **TAB** key to navigate through the elements.<br>4 Click **Ctrl + i** to open a tooltip and once done, click the **ESC** key to close it. |
| | | Navigate to the object relation chart and open and close tooltips.<br>1 In the menu, click **Environment** > **All Objects**.<br>2 Use the hierarchies in the left pane to locate the objects that you want and then click the **Metrics** tab.<br>3 Click **Show Object Relationship**.<br>4 Use the **TAB** key to navigate through the elements.<br>5 Click **Ctrl + i** to open a tooltip and once done, click the **ESC** key to close it. |

Table 1-1. Tooltips, Grid Sorting, Drag and Drop, and Combo-box with the X icon (continued)

| Component | Description | Examples |
| --- | --- | --- |
| Open Tooltips | Navigate through the alerts grid using the **TAB** key.<br>Open tooltips using the **Ctrl + i** keys. | Navigate through the inventory page and open tooltips.<br>1 In the menu, click **Administration** > **Inventory** and then click the **Objects** tab.<br>2 Use the **TAB** key to navigate to the **Relevance** column and then click **Ctrl+i** to open the tooltip. |
| | | Navigate through the alerts grid and open tooltips.<br>1 In the menu, click **Alerts** > **Triggered Alerts** and then click **All**.<br>2 Select an alert from the list to enable the Actions menu.<br>3 Use the **TAB** key to navigate to the **Importance** column and then click **Ctrl+i** to open the tooltip. |
| Grid Sorting | Sort columns that can be sorted using the **Enter** or **Space** keys. | Sort a grid.<br>1 Navigate to a column header.<br>2 Use the **Enter** or **Space** keys to sort the columns. |

Table 1-1. Tooltips, Grid Sorting, Drag and Drop, and Combo-box with the X icon (continued)

| Component | Description | Examples |
| --- | --- | --- |
| Drag and Drop | Drag and drop elements using the **TAB** and **Enter** keys.<br>**Note** If the default functionality of the **Enter** key has changed, you must use **Ctrl+Enter** instead. | Drag and drop alert symptoms.<br>1 In the menu, click **Alerts** > **Configuration** and then click **Alert Definitions**.<br>2 Click **Add** and enter the alert definition details and then click **Next**.<br>3 On the **Symptoms** tab, use the **TAB** key to navigate through the grid and click the **Enter** key on the first column to select one of the symptoms.<br>4 Use the **TAB** key again to navigate through the drop areas and then click the **Enter** key to drop the symptom.<br>5 Click **ESC** to cancel the action. |
| Combo-box with an X icon | Use the **X** icon or the **Delete** key to clear any combo-box in VMware Aria Operations. | Clear the combo-box for alerts.<br>1 In the menu, click **Alerts** > **Configuration** and then click **Alert Definitions**.<br>2 Click **Add** and enter the alert definition details and then click **Next**.<br>3 Click the **X** icon to clear it.<br>4 (Optional) Click the **Delete** key to clear it. |

# Collecting Data with Cloud Proxies in VMware Aria Operations

2

You can use cloud proxies in VMware Aria Operations to collect and monitor data from your on-premises data centers across different geo locations.

Cloud proxies provide high availability within your cloud environment, you can group two or more cloud proxies to form a collector group. The cloud proxy collector group ensures that there is no single point of failure in your cloud environment. If one of the cloud proxies experiences a network interruption or becomes unavailable, the other cloud proxy from the collector group takes charge and ensures that there is no downtime. All other user-initiated manual operations on the collector, such as to stop or restart the collector manually, do not result in automated rebalancing. For more information, see Managing Collector Groups.

**Note** When cloud proxies provide high availability within your cloud environment, the cluster can survive the loss of one data node without losing any data. However, the cloud proxy does not gurantee that the adapter instance will collect all the data during adapter instance failover (or when reassigning the adapter instance). VMware Aria Operations cloud proxy only provides additional application level data protection to ensure application level availability.

You can also use cloud proxies to rebalance the resources across the collectors in your collector group. The Rebalance option is available as part of the Edit menu in the Collector Groups page.

**Note** You can use the rebalance option before the vCenter Adapter initiates data collection. Once the data collection starts, the rebalance option is disabled.

Read the following topics next:

- Configuring Cloud Proxies in VMware Aria Operations
- Configuring Cloud Proxies in AWS and Azure
- Monitoring the Health of Cloud Proxies
- Using APIs with VMware Aria Operations Cloud Proxy
- Cloud Proxy FAQ
- Cloud Proxy Troubleshooting
- Configuring Collector Groups

# Configuring Cloud Proxies in VMware Aria Operations

Using cloud proxies in VMware Aria Operations, you can collect and monitor data from your on-prem data center or VMware cloud (VMC on AWS, VMC on Dell EMC, Azure VMware Solution, Google Cloud VMware Engine, and Oracle Cloud VMware Solution). You can deploy one or more cloud proxies in VMware Aria Operations to create a one-way communication between your end-point environment and VMware Aria Operations. Cloud proxies collect data from the end-point environment and uploads it to VMware Aria Operations. Cloud proxies can support multiple vCenter Servers or cloud accounts. For public cloud (AWS, Azure, GCP), the default collector can collect and monitor data. For more information on cloud proxies, see the topic called Cloud Proxy FAQ in the *Configuring VMware Aria Operations Guide*.

### Prerequisites

- Verify that you have an IP address, a DNS entry, and permissions to deploy OVF templates in vSphere.
- Log in to vSphere and verify that you are connected to a vCenter Server system.
- Allow outgoing HTTPS traffic for cloud proxy over port 443. For more information on firewall requirements in VMware Aria Operations see KB article 93210.
- Allow incoming traffic to cloud proxy over ports 443, 8443, 4505, and 4506 for telegraf based application monitoring.
- Allow incoming traffic to cloud proxy over port 443 for push model adapters or Suite-API on cloud proxy.
- Add a vCenter cloud account and provide an account with the following read and write privileges:
  - vCenter IP address or FQDN
  - Permissions required to install a cloud proxy on the vCenter Server.

  For more information on privileges, see, Privileges Required for Configuring a vCenter Adapter Instance.
- Using a firewall to restrict traffic by IP is not recommended since IPs can change without notice. Restricting traffic must be performed via FQDNs only.

### Procedure

1. Log in to VMware Aria Operations.
2. From the left menu, click **Data Sources > Cloud Proxies**, and then click **Add**.
3. Click the **Target Location** drop-down and select **VMware Cloud**.
4. Save the OVA path. Optionally, click **Download Cloud Proxy OVA** to download and save the OVA file locally.
   - To copy the link for the VMware Aria Operations Cloud Appliance, click the **Copy Path** icon for the Cloud Proxy OVA.

- To download and save the OVA file locally, click **Download Cloud Proxy OVA**.

5 Navigate to your vSphere, select the name of your vCenter Server cluster, and select **Deploy OVF Template** from the **Actions** menu.

6 Insert the ova link and then click **Next**.

- Paste the cloud proxy ova link in the **URL** field.
- Click the **Local File** option, browse, and select the downloaded OVA file.

7 Follow the prompts to install the OVA on your vCenter Server.

For the most current information about sizing and scaling, see Knowledge Base article 78491.

8 When prompted to enter the Unique Registration Key in the **Customize template** screen, return to the Install Cloud Proxy page in VMware Aria Operations.

9 (Optional) Activate **Data Persistence** to store data in the cloud proxy in case of connectivity issues. For more information, see Activating Data Persistence in Cloud Proxy.

**Note** The Unique Registration Key is refreshed and a new key is generated if you activate data persistence.

10 Click the **Copy Key** icon.

The Unique Registration Key expires 24 hours after generation. To avoid using an expired key, click **Regenerate Key** before proceeding. The unique registration key is used by the cloud proxy to authenticate to VMware Aria Operations.

11 Return to vSphere and paste the key in the **Unique Registration Key** text box to install the VMware Aria Operations Cloud Appliance.

12 Select **Prefer IPv6** to use IPv6 for internal communications.

13 (Optional) Set up a proxy server in the **Customize template** screen.

a Enter details in the **Network Proxy IP Address** and **Network Proxy Password** properties.

b To activate SSL, select the **Use SSL connection to proxy** check box.

c If you are using SSL, you can verify the certificate of the proxy server. Public certificate authorities are used to verify the proxy server certificate. To activate this, select the **Verify proxy's SSL cert** check box in the **Verify SSL cert** property.

d If you have a custom certificate authority, paste the root certificate authority in the **Custom CA** property to verify the certificate of the proxy server. The root certificate authority is passed on to the cloud proxy. Include the following lines when you copy the root certificate authority:

```
"-----BEGIN CERTIFICATE-----"
```

```
"-----END CERTIFICATE-----"
```

14 Click **Finish**.

The deployment takes a few minutes to finish.

15 Locate the cloud proxy you just installed, select the VMware Aria Operations Cloud Appliance, and click **Power on**.

**Note** You must power on the VMware Aria Operations Cloud Appliance within 24 hours of registering it. After 24 hours, the Unique Registration Key expires, and you must delete the VMware Aria Operations Cloud Appliance and deploy another cloud proxy.

16 Return to the Cloud Proxy page in VMware Aria Operations to view the status of the cloud proxy you just installed. For more information, see Monitoring the Health of Cloud Proxies.

17 To view the accounts that are using this connection, click the Cloud Proxy.

The communication from the cloud proxy to cloud is one way. The cloud proxy initiates this connection and if necessary, it also pulls data from cloud (like the adapters configuration or upgrade pak). The cloud proxy requires a regular Internet access over the https protocol but it does not need any special firewall configuration. The cloud proxy verifies the certificate of the cloud service it connects to and if there are transparent proxy servers which do stop SSL, it might cause connectivity problems for the cloud proxy.

The cloud proxy also supports connection through the corporate proxy server. The proxy settings are given during OVF deployment.

18 (Optional) To delete a cloud proxy, click the vertical ellipsis and then click **Delete**. For more information, see Deleting Cloud Proxies .

#### What to do next

Upgrade your cloud proxy. For more information, see the topic called Upgrading Cloud Proxy in the *Getting Started with VMware Aria Operations Guide*.

Troubleshoot any cloud proxy issues. For more information, see Cloud Proxy Troubleshooting.

## Activating Data Persistence in Cloud Proxy

Using data persistence, you can avoid data gaps in case of temporary connectivity issues. Activate data persistence to allow the cloud proxy to store data that is sent to VMware Aria Operations, if the connection fails between the cloud proxy and VMware Aria Operations. The cloud proxy stores all the persisted data that includes metrics and properties, along with time stamps.

Cloud proxy can store data for a maximum duration of one hour. If there is lack of space or if the connection fail lasts for more than an hour, the cloud proxy rotates the stored data by deleting the oldest stored data and replacing it with the most recently collected data. In case of lack of space, you can add additional storage, for more information, see KB article 2016022.

Once the connection is restored, the cloud proxy sends the stored data to VMware Aria Operations before the real-time data. The stored data is displayed before the real-time data as data is displayed in the correct time series. If the connection loss was more than one hour, there can also be a gap in the displayed data.

You can activate data persistence when you deploy a new cloud proxy. For more information, see Configuring Cloud Proxies in VMware Aria Operations . You can also activate data persistence in an existing cloud proxy. For more information, see Monitoring the Health of Cloud Proxies.

## Configuring Cloud Proxies in AWS and Azure

Using cloud proxies in VMware Aria Operations, you can collect and monitor data from your public cloud environments (Amazon Web Services and Microsoft Azure). You can deploy one or more cloud proxies on VMware Aria Operations to create a one-way communication between your end-point environment and VMware Aria Operations. The cloud proxies work as one-way remote collectors and upload data from the end-point environment to VMware Aria Operations. Cloud proxies can support multiple cloud accounts.

### Prerequisites

- Allow outgoing HTTPS traffic for cloud proxy using port 443.
- Allow outgoing traffic from the endpoints to cloud proxy using port 443.

### Procedure

1. Log in to VMware Aria Operations.
2. From the left menu, click **Data Sources > Cloud Proxy**, and then click **New**.
3. Click the **Target Location** drop-down and select one of the following.
   - **Amazon Web Services** (AWS).
     a. Click the **AWS image URL**.

        The AWS portal opens. Log in to your AWS account if prompted.

     b. Click **Launch Instance with AMI**. For more information on how to launch the instance from AMI, see the VMware KB article 88673.

     c. Click the **Regenerate Key** icon to avoid using an expired key and then click the **Copy Key** icon to copy the unique registration key for the User Data field in the Advanced section on the AWS EC2 instance creation page.

     d. Launch the instance. It will take a few minutes for VMware Aria Operations to detect the newly deployed cloud proxy after it is launched and powered up in AWS.
   - **Microsoft Azure**.
     a. Click the **Microsoft Azure URL**.

The Microsoft Azure portal opens. Log in to your Microsoft Azure account if prompted.

b Click **Create**. For more information on how to create a VM from Microsoft Azure, see the VMware KB article 88672.

c Click the **Regenerate Key** icon to avoid using an expired key and then click the **Copy Key** icon to copy the unique registration key for the User Data field in the Advanced section on the Microsoft Azure VM instance creation page.

d Launch the instance. It will take a few minutes for VMware Aria Operations to detect the newly deployed cloud proxy after it is launched and powered up in Microsoft Azure.

4 Return to the Cloud Proxy page in VMware Aria Operations to view the status of the cloud proxy you just installed. For more information, see Monitoring the Health of Cloud Proxies.

### What to do next

For AWS you must deploy one cloud proxy per account. To be able to collect metrics from different regions, you must ensure the connectivity between the AWS EC2 instances and cloud proxy.

For Microsoft Azure you must deploy one cloud proxy per account. To be able to collect metrics from different regions, you must ensure the connectivity between the Microsoft AzureVM instances and cloud proxy.

## Monitoring the Health of Cloud Proxies

You can view the status and health of your cloud proxy after you add it in VMware Aria Operations. You can then monitor the health and view alerts and metrics of your cloud proxy using the VMware Aria Operations Cloud Proxy object.

### Procedure

1 Log in to VMware Aria Operations.

2 From the left menu, click **Data Sources** > **Cloud Proxies**.

The list of cloud proxies is displayed.

| Option | Description |
|---|---|
| Name | The name of the cloud proxy. |
| IP | The IP address of the cloud proxy. |
| Status | Status of the cloud proxy. For example, the Going Online status is displayed for a few minutes when you add a new cloud proxy. Once the cloud proxy is connected to VMware Aria Operations, the status changes to Online. If the VMware Aria Operations is not connected, the Offline status is displayed. |

| Option | Description |
| --- | --- |
| Version | The version used to install the cloud proxy. |
| Accounts | The number of accounts that are created and associated with the cloud proxy. |
| Type | Displays the cloud proxy type. |
| Collector Group | Displays if the cloud proxy is part of the collector group. If the cloud proxy is part of the collector group, the name is displayed. Click the name to view the collector group details.<br>**Note** No data is displayed if the cloud proxy is a standalone cloud proxy and not a part of the collector group. |
| Network Proxy Address | The network proxy address of the cloud proxy. |
| Network Proxy Port | The network proxy port number of the cloud proxy. |
| Target | Displays the target location on which the cloud proxy is deployed. |
| Data Persistence | Displays the status of data persistence for the cloud proxy. Data persistence activates the cloud proxy to store data if the connection fails between the cloud proxy and VMware Aria Operations.<br>■ **Activated**. Cloud proxy will store data.<br>■ **Deactivated**. Cloud proxy will not store data.<br>To activate/deactivate data persistence for multiple cloud proxies, select the cloud proxies, click the horizontal ellipsis, and then select **Activate Data Persistence** or **Deactivate Data Persistence**.<br>To activate/deactivate data persistence for a single cloud proxy, select the cloud proxy, click the vertical ellipsis, and then select **Activate Data Persistence** or **Deactivate Data Persistence**.<br>**Note** When the connection is restored, the cloud proxy sends the stored data to VMware Aria Operations. The stored data is displayed before the real-time data is displayed. |

| Option | Description |
| --- | --- |
| Time Estimation | Displays the estimated time duration for which the cloud proxy persists data.<br>Cloud proxy can store data for a maximum duration of one hour. If there is a lack of space or if the connection fail lasts for more than an hour, the cloud proxy rotates the stored data by deleting the oldest stored data and replacing it with the most recently collected data.<br>**Note** The time estimate value is displayed only if data persistence is activated. |
| Filter | You can search the list of cloud proxies according to the following criteria:<br>■ Name<br>■ IP<br>■ Status<br>■ Version<br>■ Collector Group<br>■ Network Proxy Address<br>■ Network Proxy Port<br>■ Target<br>■ Data Persistence |

3 (Optional) Click the vertical ellipsis.

a Click **Activate/Deactivate Data Persistence** to activate or deactivate data persistence for the cloud proxy.

b Click **Remove** to delete the cloud proxy.

4 Click a **Cloud Proxy** name.

The **Cloud Proxy Details** open.

Table 2-1. Cloud Proxy Details

| Option | Description |
| --- | --- |
| Rename | Click the **Rename Cloud Proxy** icon to update the cloud proxy name. |
| Proxy ID | ID of the cloud proxy. |
| IP Address | IP address of the cloud proxy. |
| OVA Version | The OVA file version used to install the cloud proxy. |
| Status | Status of the cloud proxy. For example, the Getting Online status is displayed for a few minutes when you add a cloud proxy. Once the cloud proxy is connected to VMware Aria Operations, the status changes to Online. If the VMware Aria Operations is not connected, the Offline status is displayed. |

Table 2-1. Cloud Proxy Details (continued)

| Option | Description |
| --- | --- |
| Creation Date | Date of creation of the cloud proxy. |
| Last Heartbeat | Last time stamp when VMware Aria Operations ran a Health Check for this cloud proxy. When you click a cloud proxy to view its details, VMware Aria Operations sends a heartbeat to check if the cloud proxy is still reachable. |
| CPU | CPU usage. |
| Memory | Memory usage. |
| Number of Cloud Accounts | Displays the number of cloud accounts using the cloud proxy. Each cloud proxy might have one or more cloud accounts. You can also view the health and status of these cloud accounts from this details pane. |

5 If your cloud proxy is not collecting data, you can view the health of the cloud proxy. From the left menu, click **Environment > Inventory**, select the **VMware Aria Operations Cloud Proxy Object** from the list, and then click **Show Detail**.

For more details, see Inventory Tab and Inventory: List of Objects.

6 After you locate the VMware Aria Operations Cloud Proxy object, you can view the object details using the Summary tab. For more information, see Summary Tab.

7 Use the Alerts tab to monitor the health of the cloud proxy. If there are any issues, troubleshoot them using the Metrics tab.

If your cloud proxy is not working properly, an alert is displayed.

```
One or more VMware Aria Operations services on a cloud proxy are down
```

To clear this alert, perform the following steps:

- Check the network connectivity and configuration for the cloud proxy.
- Take the cloud proxy offline and then bring it online.

For more information, see Cloud Proxy FAQ and Cloud Proxy Troubleshooting.

**Note** It is recommended that you create a notification rule for this alert so that, quick remediation steps can be taken, if necessary.

8 (Optional) You can use the cloud proxy command line interface for other cloud proxy related actions. For more details, see Using the Cloud Proxy Command-Line Interface.

### What to do next

Click the **Collector Groups** tab to view the cloud proxies that are part of the collector groups. For more information on collector groups, see Managing Collector Groups.

## Deleting Cloud Proxies

You can delete a cloud proxy if there are issues with your cloud proxy. For example, you can delete the cloud proxy from the cloud proxy page in case of incorrect deployment of your cloud proxy, or if you need to deploy your cloud proxy in another cluster, or if you no longer require a cloud proxy.

### Procedure

1. Log in to VMware Aria Operations.
2. From the left menu, click **Data Sources** > **Cloud Proxy**.
3. Click the vertical ellipsis in front of the cloud proxy you want to delete and then click **Delete**.

   The confirmation wizard appears.
4. Select the **I understand the risk** check box and then click **OK**.
5. The cloud proxy is deleted.

## Upgrading Cloud Proxy

Cloud Proxies are upgraded to a compatible cluster version automatically after the cluster upgrade. Expect a downtime of one or two cycles, as cloud proxies do not collect any data during this period. Data collection resumes after the upgrade is complete. The upgrade process runs concurrently for up to ten cloud proxies. If you have more than ten cloud proxies, the upgrade process will continue sequentially and the next cloud proxy upgrade will start as soon as any one of the ten running upgrades completes or fails the upgrade process. In case the automatic upgrade fails, you can upgrade your cloud proxy manually using the CLI.

**Note** When your cloud proxies are upgraded, the cloud proxy internal certificates also get renewed. After each upgrade, the cloud proxies will have a new root-CA certificate with a five year validity period.

For more information on what data gets collected, see vSphere.

You can manually upgrade your cloud proxy Using the Cloud Proxy Command-Line Interface.

### Using the Cloud Proxy Command-Line Interface

You can use SSH to access the cloud proxy instance and use its Command-Line Interface to run the following actions:

- Manually upgrade your cloud proxy in case the automatic download of the latest binary fails. When the automatic download fails, you see a notification on the VMware Aria Operations user interface. To manually upgrade your cloud proxy instance to the latest version, see the following KB article 80590.
- Generate the support bundle.
- Gather the status of the cloud proxy's health and connectivity details.

| Command Line | Description |
| --- | --- |
| `cprc-cli -h, --help` | Displays the help message and use of the command-line interface. |
| `cprc-cli -s, --status` | Prints the cloud proxy life-cycle status, configuration details, upgrade-related information, and more. It is useful to catch necessary information related to support and troubleshooting, or check the connection to VMware Aria Operations, or check the product version number, and so on. |
| `cprc-cli -u PRODUCT_PAK, --upgrade PRODUCT_PAK` | The cloud proxy instance is enabled for an automated upgrade by default. But if the automated upgrade fails due to any exceptional issue, use this command line to upgrade your cloud proxy instance to the desired version. |
| ■ 8.3 Release `cprc-cli -sb, --generate-support-bundle`<br>■ 8.4 Release `cprc-cli -sb, --generate-support-bundle`<br>■ 8.5 Release `cprc-cli IS_HEAVY -sb, --generate-support-bundle IS_HEAVY`<br><br>The `IS_HEAVY` option should be specified as true or false. For example:<br><br>`cprc-cli -sb true`<br><br>`cprc-cli -sb false`<br><br>With the true option, the support bundle is generated with journal.ctl logs. With the false option, the support bundle is generated without journal.ctl logs | Generates the cloud proxy support bundle which is a package of logs, configurations, and status files. The support bundles are necessary for product support and troubleshooting. Generated support bundles can be found at the `/storage/core/vmware-vrops-cprc/support` location. |
| `cprc-cli -rsb SUPPORT_BUNDLE, --remove-support-bundle SUPPORT_BUNDLE` | Removes any specified support bundle. Although generated support bundle packages can be removed using the system-embedded commands, it is recommended to use this command for that action. |
| `cprc-cli -fm, --enable-fips-mode` | Enables FIPS mode for cloud proxy. |

| Command Line | Description |
| --- | --- |
| `cprc-cli -v -d all \| install \| network \| services \| connection` | Runs the troubleshoot process to check for all or any one of the following issues:<br>■ Installation. Checks if VM deployment and cloud proxy installation is successful.<br>■ Network. Checks IP configurations, DNS configurations, and Host (hostname) configurations.<br>■ Service. Checks whether the `vmware-casa.service`, `collector.service`, `haproxy.service`, `httpd-north.service`, and `httpd-south.service` are active and loaded based on the `systemctl show` command.<br>■ Connection: Checks the connection to VMware Aria Operations cluster nodes using `curl`. For the proxy server cases, checks that connection via the proxy server using its specified configuration.<br>Use the `-v` command to generate more verbose outputs. This command generates a log file that gets stored in `/var/log/support/cprc.analysis.log` location. |
| `-cs, --certificates-status` | Displays the status of the imported certificates and if the certificates have been configured correctly or not. |
| `cprc-cli -rc, --replace-certificate` | Allows you to replace the self signed certificates for cloud proxy.<br>Add the full path to the new certificate to the command. The function expects a certificate full chain containing the certificate and private key. For more information, see KB article 89583. |
| `cprc-cli -cc, --connectivity-check` | Allows you to run the connectivity check. You can check the HTTPS connectivity to a specified url. Incase the url is not specified, it checks the HTTPS connectivity to the default prerequisites<br>For SaaS configurations, the default prerequisites are the gateway FQDN and the upgrade pack url. |
| `cprc-cli -cc -r, --response` | Displays the response from the server during the connectivity check. |
| `cprc-cli -cc -rt, --routing` | Displays the routing information during the connectivity check. |
| `cprc-cli -env, --ovfenv` | Displays the content of ovfEnv.xml file. |
| `cprc-cli -rec, --recovery` | Runs the CPRC recovery script. |
| `cprc-cli -rh, --run-on-hosts` | Runs a specific script on hosts according to the configurations provided. |
| `cprc-cli -gt, --generate-template` | Generates the template.json files for -rh options configurations. |

# Using APIs with VMware Aria Operations Cloud Proxy

You can use a browser or an HTTP client program to send requests and receive responses.

## REST Client Programs

Any client application that can send HTTPS requests is an appropriate tool for developing REST applications with the VMware Aria Operations API. REST client plug-ins are available for most browsers and many IDEs. The following open-source programs are commonly used:

- cURL. http://curl.haxx.se
- Postman application. http://www.getpostman.com

In addition, VMware provides language-specific client bindings for the VMware Aria Operations API.

## About the Schema Reference

The VMware Aria Operations REST API documentation includes reference material for all elements, types, queries, and operations in the VMware Aria Operations API. It also includes the schema definition files.

Swagger based API documentation is available with the product, with the capability of making REST API calls right from the landing page.

To access the API documentation, use the URL of your VMware Aria Operations cloud proxy. For example, if the URL of your cloud proxy is `https://www.cloudproxy.vmware.com`, the API reference is available from: `https://cloudproxy.vmware.com/suite-api/doc/swagger-ui.html`.

**Note** You must use your cloud proxy IP or FDQN in the url to access the API documentation.

Language-specific client bindings are available from:

```
https://www.cloudproxy.vmware.com/vrops-cloud/suite-api/
```

## Authorize VMware Aria Operations API

To perform API calls, you must authorize your VMware Aria Operations cloud proxy in the Swagger UI using the access token. Enter the **CSP Token (apiKey)** value and click **Authorize**.

## About the VMware Aria Operations API Examples

All examples include HTTP requests and responses. These examples show the workflow and content associated with operations such as creating and querying for information about objects in your monitored environment.

Example request bodies are in JSON. Request headers required by the VMware Aria Operations API are included in example requests that are not fragments of a larger example.

Most example responses show only those elements and attributes that are relevant to the operation being discussed. Ellipses (...) indicate omitted content within response bodies.

# Cloud Proxy FAQ

This topic covers some frequently asked questions about VMware Aria Operations Proxy.

## Configuration

1 What are the prerequisites for setting up a cloud proxy account?

   Prerequisites are given in the topic, Configuring Cloud Proxies in VMware Aria Operations .

2 What does one-way connection mean?

   Outbound connections are initiated from cloud proxy to VMware Aria Operations, over `https/443`. Cloud proxy can also facilitate vCenter actions.

3 Which ports should be opened?

   The most up-to-date technical information about ports can be found on Ports and Protocol. For information on firewall requirements in VMware Aria Operations see KB article 93210.

4 Which ports should be allowed for incoming traffic to cloud proxy?

   Allow port 443 https protocol for push model adapters like application monitoring or Suite-API on cloud proxy. Allow ports 4505, 4506, and 8443 via TCP protocol for application monitoring. Allow the VRRP protocol for intercommunication between cloud proxies in a application monitoring high availability activated collector group.

5 How do I edit environment settings for cloud proxy?

   You can edit vApp options. For more information, see Edit OVF Details for a Virtual Machine.

6 How are certificates managed?

   Certificates are managed by cloud proxies. But for any additional proxy servers with SSL communication, you need to provide certificate(s).

7 What credential is used to login to cloud proxy?

   You can login as the “root” user. You are expected to set a new password on the first login to cloud proxy VM.

   SSH access is disabled by default, so the first login must be done via the vCenter console. You can run the following command to start SSH service:

   ```
   systemctl start sshd
   systemctl enable sshd
   ```

   To reset password, see the VMware KB Article, 2001476.

8 Where can I configure the local HTTP proxy for VMC on AWS?

Perform the following steps:

a Login to VMware Aria Operations and go to the Administration page.

b Go to Cloud Accounts.

c Select VMC on AWS.

d Click + next to credentials to add a credential.

e In proxy details, add details for the local HTTP proxy. (Do not add details for cloud proxy here).

For more details, see Configuring VMware Cloud on AWS in VMware Aria Operations.

9 Will I be notified if the connection between cloud proxy and VMware Aria Operations breaks down?

You can configure alerts/notifications on the *VMware Aria Operations cloud proxy* object. For more information, see Monitoring the Health of Cloud Proxies.

VMware Aria Operations automatically generates notifications for the following scenarios:

- Cloud proxy is not reachable.
- Cloud proxy is nearing sizing limits.

10 How do I change account for cloud proxy?

You can edit vApp options. For more information, see Edit OVF Details for a Virtual Machine.

11 How can I check the status of cloud proxy?

For more information, see Monitoring the Health of Cloud Proxies.

12 Should I use Remote Collector or cloud proxy for monitoring?

VMware recommends that you use cloud proxy to take advantage of the latest enhancements. Also, application monitoring, HA of collector groups, and data persistence are only supported through cloud proxy.

## Sizing

1 How should I size the cloud proxy?

For information on sizing, see the VMware KB article 78491

2 How would I know if cloud proxy is nearing sizing limit?

VMware Aria Operations customers will receive an email when cloud proxy is nearing sizing limit.

## Upgrade

1 How do I upgrade cloud proxy?

Cloud proxy is upgraded automatically. In case the upgrade fails, see the VMware KB article 80590.

## High Availability

1 Is high availability supported?

   Cloud proxy supports high availability. You can add multiple cloud proxies to a collector group. If the collecting cloud proxy fails or gets disconnected, collection can be picked up by another proxy in the group.

   **Note** Since the failover is initiated after a period of 10 minutes, few collection cycles are lost.

To troubleshoot cloud proxy issues, see Cloud Proxy Troubleshooting.

# Cloud Proxy Troubleshooting

Cloud proxy troubleshooting steps are provided to help you easily resolve issues that you may come across in VMware Aria Operations.

Before you proceed with troubleshooting, see the Cloud Proxy FAQ.

## Installation and/or First Boot Failure

```
To verify the issue, check if /var/log/firstboot contains a file named "Succeeded".
```

If not, the following problems could result in VMware Aria Operations installation and/or first boot failure:

1 Unique Registration Key used while deploying Cloud Proxy is invalid. To verify, check the cloud proxy console.

   Solution: Redeploy cloud proxy.

## Cloud Proxy VM is running, but the status is Offline in VMware Aria Operations.

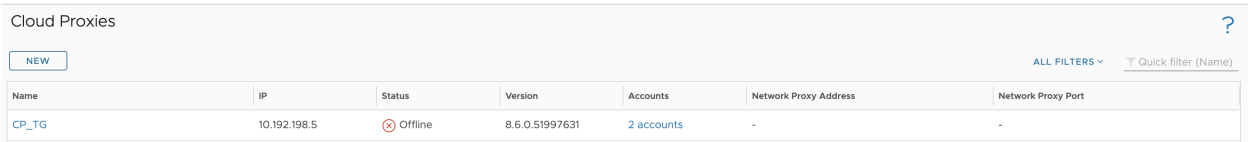


To verify the connection, use the following commands: (For the complete list of commands, please see Using the Cloud Proxy Command-Line Interface.

```
# Overall status of cloud proxy:cprc-cli -s

# Ping itself:
ip addr
ping <address>


# Ping gateway:
ip route
```

```
ping <gateway>

# Verify the connection outside the cloud proxy,
ping 8.8.8.8

Note: If you are using a network proxy,
use the /opt/vmware/share/vami/vami_config_net option#5 command
to ensure you have the correct configuration for the testings.
```

The following problems could result in VMware Aria Operations displaying the status of cloud proxy as offline.

1 Incorrect network proxy information in cloud proxy configuration.

   To verify the connection via a network proxy, use the following:

```
curl -vvv --proxy http(s)://proxy_user:proxy_pass@proxy_ip:proxy_port -H 'Accept:
application/json' -H 'Content-Type: application/json' -X GET https://<gateway url>/casa/
security/ping (gateway url example - 10238.gw.dev.vrops-ops.com)

To ignore SSL validation for a proxy server,
use curl --proxy-insecure. With SSL validation the customer can provide Proxy Server
certificate during cloud proxy deployment or re-configuration
so that provided certificate from customer can be used to check the connection with curl
with SSL certificate validation.
```

   Solution:

   a SSH to the Cloud Proxy VM and set the `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.

   b Shutdown the Cloud Proxy VM.

   c Update the network proxy configurations from the vCenter Server VM options using the vApp options Edit OVF Details for a Virtual Machine.

   d Boot the Cloud Proxy VM.

2 Incorrect Firewall settings.

   To verify access to outside URLs:

```
# For URLs to allowlist, please refer to the prerequisite section in Configuring Cloud
Proxies in VMware
                    Aria Operations
    .

# ping s3-us-west-2.amazonaws.com

# The Amazon VMware
                    Aria Operations instance can be found in this file: cat $VCOPS_BASE/
```

```
user/conf/collector/collector.properties

# Example row ttpRelayServerURL=https://10027.gw.staging.vrops-ops.com/suite-api
wget <ttpRelayServerURL>
```

Solution:

a SSH to the Cloud Proxy VM and set `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.

b Shutdown the Cloud Proxy VM.

c Allowlist URLs mentioned in the prerequisite section of Configuring Cloud Proxies in VMware Aria Operations .

d Boot the Cloud Proxy VM.

3 Required ports are not open.

To verify:

```
openssl s_client -showcerts -connect {address}:443

curl -v telnet://{address}:443

# Or, change the address to the machine you want to check:
python -c "import socket; print(socket.socket(socket.AF_INET,
socket.SOCK_STREAM).connect_ex(('127.0.0.1', 443)))"

# If you get a !=0 response, the server is not listening to the port.
```

Solution:

a SSH to the Cloud Proxy VM and set `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.

b Provide port access as mentioned in the prerequisite section of Configuring Cloud Proxies in VMware Aria Operations

c Boot the Cloud Proxy VM.

4 Invalid certificate.

To verify:

```
openssl s_client -showcerts -connect {address}:443
```

Solution:

a SSH to the Cloud Proxy VM and set `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.

b Follow the steps mentioned in VMware KB Article, 83698.

5 The logs folder `/storage/log` is running out of partition space.

Solution: Remove log files to ensure that enough space is available. Note that this is an exceptional case. In normal conditions, log files are auto archived.

6 One or more of the following services are down: `httpd-north.service`, `haproxy.service` and `collector.service`.

Solution:

- Check service status by running the following command: `systemctl status <service name>`.
- To start service, use the following command: `systemctl start <service name>`.

7 Unique Registration Key expired.

Solution: Redeploy Cloud Proxy with new Unique Registration Key.

## Cloud proxy is online, and state of Cloud Account is `Collecting`, but status is `Object Down`.

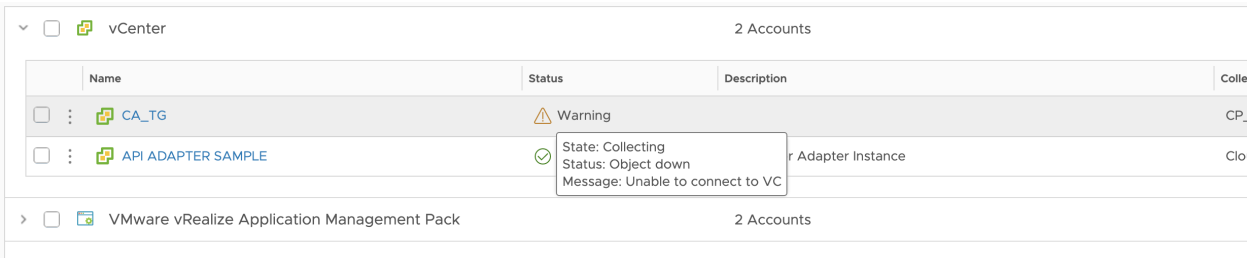


The following problem could result in VMware Aria Operations displaying the state of Cloud Account as `Collecting`, while the status is, `Object Down`.

1 Incorrect account credentials.

Solution: Check and update the credentials used while setting up the cloud account.

## Cloud proxy status is stuck in `Going Online`.

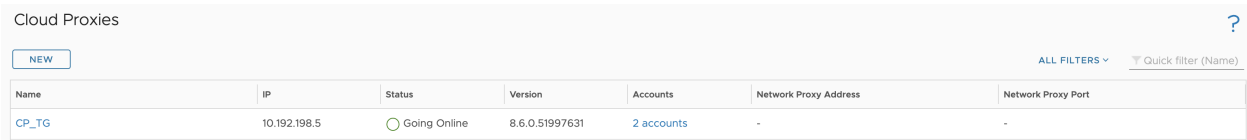


It can take up to 20 mins on first reboot, for the cloud proxy to be registered and come online. Wait for the specified time to see if cloud proxy comes online. If it still does not come online, one or more of the following services are down: `httpd-north.service`, `haproxy.service`, and `collector.service`.

Solution:

1 Check service status by running the following command: `systemctl status <service name>`

2 To start service, use the following command: `systemctl start <service name>`.

## Cloud proxy does not upgrade automatically, after the upgrade of VMware Aria Operations

There could be a few possible reasons why cloud proxy does not upgrade automatically after an upgrade of VMware Aria Operations.

1 High network latency leading to PAK download failure. Latency of >500ms is not supported.

   Solution: See the VMWare KB article 80590 on how to manually upgrade cloud proxy via CLI.

2 Upgrade status is stuck at `Running` since the previous upgrade had failed.

   Solution: Follow the steps given below to change the upgrade status.

   a Stop the casa service: `systemctl stop vmware-casa.service`.

   b Change the upgrade status from `RUNNING` to `NONE` in the following files:

```
./storage/db/vmware-vrops-cprc/status/cprc.upgrade.status
./storage/db/vmware-vrops-cprc/status/cprc.pak.status
```

   c See the VMware KB article 80590 and run the manual upgrade.

## Cloud proxy gets disconnected at regular intervals

There could be a few possible reasons why cloud proxy gets disconnected at regular intervals.

1 Check the network connectivity and latency.

2 Check if the cloud proxy VM can reach the DNS and use the NSlookup to validate the DNS connectivity.

# Configuring Collector Groups

VMware Aria Operations uses collectors like cloud proxies to manage cloud account processes such as gathering metrics from objects. You can select a collector or a collector group when configuring an cloud account.

If there are cloud proxies in your environment, you can create a collector group, and add cloud proxies to the group. When you assign an account to a collector group, the account can use any collector in the group. Use collector groups to achieve high availability in cases where the collector experiences network interruption or becomes unavailable. If this occurs, and the collector is part of a group, the total workload is redistributed among all the collectors in the group, reducing the workload on each collector. For more information, see Chapter 2 Collecting Data with Cloud Proxies in VMware Aria Operations .

You can add, edit, or remove collector groups in VMware Aria Operations, and rebalance your cloud accounts.

## Rebalancing a Cloud Account

Rebalancing of your cloud accounts is not intended to provide equally distributed cloud accounts across each collector in the collector group. The rebalancing action considers the number of resources that each cloud account collects to determine the rebalancing placement. The rebalancing happens at the cloud account, which can result in several small cloud accounts on a single collector, and a single huge cloud account on another collector, in your VMware Aria Operations instance.

Rebalancing your collector groups can add a significant load on the entire cluster. Moving cloud accounts from one collector to another collector requires that VMware Aria Operations stops the cloud account and all its resources on the source collector, then starts them on the target collector.

If a collector fails to respond or loses connectivity to the cluster, VMware Aria Operations starts automated rebalancing in the collector group. All other user-initiated manual operations on the collector, such as stopping or restarting the collector manually, do not result in automated rebalancing.

If one of the collectors fails to respond, or if it loses network connectivity, VMware Aria Operations performs automated rebalancing. In cases of automated rebalancing, to properly rebalance the collector group, you must have spare capacity on the collectors in the collector group.

## Where You Manage Collector Groups

From the left menu, click **Data Sources > Cloud Proxies**, and then click the **Collector Groups** tab.

### Table 2-2. Collector Group Summary Grid

| Options | Description |
| --- | --- |
| Collector Group toolbar | To manage collector groups, use the toolbar icons.<br>■ **Add**. Add a collector group<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>  ■ **Edit**. Modify the collector group by adding or removing cloud proxies. You can also edit a collector group to activate or deactivate application monitoring high availability.<br>  ■ **Delete**. Remove the selected collector group.<br>  ■ **Rebalance collector group**. Rebalance one collector group at a time. If you have permission to manage clusters, you can rebalance the workload across the collectors and the cloud proxies in the collector group. The rebalance action moves objects from one collector group to another to rebalance the number of objects on each collector in the collector group. If a disk rebalance is already in progress, the collector rebalance does not run<br>  ■ **Retry collector group configuration**: In case the application monitoring high availability activated collector group configuration fails, you can retry the configuration.<br>  ■ The actions are disabled when the application monitoring high availability activated collector group configuration is in progress. Once the cofiguration is complete you can perform the actions. |
| Name | The name given to the collector group when the collector group is created or updated. |
| Collector Group ID | The collector group ID created when the collector group is created. |
| Description | Description given to the collector group when the collector group is created or updated. |
| Application Monitoring HA | Displays the application monitoring high availability status of the collector group.<br>■ Activated. The collector group can be used for application monitoring.<br>■ Deactivated. The collector group cannot be used for application monitoring.<br>**Note** For default collector groups, high availability for application monitoring is deactivated by default. |
| HA Status | Displays the application monitoring high availability of the collector group.<br>■ OK. The collector group high availability configuration is successful.<br>■ In Progress. The collector group high availability configuration is in progress. This status is displayed when adding, deleting, updating, or deactivating the collector group.<br>■ Empty. If high availability is not configured, no status is displayed.<br>**Note** If the high availability status of the cloud proxy fails, the HA status of the failed cloud proxy is displayed. In case there are multiple cloud proxies in a collector group with failed HA status, the first failed message is displayed in the Collectors Group page. |

### Table 2-2. Collector Group Summary Grid (continued)

| Options | Description |
| --- | --- |
| Virtual IP | Virtual IP address of the collector group.<br>**Note** Pick the Virtual IP address that any of the cloud proxies in a collector group can own and receive traffic from. The Virtual IP address must be in the same subnet as the physical address of the cloud proxies. After you configure the Virtual IP address, ping it from a different network to ensure that it is routable. |
| Filters | Search the list of collector groups according to the following criteria:<br>■ Name<br>■ Description<br>■ Collector Name<br>■ Collector IP Address<br>■ Application Monitoring HA<br>■ HA Status<br>■ Group Virtual IP<br>■ Collector Group Id |

Click a collector group to view the collector group details.

### Table 2-3. Collector Group Details Grid

| Detail Grid Options | Description |
| --- | --- |
| Collector Group Name | Displays the name of the collector group. Click **Edit** to modify the collector group details. Click **View Description** to view the collector group description. |
| Application Monitoring High Availability Status | Displays the status of the application monitoring high availability. |
| Virtual IP | Displays the virtual IP added during the collector group creation. |
| Name | Name given to the cloud proxy when the cloud proxy was created. Click the name to view the cloud proxy details.<br>For more information on cloud proxies, see Monitoring the Health of Cloud Proxies. |
| IP Address | IP address of the cloud proxy |
| Status | Status of the cloud proxy: online or offline. |

Table 2-3. Collector Group Details Grid (continued)

| Detail Grid Options | Description |
| --- | --- |
| HA Status | Displays the cloud proxy high availability status.<br>■ OK. The cloud proxy application monitoring high availability configuration is successful.<br>■ Activation In Progress. Displayed when the cloud proxy is added to a application monitoring HA activated collector group or when the creation of a new application monitoring HA activated collector group is in progress.<br>■ Deactivation In Progress. Displayed when the deactivation of a application monitoring HA activated collector group is in progress.<br>■ Removal In Progress. Displayed when the removal of a cloud proxy from a application monitoring HA activated collector group is in progress.<br>■ Group Delete In Progress. Displayed when the deletion of the entire application monitoring HA activated collector group is in progress.<br>■ Empty. If the cloud proxy is not part of a collector group with high availability, no status is displayed.<br>■ If the cloud proxy high availability configuration fails, the failure reason is displayed. For example, "Keepalived is failed" or "Apache south is failed". |
| Type | Displays the cloud proxy type. |
| Number of Objects | Displays the number of objects collected by the cloud proxy. |
| Number of cloud accounts | Displays the number of cloud accounts using the cloud proxy. |

# Adding a Collector Group

Create a new collector group from the available cloud proxies in your VMware Aria Operations environment. A cloud proxy can only be added to only one collector group at a time.

## Where You Add New Collector Groups?

From the left menu, click **Data Sources > Cloud Proxies**, and then click the **Collector Groups** tab. Click **Add**.

Table 2-4. Add New Collector Group

| Option | Description |
| --- | --- |
| Name | Name of the collector group. |
| Description | Description of the collector group. |
| Application Monitoring High Availability | Activate this option to use high availability in application monitoring using collector groups. |

Table 2-4. Add New Collector Group (continued)

| Option | Description |
| --- | --- |
| Virtual IP | If you activate high availability for application monitoring, enter the virtual IP address.<br>**Note** Pick the Virtual IP address that any of the cloud proxies in a collector group can own and receive traffic from. The Virtual IP address must be in the same subnet as the physical address of the cloud proxies. After you configure the Virtual IP address, ping it from a different network to ensure that it is routable. |
| Collectors added to this collector group | Displays the collectors that are assigned to this collector group.<br>To assign collectors, double click a collector, or drag and drop it from the **Select Collectors** section to add it to the collector group.<br>**Note** You cannot assign a collector that is part of another collector group. To reassign a collector to another collector group, you must remove it from the existing collector group and then reassign it.<br>To remove a collector from the collector group, click the **Remove Collector** icon. |
| Select Collectors | Displays a list of cloud proxies in your VMware Aria Operations environment together with their collector name, IP address, type, collector group, and status.<br>Collectors that are assigned to this collector group appear with a check mark before the collector name. |
| Filters | You can search the list of collectors according to the following criteria:<br>■ Collector Name<br>■ IP address<br>■ Collector Group Name<br>■ Status |

# Editing Collector Groups

Edit a collector group in VMware Aria Operations by adding cloud proxies to the group, or removing the collectors that you no longer require to be part of the group. You can also edit a collector group to activate or deactivate application monitoring high availability.

## Where You Edit a Collector Group?

From the left menu, click **Data Sources > Cloud Proxies**, and then click the **Collector Groups** tab. Click the vertical ellipsis and then select **Edit**. The Edit Collector Group page opens. You can update the collector group details and then click **Save**.

## Table 2-5. Edit Collector Group Options

| Option | Description |
| --- | --- |
| Name | Name given to the collector group when the collector group is created. |
| Description | Description given to the collector group when the collector group is created. |
| Application Monitoring High Availability | Activate this option to use high availability in application monitoring using collector groups.<br>If high availability in application monitoring is activated, you can deactivate it. |
| Virtual IP | If you activate high availability for application monitoring, enter the virtual IP.<br>**Note** You cannot edit the virtual IP address for a collector group with high availability for application monitoring activated. |
| Collectors added to this collector group | Displays the collectors that are assigned to this collector group.<br>To assign collectors, double click a collector, or drag and drop it from the **Select Collector** section to add it to the collector group.<br>**Note** You cannot assign a collector that is part of another collector group. To reassign a collector to another collector group, you must remove it from the existing collector group and then reassign it.<br>To remove a collector from the collector group, click the **Remove Collector** icon. |
| Select Collectors | Displays a list of the available cloud proxies in your VMware Aria Operations environment together with their name, IP address, type, collector group, and status.<br>Collectors that are assigned to this collector group appear with a check mark before the collector name. |
| All Filters | Filter the list of collectors according to the following criteria:<br>■ Collector Name<br>■ Collector Group Name<br>■ IP Address<br>■ Status |

# Configuring Policies

3

To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

Read the following topics next:


- Policies
- Operational Policies
- Types of Policies
- Using the Policy Workspace to Create and Modify Operational Policies
- Assigning Policies


## Policies

A policy is a set of rules that you define for VMware Aria Operations to use to analyze and display information about the objects in your environment. You can create, modify, and administer policies to determine how VMware Aria Operations displays data in dashboards, views, and reports.

### How Policies Relate to Your Environment

VMware Aria Operations policies support the operational decisions established for your IT infrastructure and business units. With policies, you control what data VMware Aria Operations collects and reports on for specific objects in your environment. Each policy can inherit settings from other policies, and you can customize and override various analysis settings, alert definitions, and symptom definitions for specific object types, to support the service Level agreements and business priorities established for your environment.

When you manage policies, you must understand the operational priorities for your environment, and the tolerances for alerts and symptoms to meet the requirements for your business critical applications. Then, you can configure the policies so that you apply the correct policy and threshold settings for your production and test environments.

Policies define the settings that VMware Aria Operations applies to your objects when it collects data from your environment. VMware Aria Operations applies policies to newly discovered objects, such as the objects in an object group. For example, you have an existing VMware adapter instance, and you apply a specific policy to the group named World. When a user adds a new virtual machine to the vCenter Server instance, the VMware adapter reports the virtual machine object to VMware Aria Operations. The VMware adapter applies the same policy to that object, because it is a member of the World object group.

To implement capacity policy settings, you must understand the requirements and tolerances for your environment, such as CPU use. Then, you can configure your object groups and policies according to your environment.

- For a production environment policy, a good practice is to configure higher performance settings, and to account for peak use times.
- For a test environment policy, a good practice is to configure higher utilization settings.

VMware Aria Operations applies the policies in the priority order, as they appear in the priority column. When you establish the priority for your policies, VMware Aria Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of any active policy:

1 In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

**Note** The Reorder Policies option is activated only if there are more than one active policies.

2 In the Reorder Policies window, select the policy and drag it up or down to change the priority.

3 Click **ok** to save the changes made to the priority.

The priority for the Default Policy is always designated with the letter D, and the other active policies are prioritized with numbers 1, 2, and so on. Policy with priority 1 indicates the highest priority. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria Operations associates the highest ranking policy with that object.

Table 3-1. Configurable Policy Rule Elements

| Policy Rule Elements | Thresholds, Settings, Definitions |
| --- | --- |
| Workload | Configure symptom thresholds for Workload. |
| Time Remaining | Configure thresholds for the Time Remaining. |
| Capacity Remaining | Configure thresholds for the Capacity Remaining. |
| Maintenance Schedule | Sets a time to perform maintenance tasks. |
| Attributes | An attribute is a collectible data component. You can activate or deactivate metric, property, and super metric attributes for collection, and set attributes as key performance indicators (KPIs). A KPI is the designation of an attribute that indicates that the attribute is important in your own environment. |

Table 3-1. Configurable Policy Rule Elements (continued)

| Policy Rule Elements | Thresholds, Settings, Definitions |
| --- | --- |
| Alert Definitions | Activate or deactivate combinations of symptoms and recommendations to identify a condition that classifies as a problem. |
| Symptom Definitions | Activate or deactivate test conditions on properties, metrics, or events. |

## Privileges to Create, Modify, and Prioritize Policies

You must have privileges to access specific features in the VMware Aria Operations user interface. The roles associated with your user account determine the features you can access and the actions you can perform. To set the policy priority:

1 In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

**Note** The Reorder Policies option is activated only if there are more than one active policies.

2 In the Reorder Policies window, select the policy and drag it up or down to change the priority.

3 Click **ok** to save the changes made to the priority.

## How Upgrades Affect Your Policies

After you upgrade VMware Aria Operations from a previous version, you might find newly added or updated default settings of policies such as, new alerts and symptoms. Hence, you must analyze the settings and modify these settings to optimize them for your current environment. If you apply the policies used with a previous version of VMware Aria Operations, the manually modified policy settings remain unaltered.

## Policy Decisions and Objectives

Implementing policy decisions in VMware Aria Operations is typically the responsibility of the Infrastructure Administrator or the Virtual Infrastructure Administrator, but users who have privileges can also create and modify policies.

You must be aware of the policies established to analyze and monitor the resources in your IT infrastructure.

- If you are a Network Operations engineer, you must understand how policies affect the data that VMware Aria Operations reports on objects, and which policies assigned to objects report alerts and issues.
- If you are the person whose role is to recommend an initial setup for policies, you typically edit and configure the policies in VMware Aria Operations.
- If your primary role is to assess problems that occur in your environment, but you do not have the responsibility to change the policies, you must still understand how the policies applied to objects affect the data that appears in VMware Aria Operations. For example, you might need to know which policies apply to objects that are associated with particular alerts.

- If you are a typical application user who receives reports from VMware Aria Operations, you must have a high-level understanding of the operational policies so that you can understand the reported data values.

## Policies Library

The policies library displays the base settings, default policy, and other best practice policies that VMware Aria Operations includes. You can use the policies library to create your own policies. The policies library includes all the configurable settings for the policy elements, such as workload, capacity and time remaining, and so on.

### How the Policies Library Works

Use the options in policies library to create your own policy from an existing policy, or to override the settings from an existing policy so that you can apply the new settings to groups of objects. You can also import or export a policy and reorder the policies.

Select a policy to display its details in the right pane. The right pane displays a high-level overview of all the details and options for that policy where these details are categorized in tabs. Expand each category to view all the related details.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for metrics and properties, alerts and symptoms, capacity, compliance, workload automation, and groups and objects. In this workspace, you can also apply the policy to objects and object groups. To update the policy associated with an object or object group, the role assigned to your user account must have the Manage Association permission activated for policy management.

### Where You Manage the Policies Library

To manage the policies library, from the left menu, click **Configure** > **Policies** > **Policy Definition**. The policies library appears and lists the policies available to use for your environment.

### Table 3-2. Policy Library Tab Options

| Option | Description |
| --- | --- |
| Toolbar | Use the toolbar selections to take action in the policies library.<br>■ Add. Create a policy from an existing policy.<br>■ Edit. Customize the policy so that you can override settings for VMware Aria Operations to analyze and report data about the associated objects.<br>■ Delete. Remove a policy from the list.<br>■ Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to `D`, which gives that policy the highest priority.<br>■ Export. Downloads the policy.<br>■ Import. Allows you to import policies. To import:<br>  ■ Click the **Import** option from the horizontal ellipsis.<br>  ■ Click **Browse** and select the file to import.<br>  ■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>  ■ Click **Import** to import the policy, and click **Done**.<br>  **Note** To import or export a policy, the role assigned to your user account must have the Import or Export permissions activated for policy management.<br>■ Reorder Policies. Change the priority of the active policies. |
| Filters | Limits the list based on the text you type.<br>You can also filter by:<br>■ Name<br>■ Description<br>■ Modified By |

Table 3-2. Policy Library Tab Options (continued)

| Option | Description |
| --- | --- |
| Policies library data grid | VMware Aria Operations displays the high-level details for the policies.<br>■ Name. Name of the policy as it appears in the Add or Edit Policy workspace, and in areas where the policy applies to objects, such as in Custom Groups.<br>■ Status: Indicates whether the policy is active or inactive.<br>■ Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects.<br>■ Last Modified. Date and time that the policy was last modified. |
| Policies library > Right Pane | The right pane displays the name and description of the policy from which the settings are inherited, the policy priority, and the option to edit the policy. From the right pane, you can view the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.<br>■ Metrics and Properties: Displays all the attribute types included in the policy. Attribute type includes, metrics properties, and super metrics.<br>■ Alerts and Symptoms: Displays all the alert and symptom definitions included in the policy. The Alert Definitions tabs display an overview of the alert definition, criticality, symptom, and state. The Symptoms Definitions tab displays an overview of the symptom name, criticality, and the metric name.<br>■ Capacity: Displays an overview of all the thresholds of the objects included in the policy.<br>■ Compliance: Displays the compliance thresholds inherited from the base policy or set while creating the policy.<br>■ Workload Automation: Displays the details of the workload optimized in your environment per your definition.<br>■ Groups and Objects: Displays the object or object groups associated with the selected policy and the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it is shown here. |

# Operational Policies

Determine how to have VMware Aria Operations monitor your objects, and how to notify you about problems that occur with those objects.

VMware Aria Operations Administrators assign policies to objects or object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with objects or object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Activate and deactivate alerts.
- Control data collections by persisting or not persisting metrics on the objects in your environment.
- Configure the product analytics and thresholds.

- Monitor objects and applications at different service levels.
- Prioritize policies so that the most important rules override the defaults.
- Understand the rules that affect the analytics.
- Understand which policies apply to objects or object groups.

VMware Aria Operations includes a library of built-in active policies that are already defined for your use. VMware Aria Operations applies these policies in priority order.

When you apply a policy to an object or an object group, VMware Aria Operations collects data from the objects based on the thresholds, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that are activated in the policy.

The following examples of policies might exist for a typical IT environment.

- Maintenance: Optimized for ongoing monitoring, with no thresholds or alerts.
- Critical Production: Production environment ready, optimized for performance with sensitive alerting.
- Important Production: Production environment ready, optimized for performance with medium alerting.
- Batch Workloads: Optimized to process jobs.
- Test, Staging, and QA: Less critical settings, fewer alerts.
- Development: Less critical settings, no alerts.
- Low Priority: Ensures efficient use of resources.
- Default Policy: Default system settings.

# Types of Policies

There are three types of policies such as default policies, custom policies, and policies that are offered with VMware Aria Operations.

## Custom Policies

You can customize the default policy and base policies included with VMware Aria Operations for your own environment. You can then apply your custom policy to an individual object or groups of objects, such as the objects in a cluster, or virtual machines and hosts, or to a group that you create to include unique objects and specific criteria.

You must be familiar with the policies so that you can understand the data that appears in the user interface, because policies drive the results that appear in the VMware Aria Operations dashboards, views, and reports.

To determine how to customize operational policies and apply them to your environment, you must plan ahead. For example:

- Must you track CPU allocation? If you overallocate CPU, what percentage must you apply to your production and test objects?
- Will you overallocate memory or storage? If you use High Availability, what buffers must you use?
- How do you classify your logically defined workloads, such as production clusters, test or development clusters, and clusters used for batch workloads? Or, do you include all clusters in a single workload?
- How do you capture peak use times or spikes in system activity? In some cases, you might need to reduce alerts so that they are meaningful when you apply policies.

When you have privileges applied to your user account through the roles assigned, you can create and modify policies, and apply them to objects. For example:

- Create a policy from an existing base policy, inherit the base policy settings, then override specific settings to analyze and monitor your objects.
- Use policies to analyze and monitor vCenter Server objects and non-vCenter Server objects.
- Set custom thresholds for capacity settings on all object types to have VMware Aria Operations report on workload, and so on.
- Activate specific attributes for collection, including metrics, properties, and super metrics.
- Activate or deactivate alert definitions and symptom definitions in your custom policy settings.
- Apply the custom policy to an individual object or groups of objects.

When you use an existing policy to create a custom policy, you override the policy settings to meet your own needs. You set the allocation and demand, the overcommit ratios for CPU and memory, and the thresholds for capacity risk and buffers. To allocate and configure what your environment is actually using, you use the allocation model and the demand model together. Depending on the type of environment you monitor, such as a production environment versus a test or development environment, whether you over allocate at all and by how much depends on the workloads and environment to which the policy applies. You might be more conservative with the level of allocation in your test environment and less conservative in your production environment.

When you establish the priority for your policies, VMware Aria Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria Operations associates the highest ranking policy with that object.

Your policies are unique to your environment. Because policies direct VMware Aria Operations to monitor the objects in your environment, they are read-only and do not alter the state of your objects. For this reason, you can override the policy settings to fine-tune them until VMware Aria Operations displays the results that are meaningful and that affect for your environment. For example, you can adjust the capacity buffer settings in your policy, and then view the data that appears in the dashboards to see the effect of the policy settings.

## Default Policy in VMware Aria Operations

The default policy is a set of rules that applies to most of your objects.

The Default policy is marked with the letter D in the Priority column and can apply to any number of objects.

All the Default policies appear in the Default Policy group in the policies library, even if that policy is not associated with an object group. When an object group does not have a policy applied, VMware Aria Operations associates the Default policy with that group.

A policy can inherit the Default policy settings, and those settings can apply to various objects under several conditions.

The policy that is set to Default always takes the lowest priority. If you attempt to set two policies as the Default policy, the first policy that you set to Default is initially set to the lowest priority. When you set the second policy to Default, that policy then takes the lowest priority, and the earlier policy that you set to Default is set to the second lowest priority.

You can use the Default policy as the base policy to create your own custom policy. You modify the default policy settings to create a policy that meets your analysis and monitoring needs. When you start with the Default policy, your new policy inherits all the settings from the Default base policy. You can then customize your new policy and override these settings.

The data adapters and solutions installed in VMware Aria Operations provide a collective group of base settings that apply to all objects. In the policy navigation tree in the policies library, these settings are called Base Settings. The Default policy inherits all the base settings by default.

## Policies Provided with VMware Aria Operations

VMware Aria Operations includes sets of policies that you can use to monitor your environment, or as the starting point to create your own policies.

Verify that you are familiar with the policies provided with VMware Aria Operations so that you can use them in your own environment, and to include settings in new policies that you create.

### Where You Find the Policies Provided with VMware Aria Operations Policies

In the menu, click **Configure** > **Policies** > **Policy Definition**, to see the policies provided with VMware Aria Operations.

### Policies That VMware Aria Operations Includes

All policies exist under the Base Settings, because the data adapters and solutions installed in your VMware Aria Operations instance provide a collective group of base settings that apply to all objects. In the policies library, these settings are called Base Settings.

The Base Settings policy is the umbrella policy for all other policies, and appears at the top of the policy list in the policies library. All the other policies reside under the Base Settings, because the data adapters and solutions installed in your VMware Aria Operations instance provide a collective group of base settings that apply to all objects.

The configuration based policy set includes policies provided with VMware Aria Operations that you use for specific settings on objects to report on your objects. This set includes several types of policies:

- Efficiency alerts policies for infrastructure objects and virtual machines
- Health alerts policies for infrastructure objects
- Overcommit policies for CPU and Memory
- Risk alerts policies for infrastructure objects and virtual machines

The Default Policy includes a set of rules that applies to most of your objects.

## Using the Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the policy workspace to create local policies quickly, and update the settings in existing policies. Select a base policy to use as the source for your local policy settings, and modify the thresholds and settings used for analysis and collection of data from objects or object groups in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated objects or object groups.

#### Prerequisites

Verify that objects or object groups exist for VMware Aria Operations to analyze and collect data, and if they do not exist, create them. .

#### Procedure

1. From the left menu, click **Configure** > **Policies** > **Policy Definition**.
2. Click **Add** to add a policy or you can select a policy and click **Edit Policy** to edit an existing policy.

   You can add and edit policies and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.

3 In the Create Policies workspace, assign a name to the policy, and enter the description.

   Give the policy a meaningful name and description so that all users know the purpose of the policy.

4 From the **Inherit From** drop-down, select one or more policies to use as a baseline to define the settings for your new local policy.

   You can use any of the policies provided with VMware Aria Operations as a baseline source for your new policy settings.

5 Click **Create Policy**.

   The Create Policies workspace provides the options to customize your policy.

6 Click **Metrics and Properties**. In this workspace, select the metric, property, or super metric attributes to include in your policy.

   VMware Aria Operations collects data from the objects in your environment based on the metric, property, or super metric attributes that you include in the policy.

   a Click **Save** and return to the create policies workspace.

7 Click **Alerts and Symptoms**. In this workspace, select the alert definitions and symptom definitions, and activate or deactivate them as required for your policy.

   VMware Aria Operations identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

   a Click **Save** and return to the create policies workspace.

8 Click **Capacity**. In this workspace, select and override the situational settings such as committed projects to calculate capacity, time remaining, and other detailed settings.

   a Click **Save** and return to the create policies workspace.

9 Click **Compliance**. In this workspace, set the compliance threshold required for your policy.

   a Click **Save** and return to the create policies workspace.

10 Click **Workload Automation**. In this workspace, select the optimization settings required for your policy.

   Click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

   a Click **Save** and return to the create policies workspace.

11 Click **Groups and Objects**. In this workspace, select one or more groups and objects to which the policy applies.

   VMware Aria Operations monitors the objects according to the settings in the policy that is applied to the object or the object group, triggers alerts when thresholds are violated, and reports the results in the dashboards, views, and reports. If you do not assign a policy to one or more objects or object groups, VMware Aria Operations does not assign the settings in that policy to any objects, and the policy is not active. For an object or an object group that dos not have a policy assigned, VMware Aria Operations associates the object group with the Default Policy.

   Filter the object types, and modify the settings for those object types so that VMware Aria Operations collects and displays the data that you expect in the dashboards and views.

   a Click **Save** and return to the create policies workspace.

#### What to do next

After VMware Aria Operations analyzes and collects data from the objects in your environment, review the data in the dashboards and views. If the data is not what you expected, edit your local policy to customize and override the settings until the dashboards display the data that you need.

## Policy Workspace in VMware Aria Operations

The policy workspace allows you to quickly create and modify policies. To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more objects or object groups.

### How the Policy Workspace Works

Every policy includes a set of packages, and uses the defined problems, symptoms, metrics, and properties in those packages to apply to specific objects or object groups in your environment. You can view the details for the settings inherited from the base policy, and display specific settings for certain object types. You can override the settings of other policies, and include additional policy settings to apply to the object types.

Use the **Add** and **Edit** options to create policies and edit existing policies.

### Where You Create and Modify a Policy

To create and modify policies, from the left menu, click **Configure** > **Policies** > **Policy Definition**, and then click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit the policy. The policy workspace is where you select the base policies, and customize and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can apply the policy to objects or object groups.

To remove a policy from the list, select the policy, click the horizontal ellipse, and select **Delete**.

## Policy Workspace Options

The policy workspace includes a step-by-step workflow to create and edit a policy, and apply the policy to custom object groups.

### What to read next

- Getting Started Details

  When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

- Select the Inherited Policy Details

  You can use any of the policies provided with VMware Aria Operations as a baseline source for your policy settings when you create a policy.

- Capacity Details

  You can filter the object types, and modify the settings for those object types so that VMware Aria Operations applies these settings. The data that you expect then appears in the dashboards and views.

- Maintenance Schedule Details

  You can set a time to perform maintenance tasks for each policy.

- Compliance Details

  Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the compliance for the object types in your policy.

- Workload Automation Details

  You can set the workload automation options for your policy, so that VMware Aria Operations can optimize the workload in your environment as per your definition.

- Configuring vCenter Pricing Details

  You can add and assign new pricing cards to vCenter and Clusters in VMware Aria Operations. The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

- Metrics and Properties Details

  You can select the attribute type to include in your policy so that VMware Aria Operations can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You activate or deactivate each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

- Alert and Symptom Details

  You can activate or deactivate alert and symptom definitions to have VMware Aria Operations identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

- Groups and Objects details

  You can assign your local policy to one or more objects or groups of objects to have VMware Aria Operations analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

## Getting Started Details

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

### Where You Assign the Policy Name and Description

To add a name and description to a policy, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy. Select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. The name and description appear in the Create or Edit policy workspace.

Table 3-3. Name and Description Options in the Create or Edit Policy Workspace

| Option | Description |
| --- | --- |
| Name | Name of the policy as it appears in the Create or Edit Policy screens, and in areas where the policy applies to objects, such as Custom Groups. |
| Description | Meaningful description of the policy. For example, use the description to indicate which policy is inherited, and any specific information that users must understand the relationship of the policy to one or more groups of objects. |
| Inherit From | The base policy that is used as a starting point. All settings from the base policy will be inherited as default settings in your new policy. You can override these settings to customize the new policy.<br>Select a base policy to inherit the policy settings as a starting point for your new policy. |

## Select the Inherited Policy Details

You can use any of the policies provided with VMware Aria Operations as a baseline source for your policy settings when you create a policy.

In the policy content area, you can perform the following actions:

- View the packages and elements for the inherited policy and additional policies that you selected to override the settings.
- Compare the differences in settings highlighted between these policies.
- Display object types.

To create a policy, select a base policy to inherit your new custom policy inherits settings. To override some of the settings in the base policy according to the requirements for the service level agreement for your environment, you can select and apply a separate policy for a management pack solution. The override policy includes specific settings defined for the types of objects to override, either manually or that an adapter provides when it is integrated with VMware Aria Operations. The settings in the override policy overwrite the settings in the base policy that you selected.

When you select and apply a policy to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the policy settings cards.

Click each card to display the inherited policy configuration, and your policy, and displays a preview of the selected policy settings. When you select one of the policy cards, you can view the number of activated and deactivated alert definitions, symptom definitions, metrics and properties, and the number of activated and deactivated changes.

When you select the Groups and Objects card, you select the objects to view so that you can see which policy elements apply to the object type. For example, when you select the StorageArray object type, the workspace displays the local packages for the policy and the object group types with the number of policy elements in each group.

You can preview the policy settings for all object types, only the object types that have settings changed locally, or settings for new object types that you add to the list, such as Storage Array storage devices.

### Where You Select and Override Base Policies Settings

To select a base policy to use as a starting point for your own policy, and to select a policy to override one or more settings that your policy inherits from the base policy, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy. In the Create policies workspace, add a name and description for the policy and from the **Inherit From** drop-down, select the base policy. The policy configuration, objects, and preview appear in cards below this drop-down.

## Capacity Details

You can filter the object types, and modify the settings for those object types so that VMware Aria Operations applies these settings. The data that you expect then appears in the dashboards and views.

### How the Capacity Workspace Works

When you turn on and configure the Capacity settings for a policy, you can override the settings for the policy elements that VMware Aria Operations uses to trigger alerts and display data. These types of settings include symptom thresholds based on alerts, situational settings such as committed projects to calculate capacity and time remaining, and other detailed settings.

Policies focus on objects and object groups. When you configure policy settings for your local policy, you must consider the object type and the results that you expect to see in the dashboards and views. If you do not change these settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

### Where You Set the Policy Capacity Settings

To set the capacity settings for your policy, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy.

In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The capacity settings for host systems, virtual machines, and other object types that you select appears in the workspace.

You can also edit the capacity settings while working on the objects under the Environment Tab. In the **Capacity** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Capacity Setting**.

### Table 3-4. Capacity Settings in the Create or Edit Policy Workspace

| Option | Description |
|---|---|
| Time Remaining Calculations | You can set the risk level for the time that is remaining when the forecasted total need of a metric reaches usable capacity.<br>■ Conservative. Select this option for production and mission-critical workloads.<br>■ Aggressive. Select this option for non-critical workloads.<br>■ Peak focused. Select this option if you want to use only the upper range of the data. The projection will be based on the high utilization points.<br>■ Conservativeness Strength. You can tune the level of conservativeness from one to five, with level 1 being the least conservative and level 5 being the most conservative. By default, the level of conservativeness is set to three. The upper bound will vary based on the level of conservativeness that you choose. Modifying the level of conservativeness will make the projection bounds narrower or wider. Higher the level, the wider the bounds and more conservative the projections for Time Remaining, Capacity Remaining, and Recommended Size.<br>**Note**<br>■ Rare and momentary peaks may be considered outliers and may not impact the projections even at the most conservative risk level.<br>■ Conservativeness strength is supported by the Conservative risk level. The Aggressive risk level produces Time Remaining, Capacity Remaining, Recommended Size metrics always based on the mean of the projection.<br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |
| Business Hours Schedule | Configure business hours as per your time zone, for calculation of capacity analysis and projections. VMware Aria Operations considers the business hours for all objects using the current policy.<br>During non-business hours, VMs could be running other data center activities such as OS upgrades, virus scans, etc after working hours, and hence may not appear to be idle. When you mark business hours schedules, VMware Aria Operations can analyze after hours metrics for inventory, compliance, troubleshooting and other purposes. The reclamation and right sizing analysis and recommendations are based on the business hours and ignore spikes after business hours.<br>Since the business hours schedule are based on policies, different objects can have different business hours. The capacity charts will be based on business hours.<br>**Note** You can set business hours schedule for VMs and clusters only.<br>**Note** After you specify business hours, the capacity forecast for the object will be based on the business hours and not 24 hours.<br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |
| Filters | Select the object type by which you want to filter. You can filter by Object Types, Local Changes, and Unsaved Changes. |

Table 3-4. Capacity Settings in the Create or Edit Policy Workspace (continued)

| Option | Description |
| --- | --- |
| Capacity Settings | Select an object to view the policy elements and settings for the object type so that you can have VMware Aria Operations analyze the object type.<br>You can view and modify the settings for the following policy elements:<br>■ Allocation Model<br>■ Custom Profile<br>■ Capacity Buffer<br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |
| Criticality Thersholds | You can view and modify the threshold settings for the following policy elements:<br>■ Time Remaining<br>■ Capacity Remaining<br>■ Workload<br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |

Click **Save** to save the changes.

The local changes made will appear under **Policies > Default Policy > Capacity** section. You can also view the preview of changes in the Capacity card.

### Policy Workload Element

Workload is a measurement of the demand for resources on an object. You can turn on and configure the settings for the Workload element for the object types in your policy.

#### How the Workload Element Works

The Workload element determines how VMware Aria Operations reports on the resources that the selected object group uses. The resources available to the object group depend on the amount of configured and usable resources.

- A specific amount of physical memory is a configured resource for a host system, and a specific number of CPUs is a configured resource for a virtual machine.
- The usable resource for an object or an object group is a subset of, or equal to, the configured amount.
- The configured and usable amount of a resource can vary depending on the type of resource and the amount of virtualization overhead required, such as the memory that an ESX host machine requires to run the host system. When accounting for overhead, the resources required for overhead are not considered to be usable, because of the reservations required for virtual machines or for the high availability buffer.

### Where You Override the Policy Workload Element

To view and override the policy workload capacity setting, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The workload settings for the object type that you have selected appear in the workspace.

View the Workload policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 3-5. Policy Workload Element Settings in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Workload | Allows you to set the number of collection cycles it takes to trigger or clear an alert. |

## Policy Time Remaining Element

The Time remaining element is a measure of the amount of time left before your objects run out of capacity.

### How the Time Remaining Element Works

The Time Remaining element determines how VMware Aria Operations reports on the available time until capacity runs out for a specific object type group.

- The time remaining indicates the amount of time that remains before the object group consumes the capacity available. VMware Aria Operations calculates the time remaining as the number of days remaining until all the capacity is consumed.
- To keep the Time Remaining more than the critical threshold setting or to keep it green, your objects must have more days of capacity available.

### Where You Override the Policy Time Remaining Element

To view and override the policy Time Remaining capacity setting, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The time remaining settings for the object type that you have selected appear in the workspace.

View the Time Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 3-6. Policy Time Remaining Element Settings in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Time Remaining | Allows you to set the number of days until capacity is projected to run out based on your current consumption trend. |

## Policy Capacity Remaining Element

Capacity is a measurement of the amount of memory, CPU, and disk space for an object. You can turn on and configure the settings for the Capacity Remaining element for the object types in your policy.

### How the Capacity Remaining Element Works

The Capacity Remaining element determines how reports on the available capacity until resources run out for a specific object type group.

- The capacity remaining indicates the capability of your environment to accommodate workload.
- Usable capacity is a measurement of the percentage of capacity available, minus the capacity affected when you use high availability.

### Where You Override the Policy Capacity Remaining Element

To view and override the policy Capacity Remaining analysis setting, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The capacity remaining settings for the object type that you have selected appears in the workspace.

View the Capacity Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

Table 3-7. Policy Capacity Remaining Element Settings in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Capacity Remaining | Allows you to set the percentage at which the capacity remaining alerts must be triggered. |

## Policy Allocation Model Element

Allocation model defines how much CPU, memory, or disk space is allocated to objects in a datastore, cluster or datastore cluster. In the policy, you can turn on the Allocation Model element and configure the resource allocation for the objects.

### How the Allocation Model Element Works

The Allocation Model element determines how calculates capacity when you allocate a specific amount of CPU, memory, and disk space resource to datastores, clusters or datastore clusters. You can specify the allocation ratio for either one, or all of the resource containers of the cluster. Unlike the demand model, the allocation model is used for capacity calculations only when you turn it on in the policy.

The allocation model element also affects the reclaimable resources for memory and storage in Reclaim page. When you turn on the Allocation Model element in the policy, the tabular representation of the VMs and snapshots in the selected data center from which resources can be reclaimed displays reclaimable memory and disk space based on the overcommit values.

### Where You Override the Allocation Model Element

To view and override the policy workload analysis setting, from the left menu, click **Configure > Policies > Policy Definition**.

Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card.

The allocation model settings for the object type that you selected appear in the workspace.

Click the unlock icon next to Allocation Model to set the overcommit ratios.

| Option | Description |
| --- | --- |
| Set overcommit ratio, to enable Allocation Model | Allows you to set the overcommit ratio for CPU, memory, or disk space. Select the check box next to the resource container you want to edit and change the overcommit ratio value. |

### Policy Custom Profile Element

The custom profile element lets you apply a custom profile which shows how many more of a specified object can fit in your environment depending on the available capacity and object configuration.

### Where You Define the Custom Profiles

To define a custom profile, from the left menu click **Configure > Custom Profiles**, and then click **Add** to define a new custom profile.

### Where You Select the Custom Profile Element

To view and override the policy Custom Profile analysis setting, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The custom profile element for the object types such as datastores, clusters and datastore clusters that you selected appear in the workspace. Click the lock icon to unlock the section and make changes.

### Policy Capacity Buffer Element

The capacity buffer element lets you add buffer for capacity and cost calculation. For vCenter Server objects, you can add buffer to CPU, Memory, and Disk Space for the Demand and Allocation models. You can add capacity buffer to datastores, clusters and datastore clusters. The values that you define here affect the cluster cost calculation. The time remaining, capacity remaining, and recommended values are calculated based on the buffer. For WLP, capacity buffer is first considered and then the headroom that you have defined is considered.

#### Where You Define the Capacity Buffer

To view and override the policy Capacity Buffer analysis setting, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The Capacity Buffer for the object type that you selected appears in the workspace. Click the lock icon to unlock the section and make changes.

#### How the Capacity Buffer Element Works

The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required. The value of the usable capacity reduces by the buffer amount that you specify here. The default buffer value is zero. If you are upgrading from a previous version of VMware Aria Operations, the buffer values are carried forward to the new version.

The capacity buffer value that you specify for the Allocation model is considered only if you have activated allocation model in the policy.

Starting from version 8.6, capacity buffer is depreciated from cluster compute resources. The overcommit ratio setting (from the allocation model) and buffer settings, if set for the datastore object, takes precedence for the disk space related to datastore cluster and cluster objects. If these settings are not set, then, from a cost calculation perspective, the settings of datastore cluster and cluster (if the settings are missing for the datastore cluster as well), are used. The allocation and buffer settings made on the cluster does not impact the underlying datastores (as they do not inherit these settings), and the same works vice-versa, settings made for datastores are not propagated to the cluster.

The following tables display the capacity buffer that you can define based on the vCenter Adapter object types:

| Object Type | Valid Models for Capacity Buffer |
|---|---|
| CPU | Demand<br>Allocation |
| Memory | Demand<br>Allocation |
| Disk Space | Demand<br>Allocation |

## Maintenance Schedule Details

You can set a time to perform maintenance tasks for each policy.

### Where You Override the Policy Maintenance Schedule Element

To view and override the policy Maintenance Schedule analysis setting, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Maintenance Schedule**.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 3-8. Policy Maintenance Schedule Element Settings in the Create or Edit Policies Workspace**

| Option | Description |
| --- | --- |
| Select Object Type | Select the object type by which you want to filter. |
| Filters | You can filter by Local Changes and Unsaved Changes. Select **Yes** or **No** from the drop-down and click **Apply** to apply the filters. |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Maintenance Schedule | Sets a time to perform maintenance tasks. During maintenance, VMware Aria Operations does not calculate analytics. |

## Compliance Details

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the compliance for the object types in your policy.

### Where You Override the Policy Compliance

To view and override the policy compliance setting, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Compliance**

View the compliance thresholds and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 3-9. Compliance Settings in the Create or Edit Policies Workspace**

| Option | Description |
| --- | --- |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Compliance | Allows you to set the compliance score threshold based on the number of violations against those standards. |

## Workload Automation Details

You can set the workload automation options for your policy, so that VMware Aria Operations can optimize the workload in your environment as per your definition.

### How the Workload Automation Workspace Works

You click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

### Where You Set the Policy Workload Automation

Access this screen through the Policies pages:

1 Select **Configure** from the left menu, then select **Policies**.

2 Click the **Policy Definition** card.

3 Select a policy that you want to modify. Ideally, this should be an active policy. Or, click the **ADD** button to add a new policy.

4 Select the **Workload Automation** card to review the changes, or click **EDIT POLICY** to make changes.

Table 3-10. Workload Automation in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Workload Optimization | Select a goal for workload optimization.<br>Select **Balance** when workload performance is your first goal. This approach proactively moves workloads so that the resource utilization is balanced, leading to maximum headroom for all resources.<br>Select **Moderate** when you want to minimize the workload contention.<br>Select **Consolidate** to proactively minimize the number of clusters used by workloads. You might be able to repurpose resources that are freed up. This approach is good for cost optimization, while making sure that performance goals are met. This approach might reduce licensing and power costs. |
| Cluster Headroom | Headroom establishes a required capacity buffer, for example, 20 percent. It provides you with an extra level of control and ensures that you have extra space for growth inside the cluster when required. Defining a large headroom setting limits the systems opportunities for optimization.<br>**Note** vSphere HA overhead is already included in useable capacity and this setting does not impact the HA overhead. |

Table 3-10. Workload Automation in the Create or Edit Policies Workspace (continued)

| Option | Description |
| --- | --- |
| Change Datastore | Click the lock icon to select one of the following options:<br>■ Do not allow Storage vMotion.<br>■ Allow Storage vMotion. This is selected by default.<br>Using this option, you can select what type of virtual machines VMware Aria Operations moves first to address workload. |
| Target Network Policy Setting for WLP | Click the lock icon to select the following option:<br>■ Generate a Target Network mapping<br>When you select this checkbox, the Workload Placement algorithm in VMware Aria Operations will automatically choose compatible target network, while making the decision to move the VM for the optimization. For choosing compatible network WLP engine will consider the segment path and logical switch UUID of the Distributed Port Group.<br>Workload Optimization Across networks is supported when the optimization candidate clusters are assigned with different port groups (configured with NSX). These port groups configured via NSX have same segmentID and Logical Switch UUID. To enable this ability, check the respective setting in VMware Aria Operations Workload Automation policy settings.<br>**Note** Segment ID and logical switch UUID properties are published on the VC port groups by NSX. So Worload Placement cannot provide a target network if it is not a NSX configuration and those properties are missing.<br>This setting is not selected by default. |

## Configuring vCenter Pricing Details

You can add and assign new pricing cards to vCenter and Clusters in VMware Aria Operations. The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

If you want to copy the vCenter pricing settings from the policy currently being edited to another policy, click **Copy local changes to other policy** and select the policy to which you want to copy the settings. The copied pricing configuration will override any existing local pricing configuration, in the target policy.

### Procedure

1. From the left menu, click **Configure > Policies > Policy Definition**, and then select the required policy or click **Add** to add a new policy.
2. Select the required policy, and in the right pane, click **Edit Policy**.
3. In the <policy name> [Edit] workspace, click the **VC Pricing** card.
4. Click the Lock icon to override parent policy settings.
5. Select if you want to activate or deactivate the pricing engine.

6 Configure **Basic Charges**: Click the Lock icon to edit the parent policy settings. Pricing can be performed either on a cost basis or independent of it by specifying rate cards. The factor entered here is multiplied by the cost calculated as a derivative of cost drivers.

   a Based on Cost/Based on Rate: Select if you want to pricing card to be cost-based or rate-based.

      The following options appear if you select the **Based on Cost** option:

      - CPU Cost: Enter a valid CPU cost factor.
      - Memory Cost: Enter a valid memory cost factor.
      - Storage Cost: Enter a valid storage cost factor.
      - Additional Cost: Enter a valid additional cost factor.

      The following options appear if you select the **Based on Rate** option:

      - CPU Rate: Enter the CPU Rate per vCPU, the charging period, and how to charge for the resources.
      - Memory Rate: Enter the memory rate per GB, the charging period, and how to charge for the resources.
      - Storage Rate: Enter the storage rate per GB, the charging period, and how to charge for the resources.

7 Configure **Guest OS Rate**: Click **Guest OS Rate** in the left pane and then click the Lock icon to edit the parent policy settings. These are additional charges that have to be included based on the operating system running on the virtual machine. The name of the operating system should match exactly as discovered by VMware Tools.

   a Click **Create Guest OS Rate** and enter the following details:

      - Guest OS Name: Enter a guest OS name.
      - Charge Period: The Charge Period indicates the frequency of charging.
      - Base Rate: Enter a base rate.

      The guest OS rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

8 **Tags**: Click **Tags** in the left pane and then click the Lock icon to edit the parent policy settings. Tag-based charges can be used to charge for value-added services such as antivirus database disaster recovery and other applications. These applications are to be represented as vCenter tags on the VMs for these charges to work.

   a Recurring Charges: Recurring charges represent repeating charges such as monthly license fees for antivirus software. Click **Add Recurring Tag** and enter the following details:

      - Tag Category: Enter a tag key.
      - Tag Value: Enter a tag value.

- Base Rate: Enter a base rate.
- Charge Period: The Charge Period indicates the frequency of charging.
- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.

The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

b One Time Tag: Tag-based one-time charges can be used to represent incidental charges such as charges for addressing a support ticket or charges for applying an operating systems patch. Click **Add One Time Tag** and enter the following details:

- Tag Category: Enter a tag key.
- Tag Value: Enter a tag value.
- Base Rate: Enter a base rate.

The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

c Rate Factor Tag: Rate factors are multiplication factors applied to already calculated charges. For example, to add a 50% premium on storage, set a rate factor of 1.5 to storage charge. Click **Rate Factor Tag** and enter the following details:

- Tag Category: Enter a tag key.
- Tag Value: Enter a tag value.
- Charge Applies To: Select what the charge applies to.
- Rate Factor: Enter a valid number. For example, if you want to increase the price of CPU which has a tag 'Tag1-Value1' by 20% then select **`CPU Charge`** from the **Charge Applies To** drop-down list and enter **`1.2`** in **Rate Factor**.

The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

9 Configure **Overall Charges**: Click **Overall Charges** in the left pane and then click the Lock icon to edit the parent policy settings. Overall Charges are flat charges that are applied to VMs that match this policy.

a VM Setup Charges: Enter a valid setup fee. This is to charge for the setup of the VMs.

b Recurring Charges: Enter a valid number.

c Charge Period: The Charge Period indicates the frequency of charging.

### What to do next

You can assign policies to the required Organization/Organization VDC under **Configure > Policies > Policy Assignment**. For details, see Assigning Policies.

## Metrics and Properties Details

You can select the attribute type to include in your policy so that VMware Aria Operations can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You activate or deactivate each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

### How the Collect Metrics and Properties Workspace Works

When you create or customize a policy, you can override the base policy settings to have VMware Aria Operations collect the data that you intend to use to generate alerts, and report the results in the dashboards.

To define the metric and super metric symptoms, metric event symptoms, and property symptoms, from the left menu click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**.

### Where You Override the Policy Attributes

To override the attributes and properties settings for your policy, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

You can also edit the metrics and properties while working on the objects under the Environment Tab. In the **Metrics** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Metrics Collection**.

Table 3-11. Metrics and Properties Options

| Option | Description |
| --- | --- |
| Actions | Select one or more attributes and select activate, deactivate, or inherit to change the state and KPI for this policy. |
| Filter options | Deselect the options in the **Attribute Type**, **State**, **KPI**, and **DT** drop-down menus, to narrow the list of attributes.<br>■ ✔ Activated. Indicates that an attribute will be calculated.<br>■ ✔ Activated (Force). Indicates state change due to a dependency.<br>■ ⊘ Deactivated. Indicates that an attribute will not be calculated.<br>■ ✔ Inherited. Indicates that the state of this attribute is inherited from the base policy and will be calculated.<br>■ ⊘ Inherited. Indicates that the state of this attribute is inherited from the base policy and will not be calculated.<br>The KPI determines whether the metric, property, or super metric attribute is considered to be a key performance indicator (KPI) when VMware Aria Operations reports the collected data in the dashboards. Filter the KPI states to display attributes with KPI activated, deactivated, or inherited for the policy. |
| Object Type | Filters the attributes list by object type. |

Table 3-11. Metrics and Properties Options (continued)

| Option | Description |
| --- | --- |
| Page Size | The number of attributes to list per page. |
| Attributes data grid | Display the attributes for a specific object type.<br>■ Name. Identifies the name of the metric or property for the selected object type.<br>■ Type. Distinguishes the type of attribute to be either a metric, property, or super metric.<br>■ Adapter Type. Identifies the adapter used based on the object type selected, such as Storage Devices.<br>■ Object Type. Identifies the type of object in your environment, such as StorageArray.<br>■ State. Indicates whether the metric, property, or super metric is inherited from the base policy.<br>■ KPI. Indicates whether the key performance indicator is inherited from the base policy. If a violation against a KPI occurs, VMware Aria Operations generates an alert.<br>■ DT. Indicates whether the dynamic threshold (DT) is inherited from the base policy. |

## Alert and Symptom Details

You can activate or deactivate alert and symptom definitions to have VMware Aria Operations identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

### How the Alert and Symptom Definitions Workspace Works

VMware Aria Operations collects data for objects and compares the collected data to the alert definitions and symptom definitions defined for that object type. Alert definitions include associated symptom definitions, which identify conditions on attributes, properties, metrics, and events.

You can configure your local policy to inherit alert definitions from the base policies that you select, or you can override the alert definitions and symptom definitions for your local policy.

Before you add or override the alert definitions and symptom definitions for a policy, familiarize yourself on the available alerts and symptoms.

- To view the available alert definitions, from the left menu, click **Configure > Alerts**, and then in the right pane click **Alert Definitions**.
- To view the available symptom definitions, from the left menu, click **Configure > Alerts**, and then in the right pane click **Symptom Definitions**. Symptom definitions are available for metrics, properties, messages, faults, smart early warnings, and external events.

A summary of the number of problem and symptoms that are activated and deactivated, and the difference in changes of the problem and symptoms as compared to the base policy, appear in the Analysis Settings pane of the policies workspace.

### Where You Override the Alert Definitions and Symptom Definitions

To override the alert definitions and symptom definitions for your policy, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The definitions appear in the workspace.

You can also edit the alert settings while working on the objects under the Environment Tab. In the **Alerts** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Alerts State**.

### Policy Alert Definitions and Symptom Definitions

You can override the alert definitions and symptom definitions for each policy.

### What to read next

- Policy Alert Definitions

  Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can activate or deactivate the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

- Policy Symptom Definitions

  Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can activate or deactivate the symptom definitions in your policy.

## Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can activate or deactivate the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

### How the Policy Alert Definitions Work

VMware Aria Operations uses problems to trigger alerts. A problem manifests when a set of symptoms exists for an object, and requires you to take action on the problem. Alerts indicate problems in your environment. VMware Aria Operations generates alerts when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert occurs, VMware Aria Operations presents the triggering symptoms for you to take action.

Some of the alert definitions include predefined symptoms. When you include symptoms in an alert definition, and activate the alert, an alert is generated when the symptoms are true.

The Alert Definitions pane displays the name of the alert, the number of symptoms defined, the adapter, object types such as host or cluster, and whether the alert is activated as indicated by **Local**, deactivated as indicated by **not Local**, or inherited. Alerts are inherited with a green checkmark by default, which means that they are activated.

You can automate an alert definition in a policy when the highest priority recommendation for the alert has an associated action.

To view a specific set of alerts, you can select the badge type, criticality type, and the state of the alert to filter the view. For example, you can set the policy to send fault alerts for virtual machines.

### Where You Modify the Policy Alert Definitions

To modify the alerts associated with policies, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 3-12. Alert Definitions in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Object Type | Filters the alert definitions list by object type. |
| Filters | Limits the list based on the text you type.<br>You can also filter by:<br>■ Name<br>■ Criticality<br>■ Impact<br>■ State<br>■ Automate<br>■ Local Changes<br>■ Unsaved Changes<br>Impact indicates the health, risk, and efficiency badges to which the alerts apply.<br>Criticality indicates the information, critical, immediate, warning, or automatic criticality types to which the alert definition applies.<br>Automate indicates the actions that are activated for automation when an alert triggers, or actions that are deactivated or inherited. Actions that are activated for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to **Local** with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark. |
| Actions | Select one or more alert definitions and select activate, deactivate, or inherit to change the state for this policy. |

Table 3-12. Alert Definitions in the Create or Edit Policies Workspace (continued)

| Option | Description |
| --- | --- |
| Page Size | The number of alert definitions to list per page. |
| Alert Definitions data grid | Displays information about the alert definitions for the object types. The full name for Alert definition and the criticality icon appear in a tooltip when you hover the mouse over the Alert Definition name.<br>■ Alert Definition. Meaningful name for the alert definition.<br>■ State. Alert definition state, either activated, deactivated, or inherited from the base policy.<br>■ Automate. When the action is set to Local, the action is activated for automation when an alert triggers. Actions that are activated for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to **Local** with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark.<br>■ Symptom. Number of symptoms defined for the alert.<br>■ Criticality. Indicates the criticality of the alert.<br>■ Actionable Recommendations. Only recommendations with actions in the first priority, as they are the only ones you can automate.<br>■ Adapter. Data source type for which the alert is defined.<br>■ Object Type. Type of object to which the alert applies. |

If you do not configure the package, the policy inherits the settings from the selected base policy.

### Activating the Deactivated Alerts

Several out-of-the-box alerts have been deactivated to enhance your alert experience and reduce alert noise in your environment. The alerts that are triggered for these deactivated alerts are auto-cancelled and as a result, you may experience a dip in the number of alerts triggered. However, you can still activate these alerts in specific policies.

Perform the following steps to activate the deactivated alerts:

#### Prerequisites

The reason for deactivating these alerts is that there could be an overwhelming number of alerts when alerts are turned on for all objects, making it difficult to identify the ones that need immediate attention. It is recommended to exercise caution while activating the deactivated alerts for applicable policies.

Read the KB article, KB 91410 to know the list of deactivated alerts.

#### Procedure

1. From the left menu, click **Configure > Policies > Policy Definition**, and select the required policy.
2. In the right pane, click **Edit Policy**, and then select the **Alerts and Symptoms** tile.
3. Go to Filters and enter the name of the deactivated alert, and click **Apply**. You can refer to the KB article, KB 91410 for the list of deactivated alerts.

4 Select **Activated** from the **State** drop-down list or click **Actions** > **State** > **Activated**.

5 Click **Save**.

**Note** You can also activate all the deactivated alerts at once. To do this, filter the alerts by **Deactivated** State, click on the **Select All** option, and from the **Actions** drop-down list, click **State** > **Activated**.

**Note** You can also activate the deactivated alerts by creating a separate policy, adding custom groups in that policy, and activating the alert definitions in the required policy. By doing this, the deactivated alerts are activated in the user-defined policy and will apply to the objects in the custom group. For more details on custom groups, see Managing Custom Object Groups in VMware Aria Operations.

#### Results

The deactivated alerts are now active.

### Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can activate or deactivate the symptom definitions in your policy.

#### How the Policy Symptom Definitions Work

VMware Aria Operations uses symptoms that are activated to generate alerts. When the symptoms used in an alert definition are true, and the alert is activated, an alert is generated.

When a symptom exists for an object, the problem exists and requires that you take action to solve it. When an alert occurs, VMware Aria Operations presents the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

To assess objects for symptoms, you can include symptoms packages in your policy for metrics and super metrics, properties, message events, and faults. You can activate or deactivate the symptoms to determine the criteria that the policy uses to assess and evaluate the data collected from the objects to which the policy applies. You can also override the threshold, criticality, wait cycles, and cancel cycles.

The Symptoms pane displays the name of the symptom, the associated management pack adapter, object type, metric or property type, a definition of the trigger such as for CPU usage, the state of the symptom, and the trigger condition. To view a specific set of symptoms in the package, you can select the adapter type, object type, metric or property type, and the state of the symptom.

When a symptom is required by an alert, the state of the symptom is activated, but is dimmed so that you cannot modify it. The state of a required symptom includes an information icon that you can hover over to identify the alert that required this symptom.

### Where You Modify the Policy Symptom Definitions

To modify the policy package of symptoms, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

Table 3-13. Symptom Definitions in the Create or Edit Policies Workspace

| Option | Description |
| --- | --- |
| Object Type | Select an object type to view the symptom definitions list by the selected object type. |
| Filters | Limits the list based on the text you type.<br>You can also filter by:<br>■ Name<br>■ Criticality<br>■ Type<br>■ State<br>■ Local Changes<br>■ Unsaved Changes |
| Actions | Select one or more symptom definitions and select activate, deactivate, or inherit to change the state for this policy. |

Table 3-13. Symptom Definitions in the Create or Edit Policies Workspace (continued)

| Option | Description |
| --- | --- |
| Page Size | The number of symptom definitions to list per page. |
| Symptom Definitions data grid | Displays information about the symptom definitions for the object types. The full name for Symptom Definition appears in a tooltip when you hover the mouse over the Symptom Definition name.<br>■ Symptom Definition. Symptom definition name as defined in the list of symptom definitions in the Content area. Click this name to view the details of the symptom.<br>■ State. Symptom definition state, either activated, deactivated, or inherited from the base policy.<br>  ■ Activated. Indicates that a symptom definition will be included.<br>  ■ Activated (Force). Indicates state change due to a dependency.<br>  ■ Deactivated. Indicates that a symptom definition not be included.<br>  ■ Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will be included.<br>  ■ Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will not be included.<br>■ Threshold. To change the threshold, you must set the State to **Activated**, set the condition to **Override**, and set the new threshold in the Override Symptom Definition Threshold dialog box.<br>■ Type. Type of object to which the alert applies. Type determines whether symptom definitions that apply to HT and DT metrics, properties, events such as message, fault, and metric, and smart early warnings appear in the list.<br>■ Criticality. Indicates the criticality.<br>■ Adapter. Data source type for which the alert is defined.<br>■ Object Type. Object type on which the symptom definition must be evaluated.<br>■ Trigger. Static or dynamic threshold, based on the number of symptom definitions, the object type and metrics selected, the numeric value assigned to the symptom definition, the criticality of the symptom, and the number of wait and cancel cycles applied to the symptom definition.<br>■ To activate the metric threshold evaluation based on real-time data collected every 20 seconds, click the **Near Real-Time Monitoring** check box.<br>■ Condition. Activates action on the threshold. When set to Override, you can change the threshold. Otherwise set to default. |

If you do not configure the package, the policy inherits the settings from the selected base policy.

## Groups and Objects details

You can assign your local policy to one or more objects or groups of objects to have VMware Aria Operations analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

### How the Groups and Objects Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more objects or groups of objects. VMware Aria Operations uses the settings in the policy to analyze and collect data from the associated objects, and displays the data in dashboards, views, and reports.

### Where You Apply a Policy to Groups and Objects

To apply the policy to an object or groups of objects, from the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**.

### Groups and Objects Options

To apply the policy to an object or groups of objects, select the check box for the groups or objects in the workspace.

You can then view the groups and objects associated with the policy. From the left menu click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**. Click the **Custom Groups** tab to apply the policy to one or more groups of objects. Click the **Objects** tabs to apply the policy to one or more objects.

For more information about how to create an object group, see the topic called **Custom Object Groups Workspace to Create a New Group**.

For more information about how to create a policy, see Policy Workspace in VMware Aria Operations.

# Assigning Policies

The **Policy Assignment** workspace displays all the policies available in your environment.

## Where You Assign Policies

To create and modify policy assignments, from the left menu, click **Configure > Policies**, and then click **Policy Assignment**.

## How the Policy Assignment Workspace Works

You can assign policies to your environment to activate controls, view, and manage your object assignment scope.

| Option | Description |
| --- | --- |
| Left Pane: Displays the list of policies available in your environment and provides different filters to understand policy distribution. | |
| Inactive/Active | You can filter policies by active or inactive status. A policy is considered to be active only when it is assigned to an object or a custom group. |
| Visualization Type | You can sort active policies based on different criteria.<br>■ Fixed: Displays active policies with the sections of the policies fixed and equal to each other.<br>■ By Assignments: Displays the policies based on the number of objects assigned.<br>■ By Affected Objects: Displays policies based on the affected object count.<br>**Note** The sorting options are available only for the active policies. |
| Type Policy Name | Type the name of the policy in the search box to filter by policy name. |
| Policies | All the policies available in your environment are displayed in the left pane. Expand the policy card to view the priority, direct assignments, custom groups, affected objects, default assignment, and hierarchy of the policy.<br>**Note** The default assignments are displayed only for the default policy.<br>When you click on a policy, the objects and custom groups associated with the selected policy are displayed under **Assigned Objects** in the right pane. |
| Right Pane: Displays details of the selected policy and allows you to drag and drop objects or custom groups to the selected policy. | |
| Selected Policy | Displays the name of the selected policy. |

| Option | Description |
| --- | --- |
| All Objects | Displays all objects and custom groups available in your environment in the **Inventory** and **Custom Groups** tabs respectively.<br>To add new objects to a policy:<br>1 In the right pane, click **All Objects** > **Inventory**.<br>2 Select the objects you want to assign to the policy. You can also search for the objects by typing the object name in the search box.<br>3 Drag the items from the **Inventory** tab and drop them into the policy card on the left pane.<br>4 In the **Assign Objects** window, select one of the following options:<br>a Only this object: Select this if you want to apply changes only to the selected objects.<br>b Include child object: Select this if you want to apply changes to the child objects. You can define the depth of change by entering a number in the **Depth** field.<br>**Note** The maximum value for the Depth field is 10.<br>5 Click **Confirm**.<br>To add custom groups to a policy:<br>1 In the right pane, click **All Objects** > **Custom Groups**.<br>2 Select the custom groups you want to assign to the policy in the **Custom Groups** tab. You can also search for the custom groups by typing the custom group name in the search box.<br>3 Drag the items from the **Custom Groups** tab and drop them into the policy card on the left pane.<br>4 Click **Confirm**. |
| Assigned Objects | Displays the objects and custom groups assigned to the selected policy.<br>■ Name: Displays the name of the object/custom group.<br>■ Assignment Type: Displays the type of assignment.<br>■ Depth: Displays the depth of the child objects that the policy affects. Click the Edit icon to change the depth.<br>■ Action: Allows you to delete an object or custom group.<br>You can also drag objects/custom groups from **Assigned Objects** and drop them into the policy card of your choice. |

# Integrating Data Sources with VMware Aria Operations

4

You can extend the monitoring capabilities of VMware Aria Operations by installing and configuring integrations in VMware Aria Operations to connect to, and analyze data from external data sources in your environment. Once connected, you can use VMware Aria Operations to monitor and manage objects in your environment. These integrations are also referenced as Management Packs or as Solutions.

An integration can be a connection to a data source, or it might include predefined dashboards, widgets, alerts, and views.

Integrations can include cloud accounts, other accounts, dashboards, reports, alerts, and other content. The cloud accounts and other accounts comprise of adapters, using which VMware Aria Operations manages communication and integration with other products, applications, and functions. When an integration is installed and the accounts are configured, you can use the VMware Aria Operations analytics and alerting tools to manage the objects in your environment.

VMware Integrations include accounts for the following:

- Storage Devices
- VMware Aria Operations for Logs
- NSX for vSphere
- Network Devices
- VCM
- VMware Aria Hub

Third-party integrations include AWS, SCOM, EMC Smarts, and many others.

Other integrations such as the VMware Management Pack for NSX for vSphere, can be added to VMware Aria Operations.

VMware Aria Operations includes integrations that are pre-installed.

VMware Aria Operations also includes integrations that are bundled with VMware Aria Operations, but not activated.

For a fresh deployment of VMware Aria Operations, the activation status of integrations are as follows:

### Table 4-1. Integrations Activation Status

| Integration Name | Activated by Default? | Can be Deactivated? |
|---|---|---|
| vCenter | Yes | No |
| VMware Cloud on AWS | Yes | No |
| Oracle Cloud VMware Solution | No | Yes |
| Google Cloud VMware Engine | No | Yes |
| Azure VMware Solution | No | Yes |
| VMware Cloud on Dell EMC | No | Yes |
| VMware Cloud Foundation | No | Yes |
| VMware Infrastructure Health | No | Yes |
| AWS | No | Yes |
| Microsoft Azure | No | Yes |
| GCP | No | Yes |
| vSAN | Yes | No |
| NSX-T | Yes | No |
| Service Discovery | Yes | No |
| OS and Application Monitoring | Yes | No |
| Cloud Management Assessments | Yes | Yes |
| VMware Aria Automation8.x | No | Yes |
| VMware Aria Hub | No | Yes |
| VMware Aria Operations for Logs | No | Yes |
| VMware Aria Operations for Networks | No | Yes |
| Ping | No | Yes |
| PCI Compliance | No | Yes |
| ISO Compliance | No | Yes |
| HIPAA Compliance | No | Yes |
| FISMA Compliance | No | Yes |
| DISA Compliance | No | Yes |

Table 4-1. Integrations Activation Status (continued)

| Integration Name | Activated by Default? | Can be Deactivated? |
| --- | --- | --- |
| CIS Compliance | No | Yes |
| VMware Aria Hub | Yes | Yes |

The list of activated and available integrations, accounts, Repository, and Marketplace are all available from a central **Integrations** page in VMware Aria Operations. This page can be accessed from the left menu by clicking **Data Sources > Integrations**.

# Non-Native Management Packs and Management Pack Builder

If there's a technology without a current, native, management pack, you have the option to install a non-native pack or build your own using VMware Aria Operations Management Pack Builder.

VMware Aria Operations Management Pack Builder is a stand-alone appliance that enables the creation of custom management packs for use in VMware Aria Operations. This is a no-code solution for bringing in data from an external API and either creating new resources or extending your VMware and third party resources with new Data, Relationships, and Events.

| Non-Native Management Pack | User Guide |
| --- | --- |
| Aggregator | VMware Aria Operations Management Pack for Aggregator |
| Apache Hadoop | VMware Aria Operations Management Pack for Apache Hadoop |
| Apache Tomcat | VMware Aria Operations Management Pack for Apache Tomcat |
| AppDynamics | VMware Aria Operations Management Pack for AppDynamics |
| Arista EOS | VMware Aria Operations Management Pack for Arista EOS |
| Automation Orchestrator | VMware Aria Operations Management Pack for Automation Orchestrator |
| Cisco HyperFlex | VMware Aria Operations Management Pack for Cisco HyperFlex |
| Cisco Networking | VMware Aria Operations Management Pack for Cisco Networking |
| Cisco UCS | VMware Aria Operations Management Pack for Cisco UCS |
| Citrix ADC | VMware Aria Operations Management Pack for Citrix ADC |
| Citrix Virtual Apps and Desktops | VMware Aria Operations Management Pack for Citrix Virtual Apps and Desktops |

| Non-Native Management Pack | User Guide |
| --- | --- |
| Cloud Director | VMware Aria Operations Management Pack for Cloud Director |
| Cloud Director Availability | VMware Aria Operations Management Pack for Cloud Director Availability |
| Cohesity | VMware Aria Operations Management Pack for Cohesity |
| Datadog | VMware Aria Operations Management Pack for Datadog |
| Dell EMC Isilon | VMware Aria Operations Management Pack for Dell EMC Isilon |
| Dell EMC OpenManage | VMware Aria Operations Management Pack for Dell EMC OpenManage |
| Dell EMC PowerEdge | VMware Aria Operations Management Pack for Dell EMC PowerEdge |
| Dell EMC PowerMax and VMAX | VMware Aria Operations Management Pack for Dell EMC PowerMax and VMAX |
| Dell EMC Unity | VMware Aria Operations Management Pack for Dell EMC Unity |
| Dell EMC VNX (Block and File) | VMware Aria Operations Management Pack for Dell EMC VNX |
| Dell EMC XtremIO | VMware Aria Operations Management Pack for Dell EMC XtremIO |
| Docker | VMware Aria Operations Management Pack for Docker |
| Dynatrace | VMware Aria Operations Management Pack for Dynatrace |
| F5 BIG-IP | VMware Aria Operations Management Pack for F5 BIG-IP |
| Flowgate | VMware Aria Operations Management Pack for Flowgate |
| HCX | VMware Aria Operations Management Pack for HCX |
| Horizon | VMware Aria Operations Management Pack for Horizon |
| HPE 3Par StoreServ | VMware Aria Operations Management Pack for HPE 3Par StoreServ |
| HPE Nimble | VMware Aria Operations Management Pack for HPE Nimble |
| HPE OneView | VMware Aria Operations Management Pack for HPE OneView |
| HPE ProLiant | VMware Aria Operations Management Pack for HPE ProLiant |
| IBM HMC | VMware Aria Operations Management Pack for IBM HMC |
| IBM PowerVC | VMware Aria Operations Management Pack for IBM PowerVC |

| Non-Native Management Pack | User Guide |
| --- | --- |
| IBM SVC and Storwize | VMware Aria Operations Management Pack for IBM SVC and Storwize |
| Identity Manager | VMware Aria Operations Management Pack for Identity Manager |
| Integrated OpenStack | VMware Aria Operations Management Pack for Integrated OpenStack |
| Kubernetes | VMware Aria Operations Management Pack for Kubernetes |
| Lenovo Compute | VMware Aria Operations Management Pack for Lenovo Compute |
| Microsoft Azure SQL | VMware Aria Operations Management Pack for Microsoft Azure SQL |
| Microsoft Hyper-V | VMware Aria Operations Management Pack for Microsoft Hyper-V |
| Microsoft IIS Server | VMware Aria Operations Management Pack for Microsoft IIS Server |
| Microsoft SCOM | VMware Aria Operations Management Pack for Microsoft SCOM |
| Microsoft SQL Server | VMware Aria Operations Management Pack for Microsoft SQL Server |
| Mongo DB | VMware Aria Operations Management Pack for Mongo DB |
| MySQL | VMware Aria Operations Management Pack for MySQL |
| Nagios | VMware Aria Operations Management Pack for Nagios |
| NetApp E-Series | VMware Aria Operations Management Pack for NetApp E-Series |
| NetApp FAS/AFF | VMware Aria Operations Management Pack for NetApp FAS/AFF |
| NetApp Solidfire | VMware Aria Operations Management Pack for NetApp SolidFire |
| Network Devices | VMware Aria Operations Management Pack for Network Devices |
| New Relic | VMware Aria Operations Management Pack for New Relic |
| NSX Advanced Load Balancer | VMware Aria Operations Management Pack for NSX Advanced Load Balancer |
| Nutanix | VMware Aria Operations Management Pack for Nutanix |
| Oracle Database | VMware Aria Operations Management Pack for Oracle Database |
| Oracle Enterprise Manager | VMware Aria Operations Management Pack for Oracle Enterprise Manager |

| Non-Native Management Pack | User Guide |
|---|---|
| Palo Alto Networks | VMware Aria Operations Management Pack for Palo Alto Networks |
| PostgreSQL | VMware Aria Operations Management Pack for PostgreSQL |
| Pure Storage FlashArray | VMware Aria Operations Management Pack for Pure Storage FlashArray |
| Red Hat JBoss EAP | VMware Aria Operations Management Pack for Red Hat JBoss EAP |
| SAP | VMware Aria Operations Management Pack for SAP |
| SAP HANA | VMware Aria Operations Management Pack for SAP HANA |
| SDDC Health Monitoring | VMware Aria Operations Management Pack for SDDC Health Monitoring |
| ServiceNow | VMware Aria Operations Management Pack for ServiceNow |
| Skyline | VMware Aria Operations Management Pack for Skyline |
| SNMP | VMware Aria Operations Management Pack for SNMP |
| SolarWinds | VMware Aria Operations Management Pack for SolarWinds |
| Storage Devices | VMware Aria Operations Management Pack for Storage Devices |
| VMware Aria Cost powered by CloudHealth | VMware Aria Operations Management Pack for VMware Aria Cost powered by CloudHealth |
| VMware Aria Operations Compliance Pack for VMware Sovereign Cloud | VMware Aria Operations Compliance Pack for VMware Sovereign Cloud |
| VMware Aria Operations for Applications | VMware Aria Operations Management Pack for VMware Aria Operations for Applications |
| VMware Aria Operations Management Pack Builder | VMware Aria Operations Management Pack Builder |
| VMware Site Recovery Manager | VMware Aria Operations Management Pack for VMware Site Recovery Manager |
| VMware Tanzu Application Service | VMware Aria Operations Management Pack for VMware Tanzu Application Service |
| vSphere Replication | VMware Aria Operations Management Pack for vSphere Replication |

## Upgrade Considerations

The native integrations in VMware Aria Operations are reinstalled if VMware Aria Operations is upgraded.

If a new version of an integration is released, VMware sends an email to the organization owner with the new version information and the request to upgrade to the latest version. Integrations which require an update have the Upgrade Available sign on the top right corner of the integration card. With every release of a new version of an integration you must update to the latest version manually. If you fail to update to the new version by the time the next update for that integration is released, then VMware will send out a notice and update the pack to maintain compatibility since VMware supports the current version and one prior version on its cloud infrastructure.

Read the following topics next:

- The Integrations Page in VMware Aria Operations
- vSphere
- VMware Cloud on AWS
- VMware Cloud on AWS Outposts
- Azure VMware Solution
- Oracle Cloud VMware Solution
- Google Cloud VMware Engine
- VMware Cloud Foundation
- VMware Infrastructure Health
- VMware Cloud on Dell EMC
- Amazon Web Services
- Microsoft Azure
- OS and Application Monitoring
- Monitoring Application Services and Operating Systems using Open Source Telegraf
- Monitoring Physical Servers
- Service and Application Discovery
- Configuring Business Applications
- Application Integration
- Configuring Ping Adapter Instances
- VMware Aria Automation
- vSAN
- VMware Aria Operations for Networks
- NSX
- Configuring Cloud Federation Adapter
- Google Cloud Platform

- Integration of VMware Aria Operations and VMware Aria Hub

# The Integrations Page in VMware Aria Operations

Activate integrations, add, import, or export accounts from the **Integrations** page in VMware Aria Operations. Integrations are also referenced as Management Packs in VMware Aria Operations.

## How Integrations Work

Integrations can include dashboards, reports, alerts, and accounts. Each account contains adapters using which VMware Aria Operations integrates with other products, applications, and monitors their function and health.

## Where You Find Integrations

In the menu, click **Data Sources > Integrations**. By default, the **Accounts** tab opens. Use this tab to add, import, or export accounts. For more information, see The Accounts Tab.

Click the **Repository** tab to view the installed and available integrations. For more information, see The Repository Tab.

**Note** For more information on pre-activated integrations, see Chapter 4 Integrating Data Sources with VMware Aria Operations.

## Data Collection Notifications

The **Data Collection** bell icon on the menu provides quick access to status and critical notifications related to data collections. The icon indicates whether notifications exist, and whether any of them are critical. 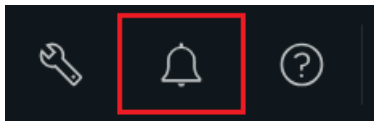

The list displays notifications about the data collections that are in progress, and indicates whether any of them have critical issues. The list groups the data collection notifications that are in progress into a single entry at the bottom of the list. To view the details about a collection, expand the notification.

Each notification displays the status of the last or current data collection, the associated adapter instance, and the time since the collection completed or an issue was identified. You can click a notification to open the Integrations page, where you can see further details, and manage adapter instances.

If problems occur with the data collections, VMware Aria Operations identifies those problems during each 5-minute collection cycle.

## Failed Integration Installation

If a integration installation fails, plug-ins related to the integration might appear in the Plug-ins page of VMware Aria Operations, even though the integration is not installed and does not appear on the Integrations page. When the integration installation fails, reinstall the integration.

## The Repository Tab

The repository tab lists all the installed and available integrations in VMware Aria Operations. Integrations are also referenced as Management Packs and Solutions.

### Where You Find the Repository Tab

From the left menu, click **Data Sources** > **Integrations**. Click the **Repository** tab. The page displays installed integrations as clickable cards.

### What Can You Do in the Repository Tab

In the **Repository** tab, when you click on an integration card, the details page is displayed.

The management pack details page consists of the following three tabs:

- Overview. Describes the management pack and its purpose.
- Metrics. Displays a list of metrics and properties for the various object types as defined by the adapter.

  **Note** This list of metrics and properties can display additional metrics and properties after account creation since other sources can also publish them. The actual list of published metrics and properties depends from policy customizations and monitored environment.

- Content. Displays the various content types and the data defined by that management pack.

Table 4-2. Management Pack Details Page Options

| Options | Descriptions |
| --- | --- |
| Name | Name of the management pack, the name of the vendor or manufacturer who created the management pack, and version number. |
| Activate | Installs the management pack. You can configure cloud management packs after activation from the **Repository** or **Accounts** tab of the **Integrations** page.<br>The activation starts only if all the cluster's nodes are accessible.<br>**Note** Some of the pre-Installed management packs are activated by default. You can configure them from the Management Pack Details Page or the **Accounts** tab of the **Integrations** page |
| Add Account | For more information on the accounts which are activated by default, see, Chapter 4 Integrating Data Sources with VMware Aria Operations. |

Table 4-2. Management Pack Details Page Options (continued)

| Options | Descriptions |
| --- | --- |
| **Horizontal Ellipses > Reset Default Content** | This option is only available for the vCenter management pack.<br>After you update your instance of VMware Aria Operations and select the option to overwrite, alert definitions and symptom definitions, you must overwrite your existing compliance alert definitions. |
| **Horizontal Ellipses > Deactivate** | You can uninstall a particular integration with its associated data, metadata, and the out of the box content. Select **I understand the risk and agree** to uninstall an integration. |

# The Accounts Tab

You can view and add accounts from the **Accounts** tab on the Integration page in VMware Aria Operations. You can add accounts for pre-activated accounts or you can add accounts after activating integrations from the repository tab.

## Where You Find the Accounts Tab

From the left menu, click **Data Sources > Integrations**. The **Accounts** tab opens by default. The accounts tab lists all the accounts configured for the different integrations available in VMware Aria Operations. You can filter them the accounts as per type to view accounts specific to an integration type.

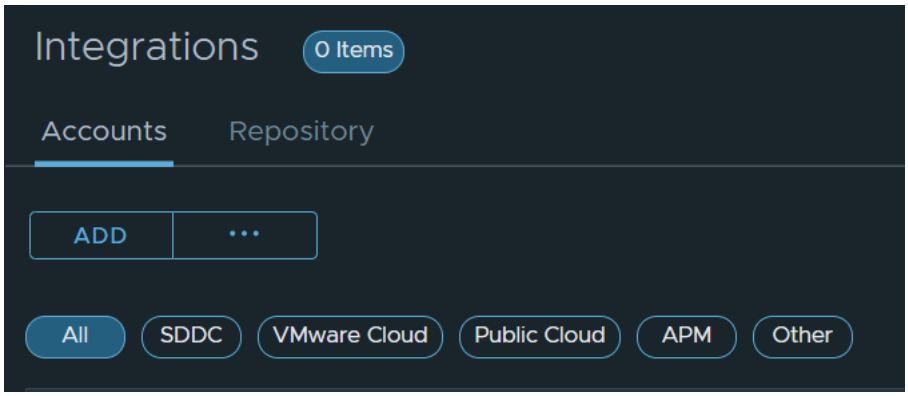


## What Can You Do in the Accounts Tab

You can add and configure accounts associated with integrations that are pre-activated or those that you activate in VMware Aria Operations. After you have configured the account, VMware Aria Operations can collect data from or send data to the target system. You can access the accounts page at any time to modify your adapter configurations.

Table 4-3. Accounts Page Options

| Option | Description |
| --- | --- |
| Add | Click **Add** to configure accounts for pre-activated integrations. For more information on pre-activated integrations, see Chapter 4 Integrating Data Sources with VMware Aria Operations.<br>To add accounts for any of the other available integrations, select the account type and click **Yes** in the **Installation Required** wizard.<br>**Note** You can also activate the integration from the Repository page and then add an account from the Accounts tab. For more information on activating available integrations from the Repository tab, see The Repository Tab. |
| Horizontal Ellipses | As a VMware Aria Operations admin, you can backup the adapter configurations before upgrading. Click the horizontal ellipses and then click **Export Accounts** to export all the adapter configurations. To import them into a different VMware Aria Operations instance, click **Import Accounts**. For more information, see Exporting and Importing Accounts. |

After configuring accounts you can view the account details and edit them in the Accounts page.

## Table 4-4. Accounts Grid Options

| | |
|---|---|
| Vertical Ellipses | To view specific account details, expand the account type. Click the > icon next to the account and then click the vertical ellipses to change the account configurations.<br>■ **Edit**. Allows you to edit the integrated adapter instance.<br>■ **Delete**. Removes the adapter instance and clears the objects associated with the instance from the system, including historical data and role assignments.<br>**Note** When you delete an account, you can choose to delete any related objects by selecting the checkbox, **Delete related objects**. If you do not wish to delete the related objects immediately, leave the check box unselected. The related objects are kept in the inventory for the duration of the retention period specified in the Global Settings page. If you recover an adapter instance before the end of the retention period, the related objects are unmarked and not deleted.<br>■ **Delete All**. Available for adapter instances with multiple sub-adapter configurations (such as the vCenter adapter). It removes the adapter instance along with its child adapter instances.<br>**Note** When you delete an account, you can choose to delete any related objects by selecting the checkbox, **Delete related objects**. If you do not wish to delete the related objects immediately, leave the check box unselected. The related objects are kept in the inventory for the duration of the retention period specified in the Global Settings page. If you recover an adapter instance before the end of the retention period, the related objects are unmarked and not deleted.<br>■ **Start/ Stop Collecting**. Starts or Stops the data collection process.<br>■ **Start/ Stop Collecting All**. Available for adapter instances with multiple sub-adapter configurations (such as the vCenter adapter). It starts or stops the data collection process of the adapter instance along with its child adapter instances.<br>■ **Object details**. Open the object in the object browser. |
| Name | Name that the vendor or manufacturer gave to the solution. |
| Status | Indicates the status of the solution and whether the adapter is collecting any data. If the status displays a green tick with the text OK, it means that the solution is collecting data. |
| Description | Typically, an indication of what the solution monitors or what data source its adapter connects to. |

Table 4-4. Accounts Grid Options (continued)

| | |
|---|---|
| Collector | Indicates the status of the solution. Data receiving shows that the solution is collecting data. |
| Filters | You can search the list of accounts according to the following criteria:<br>■ Name<br>■ Description<br>■ Solution<br>■ Adapter<br>■ Collector |

## Adding Accounts

Accounts specify the connection info for the public and hybrid cloud you want to monitor. You can add and configure accounts associated with integrations that are provided with or that you add to VMware Aria Operations. After you have configured the account, VMware Aria Operations can communicate with the target system. You can access the **Accounts** tab in the **Integrations** page at any time to modify your adapter configurations.

From the left menu, click **Data Sources > Integrations**. In the Accounts tab, click **Add** and select the integration you want to manage.

To manage accounts for the vSphere integration, see Configure a vCenter Server Cloud Account.

To manage the AWS cloud account, see Configuring the AWS Cloud Account.

To manage the Microsoft Azure cloud account, see Configuring the Microsoft Azure Cloud Account.

To manage the VMware Cloud on AWS cloud account, see Configuring the VMware Cloud on AWS Cloud Account.

### Prerequisites

**Note** Configure cloud proxy and activate the integration before adding and configuring accounts. For more information, see Configuring Cloud Proxies in VMware Aria Operations and The Repository Tab.

## Exporting and Importing Accounts

As a VMware Aria Operations admin, you can backup the adapter configurations before upgrading, export all the adapter configurations, and import them into a different VMware Aria Operations instance. You can export the adapter configurations from VMware Aria Operations on-prem to VMware Aria Operations, VMware Aria Operations to VMware Aria Operations on-

prem, VMware Aria Operations on-prem to on-prem instance, and VMware Aria Operations to cloud instance.

**Note** Users with "Export" permission can export and users with "Import" permission can import the adapter configurations.

Procedure

1 Export the adapter configuration.

   a From the left menu, click **Data Sources > Integrations**.

   b In the Accounts tab, select the adapter configuration(s) that you want to export, click the horizontal ellipses, and select **Export Accounts**.

   c Setup a new password to export data. The password should be at least 14 characters long and must include at least one numerical character, upper and lower case character, and special character.

   d Click **Export**.

   The adapter configurations are exported in .zip format. A password is used to encrypt the data. Use the same password while importing this file.

2 Import the adapter configuration.

   **Note** Before importing the content, ensure that you have exported the adapter configurations.

   a From the left menu, click **Data Sources > Integrations**.

   b Click the horizontal ellipses and select **Import Accounts**.

   c Click **Browse** to select the `.zip` file and enter the password that you had set while exporting the content.

   d If there is a conflict while importing the adapter configurations, you can either overwrite the existing adapter configurations or skip the import, which is the default option.

   e Click **Import** to import adapter configurations to the destination setup.

## Configuring Credentials in Integrations

Credentials are the user accounts that VMware Aria Operations uses to activate one or more integrations and associated adapters, and to establish communication with the target data sources. The credentials are supplied when you configure each adapter. You can create new credentials and modify the credential settings outside the adapter configuration or during adapter configuration to accommodate changes to your environment.

If you are modifying credentials to accommodate changes based on your password policy, you must validate your connection so that the adapters configured with these credentials begin using the new user name and password to communicate between VMware Aria Operations and the target system.

Another use of credential management is to remove misconfigured credentials. If you delete valid credentials that were in active use by an adapter, you disable the communication between the two systems. You cannot delete an active credential that is being used by an adapter instance.

If you need to change the configured credential to accommodate changes in your environment, you can edit the credential settings without being required to configure a new adapter instance for the target system. To edit credential settings, from the left menu, click **Data Sources > Integrations**. In the Accounts tab, click **Credentials** .

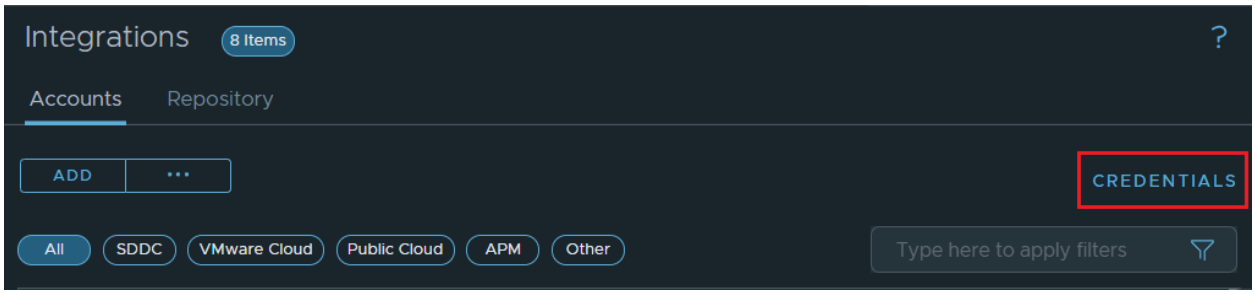


You can only view, add, modify, or delete credentials that you create or that are assigned to you. You can view unassigned credentials only if you have the required permissions.

**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The administrator must assign the credentials from **Orphan and Unassigned** page. For more information, see Managing Orphaned and Unassigned Content.

## Where You Find Credentials

From the left menu, click **Data Sources > Integrations**. In the **Accounts** tab, click the **Credentials** link on the upper right side.

Table 4-5. Credentials Options

| Option | Description |
| --- | --- |
| Toolbar options | Manages the selected credential.<br>■ **Add**. Add new credentials for an adapter type that you can later apply when configuring an adapter.<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>  ■ **Edit**. Modify the selected credentials, usually when the user name and password require a change. The change is applied to the current adapter credentials after you validate the connection and the data source continues to communicate with VMware Aria Operations.<br>  ■ **Delete** . Remove the selected credentials from VMware Aria Operations. If you have an adapter that uses these credentials, the communication fails and you cease monitoring the objects that the adapter was configured to manage. Commonly used to delete misconfigured credentials. You cannot delete an active credential that is being used by an adapter instance. |
| Filtering options | Limits the displayed credentials based on the adapter or credential types. |
| Credential name | Description of user-defined name that you provide to manage the credentials. Not the account user name. |
| Adapter Type | Adapter type for which the credentials are configured. |
| Credential Type | Type of credentials associated with the adapter. Some adapters support multiple types of credentials. For example, one type might define a user name and password, and another might define a pass code and key phrase. |

## Manage Credentials

To configure or reconfigure credentials that you use to activate an adapter instance, you must provide the collection configuration settings, for example, user name and password, that are valid on the target system. You can also modify the connection settings for an existing credential instance.

### Where You Manage Credentials

From the left menu, click **Data Sources > Integrations**. In the **Accounts** tab, click the **Credentials** link on the upper right side.

### Manage Credentials Options

The Manage Credentials dialog box is used to add new or modify existing adapter credentials. The dialog box varies depending on the type of adapter and whether you are adding or editing. The following options describe the basic options. Depending on the , the options other than the basic ones vary.

You can only view, add, modify, or delete credentials that you create or that are assigned to you. You can view unassigned credentials only if you have the required permissions.

**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The administrator must assign the credentials from **Orphan and Unassigned** page. For more information, see Managing Orphaned and Unassigned Content.

Table 4-6. Manage Credential Add or Edit Options

| Option | Description |
| --- | --- |
| Adapter Type | Adapter type for which you are configuring the credentials. |
| Credential Kind | Credentials associated with the adapter. The combination of adapter and credential type affects the additional configuration options. |
| Credential Name | Descriptive name by which you are managing the credentials. |
| User Name | User account credentials that are used in the adapter configuration to connect VMware Aria Operations to the target system. |
| Password | Password for the provided credentials. |

# vSphere

The vSphere solution connects VMware Aria Operations to one or more vCenter Server instances. You collect data and metrics from those instances, monitor them, and run actions in them.

VMware Aria Operations evaluates the data in your environment, identifying trends in object behavior, calculating possible problems and future capacity for objects in your system based on those trends, and alerting you when an object exhibits defined symptoms.

## Configuring the vSphere Solution

The vSphere solution is installed together with VMware Aria Operations. The solution provides the vCenter Server adapter which you must configure to connect VMware Aria Operations to your vCenter Server instances.

## How Adapter Credentials Work

The vCenter Server credentials that you use to connect VMware Aria Operations to a vCenter Server instance, determines what objects VMware Aria Operations monitors. Understand how these adapter credentials and user privileges interact to ensure that you configure adapters and users correctly, and to avoid some of the following issues.

- If you configure the adapter to connect to a vCenter Server instance with credentials that have permission to access only one of your three hosts, every user who logs in to VMware Aria Operations sees only the one host, even when an individual user has privileges on all three of the hosts in the vCenter Server.
- If the provided credentials have limited access to objects in the vCenter Server, even VMware Aria Operations administrative users can run actions only on the objects for which the vCenter Server credentials have permission.
- If the provided credentials have access to all the objects in the vCenter Server, any VMware Aria Operations user who runs actions is using this account.

## Controlling User Access to Actions

Use the vCenter Server adapter to run actions on the vCenter Server from VMware Aria Operations. If you choose to run actions, you must control user access to the objects in your vCenter Server environment. You control user access for local users based on how you configure user privileges in VMware Aria Operations. If users log in using their vCenter Server account, then the way their account is configured in vCenter Server determines their privileges.

For example, you might have a vCenter Server user with a read-only role in vCenter Server. If you give this user the VMware Aria Operations Power User role in vCenter Server rather than a more restrictive role, the user can run actions on objects because the adapter is configured with credentials that has privileges to change objects. To avoid this type of unexpected result, configure local VMware Aria Operations users and vCenter Server users with the privileges you want them to have in your environment.

To configure a vCenter Server cloud account, see Configure a vCenter Server Cloud Account in VMware Aria Operations.

## Configure a vCenter Server Cloud Account in VMware Aria Operations

To manage your vCenter Server instances in VMware Aria Operations, you must configure a cloud account for each vCenter Server instance. The cloud account requires the credentials that are used for communication with the target vCenter Server.

**Note** Any cloud account credentials you add are shared with other cloud account administrators and VMware Aria Operations collector hosts. Other administrators might use these credentials to configure a new cloud account or to move a cloud account to a new host.

### Prerequisites

- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data, see Privileges Required for Configuring a vCenter Adapter Instance. If the provided credentials have limited access to objects in vCenter Server, all users, regardless of their vCenter Server privileges see only the objects that the provided credentials can access. At a minimum, the user account must have Read privileges and the Read privileges must be assigned at the data center or vCenter Server level.
- Ensure you configure your cloud proxy in VMware Aria Operations, see Chapter 2 Collecting Data with Cloud Proxies in VMware Aria Operations.

### Procedure

1. From the left menu, click **Data Sources > Integrations > Accounts** tab.
2. Click **Add**.
3. On the Accounts Type page, click **vCenter**.
4. Enter a display name and description for the cloud account.
   - Display name. Enter the name for the vCenter Server instance as you want it to appear in VMware Aria Operations. A common practice is to include the IP address so that you can readily identify and differentiate between instances.
   - Description. Enter any additional information that helps you manage your instances.
5. (Optional) Select the **Physical Data Center** you want to associate with the vCenter Server cloud account.

   If no physical data center is created or if you want to create a new physical data center for your cloud account, you can add a new physical data center. For more information, see Adding Physical Data Centers in VMware Aria Operations
6. In the vCenter Server text box, enter the FQDN or IP address of the vCenter Server instance to which you are connecting.

   The vCenter Server FQDN or IP address must be reachable from the selected cloud proxy.

7 To add credentials for the vCenter Server instance, click the **Add** icon, and enter the required credentials. The vCenter credential must have `Performance > Modify intervals` permission activated in the target vCenter to collect VM guest metrics.

   Optionally, you can use alternate user credentials for actions. Enter an **Action User Name** and **Password**. If you do not enter an action user name and password, the default user specified is considered for actions.

   **Note** Credentials are stored in VMware Aria Operations and can be used for one or more instances of the vCenter Server.

   **Note** To monitor application services and operating systems, it is recommended that you enter action credentials with guest operations privileges such as `guest operation alias modification`, `guest operation alias query`, `guest operation modifications`, `guest operation program execution`, `guest operation queries`.

8 The collector for VMware Aria Operations is the cloud proxy. Specify the cloud proxy you just deployed as the collector for this vCenter cloud account.

9 To collect data at an interval of 20 seconds from the vCenter Server, click **Activate** for Near Real-Time Monitoring.

   The Near Real-Time Monitoring dialog box appears. Click the **I understand the affects of near real-time monitoring** check box and then click **OK**.

10 The cloud account is configured to run actions on objects in the vCenter Server from . If you do not want to run actions, deselect **Activate** for Operational Actions.

11 Click **Validate Connection** to validate the connection with your vCenter Server instance.

12 In the **Review and Accept Certificate** dialog box, review the certificate information.

   - If the certificate presented in the dialog box matches the certificate for your target vCenter Server, click **OK**.
   - If you do not recognize the certificate as valid, click **Cancel**. The test fails and the connection to vCenter Server is not completed. You must provide a valid vCenter Server URL or verify the certificate on the vCenter Server is valid before completing the adapter configuration.

13 To modify the advanced options regarding collectors, object discovery, or change events, expand the **Advanced Settings**.

   For information about these advanced settings, see Cloud Account Information - vSphere Account Options.

14 Click **Manage Registrations,** enter the credentials to register your VMware Aria Operations instance in vCenter Server, and then click **Register**.

   - Enter username and password for a user with vCenter Server register/unregister privileges.

- Select the collection credentials checkbox.

15 Click **Add** to save the configurations.

The vCenter Server adapter instance gets saved.

#### Results

The cloud account is added to the list. VMware Aria Operations begins collecting metrics, properties, and events from the vCenter Server instance. Depending on the number of managed objects, the initial collection can take more than one collection cycle. A standard collection cycle begins every five minutes.

#### What to do next

You can activate vSAN Configuration for your cloud account. For more information, see Configure a vSAN Adapter Instance.

You can use the vCenter Server for service discovery, see Configure Service Discovery.

## Privileges Required for Configuring a vCenter Adapter Instance

To configure your vCenter Adapter instance in VMware Aria Operations, you need sufficient privileges to monitor and collect data and to perform vCenter Server actions. You can configure these permissions as a single role in vCenter Server to be used by a single service account or configure them as two independent roles for two separate service accounts.

The vCenter Adapter instance monitors and collects data from vCenter Server and the vCenter Action adapter performs some actions in vCenter Server. So, for monitoring or collecting vCenter Server inventory and their metrics and properties, the vCenter Adapter instance needs credentials with the following privileges activated in vCenter Server.

**Note** The vCenter Server System Roles is created as a Read Only role with three system-defined privileges:: **System.Anonymous**, **System.View**, and **System.Read**. See, Using Roles to Assign Privileges.

### Table 4-7. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection

| Task | Privilege |
| --- | --- |
| Property Collection | **System > Anonymous**<br>**Note** This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Objects Discovery<br>Events Collection | **Profile-Driven Storage > View**<br>**Storage views > View**<br>**Profile-Driven Storage > Profile-Driven Storage View**<br>**Datastore > Browse Datastore**<br>**System > View**<br>**Note** This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Performance Metrics Collection | **Performance > Modify intervals**<br>**System > Read**<br>**Note** This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Service Discovery | For credential-based service discovery<br>**Virtual Machine > Guest Operations > Guest Operation alias modification**<br>**Virtual Machine > Guest Operations > Guest Operation alias query**<br>**Virtual Machine > Guest Operations > Guest Operation modifications**<br>**Virtual Machine > Guest Operations > Guest Operation program execution**<br>**Virtual Machine > Guest Operations > Guest Operation queries** |
| | For credential-less service discovery<br>**Virtual machine > Service configuration > Manage service configurations**<br>**Virtual machine > Service configuration > Modify service configuration**<br>**Virtual machine > Service configuration > Query service configurations**<br>**Virtual machine > Service configuration > Read service configuration** |
| VC Plugin | **Extension > Register extension**<br>**Extension > Unregister extension**<br>**Extension > Update extension** |
| Orphaned Disk | **Datastore > Browse datastore** |

### Table 4-7. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection (continued)

| Task | Privilege |
| --- | --- |
| Authentication on VMware Aria Operations using VC User and apply actions | **privilege.Global.com.vmware.label > VMware Aria Operations Read Only Role**<br>**privilege.Global.com.vmware.label > VMware Aria Operations Power User Role** |
| Optimize Container<br>Schedule Optimize Container<br>Automate Optimize Container | ■ **AutoDeploy -> Rule -> Create**<br>■ **AutoDeploy -> Rule -> Delete**<br>■ **AutoDeploy -> Rule -> Edit**<br>■ **AutoDeploy -> RuleSet -> Activate**<br>■ **AutoDeploy -> RuleSet -> Edit**<br>■ **Datastore -> Allocate Space**<br>■ **Global -> Global tag**<br>■ **Global -> System tag**<br>■ **Host -> Inventory -> Manage Cluster Lifecyle**<br>■ **Host -> Inventory -> Modify cluster**<br>■ **Resource -> Assign virtual machine to resource pool**<br>■ **Resource -> Migrate powered off virtual machine**<br>■ **Resource -> Migrate powered on virtual machine**<br>■ **Resource -> Query vMotion**<br>■ **Storage views -> Configure service**<br>■ **Storage views -> View**<br>■ **Virtual machine -> Edit Inventory > Move**<br>Privilege required for vCenter version 7.x:<br>■ **Profile-driven storage -> Profile-driven storage update**<br>■ **Profile-driven storage -> Profile-driven storage view**<br>Privilege required for vCenter version 8.x :<br>■ **VM storage policies -> Apply VM storage policies**<br>■ **VM storage policies -> Update VM storage policies**<br>■ **VM storage policies -> VM storage policies edit permissions**<br>■ **VM storage policies -> VM storage policies view permissions**<br>■ **VM storage policies -> View VM storage policies** |
| Provide data to vSphere Predictive DRS | **External stats provider > Update**<br>**External stats provider > Register**<br>**External stats provider > Unregister**<br>**vSphere Stats Privileges > Collect Stats Data**<br>**vSphere Stats Privileges > Modify Stats Configuration**<br>**vSphere Stats Privileges > Query Stats Data** |

### Table 4-7. Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection (continued)

| Task | Privilege |
| --- | --- |
| Tag Collection | **Global > Global tag**<br>**Global > Global health**<br>**Global > Manage custom attributes**<br>**Note** This privilege is required only if the tags are associated with custom attributes.<br>**Global > System tag**<br>**Global > Set custom attribute** |
| Monitoring and collecting data from vSphere with Tanzu | **Administrator**<br>**Note** Users with Non-Administrator or custom role must be added to the **ServiceProviderUser** group. **Administrator > Single Sign On > Users and Groups > Groups.**<br>The ServiceProviderUsers is a group in the vCenter Server Single Sign-On Domain. Members of this group can manage the vSphere with Tanzu and VMware Cloud on AWS infrastructure. |

### Table 4-8. Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions

| Task | Privilege |
| --- | --- |
| Set CPU Count for VM | **Virtual Machine > Configuration > Change CPU Count** |
| Set CPU Resources for VM | **Virtual Machine > Configuration > Change Resource** |
| Set Memory for VM | **Virtual Machine > Configuration > Change Memory** |
| Set Memory Resources for VM | **Virtual Machine > Configuration > Change Resource** |
| Delete Idle VM | **Virtual machine > Edit Inventory > Remove** |
| Delete Powered Off VM | **Virtual machine > Edit Inventory > Remove** |
| Create Snapshot for VM | **Virtual Machine > Snapshot Management > Create Snapshot** |
| Delete Unused Snapshots for Datastore | **Virtual Machine > Snapshot Management > Remove Snapshot** |
| Delete Unused Snapshot for VM | **Virtual Machine > Snapshot Management > Remove Snapshot** |
| Power Off VM | **Virtual Machine > Interaction > Power Off** |
| Power On VM | **Virtual Machine > Interaction > Power On** |
| Shut Down Guest OS for VM | **Virtual Machine > Interaction > Power Off** |

Table 4-8. Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions (continued)

| Task | Privilege |
| --- | --- |
| Move VM | ■ **Resource > Assign Virtual Machine to Resource Pool**<br>■ **Resource > Migrate Powered Off Virtual Machine**<br>■ **Resource > Migrate Powered On Virtual Machine**<br>■ **Datastore > Allocate Space**<br>■ **Virtual machine -> Edit Inventory > Move**<br>**Note** Combining these four permissions allows the service account to perform Storage vMotion and regular vMotion of an object therefore allowing VMware Aria Operations to perform the given operations. |
| Set DRS Automation | **Host > Inventory > Modify Cluster** |
| Provide data to vSphere Predictive DRS | **External stats provider > Update**<br>**External stats provider > Register**<br>**External stats provider > Unregister** |
| Reboot Guest OS for VM | **Virtual machine > Interaction > Reset** |

For more information about tasks and privileges, see Required Privileges for Common Tasks in the *vSphere Virtual Machine Administration Guide* and Defined Privileges in the *vSphere Security Guide*.

## Configure User Access for Actions

To ensure that users can run actions in VMware Aria Operations, you must configure user access to the actions.

You use role permissions to control who can run actions. You can create multiple roles. Each role can give users permissions to run different subsets of actions. Users who hold the administrator role or the default super user role already have the required permissions to run actions.

You can create user groups to add action-specific roles to a group rather than configuring individual user privileges.

### Procedure

1 From the left menu, click **Administration**, and then click the **Access Control** tile.

2 To create a role:

   a Click the **Roles** tab.

   b Click the **Add** icon, and enter a name and description for the role.

3 To apply permissions to the role, select the role, and in the Permissions pane, click the **Edit** icon.
   a Expand **Environment**, and then expand **Action**.
   b Select one or more of the actions, and click **Update**.
4 To create a user group:
   a Click the **User Groups** tab, and click the **Add** icon.
   b Enter a name for the group and a description, and click **Next**.
   c Assign users to the group, and click the **Objects** tab.
   d Select a role that has been created with permissions to run actions, and select the **Assign this role to the user** check box.
   e Configure the object privileges by selecting each adapter instance to which the group needs access to run actions.
   f Click **Finish**.

#### What to do next

Test the users that you assigned to the group. Log out, and log back in as one of the users. Verify that this user can run the expected actions on the selected adapter.

## Cloud Account Information - vSphere Account Options

To begin monitoring your environment with VMware Aria Operations, you configure the vSphere solution. The solution includes the vCenter Server cloud account that collects data from the target vCenter Server instances.

### Where You Find the Solution - vSphere

From the left menu, click **Data Sources > Integrations > Accounts** tab. Click **Add Account**, and then select the **vCenter** card.

### Account Information - vSphere Account Options

Configure and modify cloud accounts on the Account Information page.

### Table 4-9. Advanced Settings Options

| Option | Description |
| --- | --- |
| Advanced Settings | Provides options related to designating specific collectors to manage this cloud account, managing object discovery and change events. |
| Auto Discovery | Determines whether new objects added to the monitored system are discovered and added to VMware Aria Operations after the initial configuration of the cloud account.<br>■ If the value is true, VMware Aria Operations collects the information about any new objects that are added to the monitored system after the initial configuration. For example, if you add more hosts and virtual machines, these objects are added during the next collections cycle. This is the default value.<br>■ If the value is false, VMware Aria Operations monitors only those objects that are present on the target system when you configure the cloud account. |
| Process Change Events | Determines whether the cloud account uses an event collector to collect and process the events generated in the vCenter Server instance.<br>■ If the value is true, the event collector collects and publishes events from vCenter Server. This is the default value.<br>■ If the value is false, the event collector does not collect and publish events. |
| Activate Collecting vSphere Distributed Switch<br>Activate Collecting Virtual Machine Folder<br>Activate Collecting vSphere Distributed Port Group | When set to false, reduces the collected data set by omitting collection of the associated category. |
| Exclude Virtual Machines from Capacity Calculations | When set to true, reduces the collected data set by omitting collection of the associated category. |
| Maximum Number Of Virtual Machines Collected | Reduces the collected data set by limiting the number of virtual machine collections.<br>To omit data on virtual machines and have VMware Aria Operations collect only host data, set the value to zero. |
| Provide data to vSphere Predictive DRS | vSphere Predictive DRS proactively load balances a vCenter Server cluster to accommodate predictable patterns in the cluster workload.<br>VMware Aria Operations monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.<br>Predictive DRS must also be activated for the Compute Clusters managed by the vCenter Server instances monitored by VMware Aria Operations. Refer to the *vSphere Resource Management Guide* for details on activating Predictive DRS on a per Compute Cluster basis.<br>When set to true, designates VMware Aria Operations as a predictive data provider, and sends predicative data to the vCenter Server. You can only register a single active Predictive DRS data provider with a vCenter Server at a time. |
| Activate Actions | Activating this option helps in triggering the actions that are related to vCenter. |

Table 4-9. Advanced Settings Options (continued)

| Option | Description |
| --- | --- |
| Cloud Type | Provides an ability to identify the type of vCenter that is used in VMware Aria Operations. By default, the cloud type is set to Private Cloud.<br>The cloud types available are: Azure VMware Solution, Google Cloud VMware Engine, Hosted Private Cloud, Private Cloud, VMware Cloud on AWS, and VMware Cloud on Dell EMC. |
| vCenter ID | A globally unique identifier associated with the vCenter Server instance. |
| Deactivate collecting Guest File Systems with names containing | Provide comma separated list of strings. If these strings are found in any guest files system mount point name, that guest file system will not be collected. |
| Activate real time monitoring | The real time monitoring setting is deactivated by default. To collect real time data every 20 seconds change it to **true**. |
| Dynamic Thresholding | This setting is activated by default. |

You can find the vSphere Hardening Guides at http://www.vmware.com/security/hardening-guides.html.

Click **Save Settings** to finish configuration of the solution.

# VMware Cloud on AWS

VMware Cloud on AWS provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing private cloud like operating environment. To manage your VMware Cloud on AWS instances in VMware Aria Operations, you must configure a cloud account.

## Configuring VMware Cloud on AWS in VMware Aria Operations

To manage your VMware Cloud on AWS instances in VMware Aria Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS.

### Prerequisites

Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
- For data collection of bills, include either Billing Read-only or Organization Owner roles from All Organization Roles.

  **Note** The data collection of bills requires the bills to be available in the CSP.

- For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on AWS service roles.

Procedure

1 From the left menu, click **Data Sources > Integrations**.

2 On the Accounts tab, click **Add**.

3 On the Accounts Types page, click **VMware Cloud on AWS**.

4 Enter a display name and description for the cloud account.

- Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in VMware Aria Operations.
- Description. Enter any additional information that helps you manage your instances.

5 To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating a CSP API Token for VMware Cloud on AWS.

**Note** Enter the following details if you are using proxy server to access internet or public services.

- Proxy Host. A remote proxy server IP.
- Proxy Port. The port that is activated on a remote proxy server.
- Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
- Proxy Password. Password for the proxy server username.
- Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

**Note** The proxy credentials will be used by NSX adapters.

6 Determine which VMware Aria Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** The CSP token is used to access the publicly available VMware Cloud Services Portal API. It is recommended to use the **Default Collector Group** for this access. If you use a cloud proxy, ensure it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Cloud.

7 Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

   The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings > View Organization**.

8 Click **Validate Connection** to validate the connection.

9 You can monitor the costs of running your VMware Cloud on AWS infrastructure by bringing in the billing from VMware Cloud on AWS to VMware Aria Operations. To do so, activate the costing option in **Advanced Settings**.

   **Note** If the bills are not available in CSP, then the VMC infrastructure costs calculation will automatically switch from the bill-based calculation to the list-price based calculation.

10 Click **Save**.

   The page to configure the SDDC in VMware Cloud on AWS appears.

11 From the list of available SDDCs in VMware Cloud on AWS, click any one of the SDDCs that you want to monitor from VMware Aria Operations.

12 Configure the vCenter adapter:

   a Click the **vCenter** tab, and enter the required credentials.

      - Credential Name. The name by which you are identifying the configured credentials.
      - User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
      - Password. The vCenter password configured for that vCenter user name.

   b Select the required collector group.

      **Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS.

      Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met.

      For details, see Configuring Cloud Proxies in VMware Cloud.

      It is advised not to use the default collector groups as the VMware Cloud on AWS management gateway firewall rule does not allow traffic originating from any address.

      If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c If you have installed cloud proxy in VMware Cloud on AWS SDDC, the cloud proxy might not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:

- Request a new public IP in the VMware Cloud on AWS SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.
- Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.
- Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.

13 Click the **vSAN** tab. By default, the vSAN adapter is activated.

a Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.

b Select **Enable SMART data collection**, if required.

c Click **Validate Connection** to validate the connection.

14 Click the **NSX** tab. By default, the NSX adapter is activated.

a Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change the NSX adapter instance to use the private IP address of NSX manager by performing the following steps:

1 Navigate to **Environment > Inventory > Adapter Instances > NSX** and click your NSX adapter instance.

2 From the list of objects displayed, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.

3 Click **OK**.

15 Click **Save This SDDC**.

**Note** The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information, see Configure Service and Application Discovery.

The VMware Cloud on AWS account, with the configured SDDC, is added to the list.

## Known Limitations

Review the following list of feature limitations of VMware Cloud on AWS integration.

- The compliance workflows in VMware Aria Operations work for the virtual machines running on a vCenter Server in VMware Cloud on AWS. The compliance checks for VMware management objects such Hosts, vCenter, and so on, are not available.

- Workload optimization including pDRS and host-based business intent do not work because of the VMware managers cluster configurations.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Aria OperationsVMware Aria Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud does not support VMware Aria Operations SaaS plugin.
- You cannot log in to VMware Aria Operations using your VMware Cloud vCenter Server credentials.
- Credential-less service discovery is not supported on VMware Cloud on AWS.

## Generating a CSP API Token for VMware Cloud on AWS

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of Vmware Cloud on AWS.

### Prerequisites

- To configure the VMware Cloud on AWS Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.
- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

### Procedure

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.
5. Copy or save the generated token.

## Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Environment > Inventory > Adapter Instances > NSX**, and then click the user-created instance.

### Table 4-10. Object Types that NSX Discovers

| Object Type | Description |
| --- | --- |
| NSX Adapter Instance | The VMware Aria Operations Management Pack for NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

#### Procedure

1. In the menu, click **Environment > Inventory**, and then click **Adapter Instances > NSX**.
2. Select the adapter instance name to display the list of objects discovered by your adapter instance.
3. Slide the display bar to the right to view the object status.

| Object Status | Description |
| --- | --- |
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

4. Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

## Configuring VMware Cloud on AWS Government Cloud Endpoint in VMware Aria Operations

To manage your VMware Cloud on AWS instances in VMware Aria Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS.

### Prerequisites

- Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:
  - To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
  - For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on AWS service roles.
- To activate the cost calculations based on VMware Cloud on AWS GovCloud Pricing, you must modify the VMware Cloud on AWS rate card on the Cloud Providers tab in the Cost Settings page. For details on updating the rate card, see the VMware KB article 88488.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Account Types page, click **VMware Cloud on AWS**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. Under Advanced Settings, set **Environment** to **VMware Government Cloud** and **Billing Enabled** to **False**.
6. To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.
   - Credential Name. The name by which you are identifying the configured credentials.
   - CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating CSP API Token.

   **Note** Enter the following details if you are using a proxy server to access the Internet or public services.

   - Proxy Host. A remote proxy server IP.
   - Proxy Port. The port that is activated on a remote proxy server.

- Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
- Proxy Password. Password for the proxy server username.
- Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

**Note** The proxy credentials will be used by NSX adapters.

7 Determine which VMware Aria Operations collector or collector group is used in managing the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** It is recommended to use a dedicated Cloud Proxy in VMC on AWS government cloud to monitor the VMC on AWS government cloud endpoint.

Ensure that you have Internet connectivity for the collectors to work.

If you have installed cloud proxy in VMware Cloud on AWS government cloud SDDC, the cloud proxy might not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:

a Request a new public IP in the VMware Cloud on AWS government cloud SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.

b Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.

c Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.

8 Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings > View Organization**.

9 Under **Advanced Settings**, set **Billing Enabled** to **False**.

10 Click **Validate Connection** to validate the connection.

11 Click **Save**.

The page to configure the SDDC in VMware Cloud on AWS appears.

12 From the list of available SDDCs in VMware Cloud on AWS government cloud, click any one of the SDDCs that you want to monitor from VMware Aria Operations.

13 Configure the vCenter adapter:

a Click the **vCenter** tab, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
- Password. The vCenter password configured for that vCenter user name.

b Select the required collector group.

**Note** It is recommended to use a dedicated Cloud Proxy in VMC on AWS government cloud to monitor the VMC on AWS government cloud endpoint.

If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS.

Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met. For details, see "Configuring Cloud Proxies in VMware Cloud" topic in the *Getting Started with VMware Aria Operations* guide.

It is advised not to use the default collector groups as the VMware Cloud on AWS management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c If you have installed cloud proxy in VMware Cloud on AWS government cloud SDDC, the cloud proxy might not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:

- Request a new public IP in the VMware Cloud on AWS government cloud SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.
- Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.
- Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.
- Click **Next**.

14 Configure the vSAN adapter.

a Click the **vSAN** tab. By default, the vSAN adapter is activated.

b Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.

c Select **Enable SMART data collection**, if required.

d Click **Validate Connection** to validate the connection.

e Click **Next**.

15 Configure the NSX adapter.

a Click the **NSX** tab.

b Activate NSX configuration if it is deactivated.

c Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change the NSX adapter instance to use the private IP address of NSX manager by performing the following steps:

1 Navigate to **Environment > Inventory > Adapter Instances > NSX** and click your NSX adapter instance.

2 From the list of objects displayed, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.

3 Click **Next**.

16 Click **Save This SDDC**.

**Note** The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information, see Configure Service and Application Discovery.

The VMware Cloud on AWS government cloud account, with the configured SDDC, is added to the list.

## Known Limitations

- For Government Cloud Endpoint, bill-based costing is not supported and the costing defaults to rate card based costing.
- The compliance workflows in VMware Aria Operations checks for compliance only for virtual machines running on a vCenter servcer in VMware Cloud on AWS. The compliance checks for VMware management objects such as Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent does not work because VMware manages cluster configurations.

- Workload optimization for the cross-cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Aria Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud does not support the VMware Aria Operations SaaS plugin.
- You cannot log in to VMware Aria Operations using your VMware Cloud vCenter Server credentials.
- Credential-less service discovery is not supported on VMware Cloud on AWS.

## Generating CSP API Token

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of VMware Cloud on AWS.

### Prerequisites

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

### Procedure

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.
5. Copy or save the generated token.

## Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Environment > Inventory > Adapter Instances > NSX**, and then click the user-created instance.

Table 4-11. Object Types that NSX Discovers

| Object Type | Description |
| --- | --- |
| NSX Adapter Instance | The VMware Aria Operations Management Pack for NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

Procedure

1 In the menu, click **Environment > Inventory**, and then click **Adapter Instances > NSX**.

2 Select the adapter instance name to display the list of objects discovered by your adapter instance.

3 Slide the display bar to the right to view the object status.

| Object Status | Description |
| --- | --- |
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

4 Deselect the adapter instance name and expand the **Object Types** tag.

Each Object Type name appears with the number of objects of that type in your environment.

# VMware Cloud on AWS Outposts

VMware Cloud on AWS Outposts provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing private cloud like operating environment. To

manage your VMware Cloud on AWS Outposts instances in VMware Aria Operations, you must configure a cloud account.

## Configuring VMware Cloud on AWS Outposts in VMware Aria Operations

To manage your VMware Cloud on AWS Outposts instances in VMware Aria Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS Outposts.

### Prerequisites

Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMWare Cloud on AWS service roles.
- For data collection of bills, include either Billing Read-only or Organization Owner roles from All Organization Roles.

  **Note** The data collection of bills requires the bills to be available in the CSP.

- For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMWare Cloud on AWS service roles.
- To activate the cost calculations based on VMware Cloud on AWS Outposts pricing, you must modify the VMware Cloud on AWS Outposts rate card on the Cloud Providers tab in the Cost Settings page. For details on updating the rate card, see the VMware KB article KB88488.
- If you have subscribed to both VMC on AWS service and VMware Cloud on AWS Outposts service, ensure that they are in different CSP organizations. Otherwise, certain functionalities such as costing, and configuration maximums may not function as expected when both VMC on AWS service and VMware Cloud on AWS Outposts services are onboarded for monitoring within VMware Aria Operations.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Accounts Types page, click **VMware Cloud on AWS**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the VMware Cloud on AWS Outposts instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.

5 To add credentials for the VMware Cloud on AWS Outposts instance, click the **Add** icon, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating a CSP API Token for VMware Cloud on AWS Outposts.

**Note** Enter the following details if you are using proxy server to access internet or public services.

- Proxy Host. A remote proxy server IP.
- Proxy Port. The port that is available on a remote proxy server.
- Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
- Proxy Password. Password for the proxy server username.
- Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

**Note** The proxy credentials will be used by NSX adapters.

6 Determine which VMware Aria Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** The CSP token is used to access the publicly available VMware Cloud Services Portal API. It is recommended to use the **Default Collector Group** for this access. If you use a Cloud Proxy, ensure it has access to the Internet, or if the outbound internet access for the Cloud Proxy must be restricted, ensure the minimum Cloud Proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations.

If you have installed Cloud Proxy in VMware Cloud on AWS Outposts SDDC, the Cloud Proxy might not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed Cloud Proxy and allow Cloud Proxy to connect to vCenter, perform the following steps:

- Request a new public IP in the VMware Cloud on AWS Outposts SDDC where the Cloud Proxy was deployed. For more details, see Request or Release a Public IP Address.
- Add a new NAT rule for the internet that associates private IP of the Cloud Proxy with the public IP. For details, see Create or Modify NAT Rules.
- Add a firewall rule that allows incoming traffic from the public IP that was associated with Cloud Proxy VM in earlier step to vCenter.

7 Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

   The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings > View Organization**.

8 Click **Validate Connection** to validate the connection.

9 Bill-based costing is not supported in VMware Cloud on AWS Outposts. In the advanced settings, set "Billing Enabled" field to false

10 Click **Save**.

   The page to configure the SDDC in VMware Cloud on AWS Outposts appears.

11 From the list of available SDDCs in VMware Cloud on AWS Outposts, click any one of the SDDCs that you want to monitor from VMware Aria Operations.

12 Configure the vCenter adapter:

   a Click the **vCenter** tab, and enter the required credentials.

      - Credential Name. The name by which you are identifying the configured credentials.
      - User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
      - Password. The vCenter password configured for that vCenter user name.

   b Select the required collector group.

   **Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS Outposts.

   Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met.

   For details, see Configuring Cloud Proxies in VMware Aria Operations.

   It is advised not to use the default collector groups as the VMware Cloud on AWS Outposts management gateway firewall rule does not allow traffic originating from any address.

   If you have configured an HTTP proxy on your VMware Aria Operations Cloud Proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c If you have installed a Cloud Proxy in VMware Cloud on AWS Outposts SDDC, the Cloud Proxy might not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed Cloud Proxy and allow Cloud Proxy to connect to vCenter, perform the following steps:

- Request a new public IP in the VMware Cloud on AWS Outposts SDDC where the Cloud Proxy was deployed. For details, see Request or Release a Public IP Address.
- Add a new NAT rule for the internet that associates private IP of the Cloud Proxy with the public IP. For details, see Create or Modify NAT Rules.
- Add a firewall rule that allows incoming traffic from the public IP that was associated with Cloud Proxy VM in earlier step to vCenter.

13 Click the **vSAN** tab. By default, the vSAN adapter is activated.

a Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.

b Select **Enable SMART data collection**, if required.

c Click **Validate Connection** to validate the connection.

14 Click the **NSX** tab. By default, the NSX adapter is activated.

a Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change NSX adapter instance to use the private IP address of NSX manager by following the steps below:

1 Navigate to **Environment > Inventory > Adapter Instances > NSX** and click your NSX adapter instance.

2 In the list of objects shown, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.

3 Click **OK**.

15 Click **Save This SDDC**.

**Note** The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Outposts Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information about configuring the vCenter Service Discovery. see *Configure Service Discovery*.

The VMware Cloud on AWS Outposts account, with the configured SDDC, is added to the list.

## Known Limitations

Review the following list of feature limitations of VMware Cloud on AWS Outposts integration.

- VMware Cloud on AWS Outposts does not support Bill Based Costing.

- VMware Cloud on AWS Outposts does not support configuration maximums. You can ignore the displayed configuration maximums and related alerts.
- The compliance workflows in VMware Aria Operations work for the virtual machines running on a vCenter Server in VMware Cloud on AWS Outposts. The compliance checks for VMware management objects such Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent do not work because of the VMware Aria Operations cluster configurations in VMware Cloud on AWS Outposts.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Aria Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter Server interface.
- VMware Cloud on AWS Outposts does not support VMware Aria Operations plugin.
- You cannot log in to VMware Aria Operations using your vCenter Server credentials.
- Credential-less service discovery is not supported on VMware Cloud on AWS Outposts.
- If you have subscribed to both VMC on AWS service and VMware Cloud on AWS Outposts service, ensure that they are in different CSP organizations. Otherwise, certain functionalities such as costing, and configuration maximums may not function as expected when both VMC on AWS service and VMware Cloud on AWS Outposts services are onboarded for monitoring within VMware Aria Operations.

## Generating a CSP API Token for VMware Cloud on AWS Outposts

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of VMware Cloud on AWS Outposts.

### Prerequisites

- To configure the VMware Cloud on AWS Outposts Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.
- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

### Procedure

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.

4 Click **Generate**.

5 Copy or save the generated token.

## Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Environment > Inventory > Adapter Instances > NSX**, and then click the user-created instance.

Table 4-12. Object Types that NSX Discovers

| Object Type | Description |
| --- | --- |
| NSX Adapter Instance | The Aria Operations management pack for the NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

### Procedure

1 In the menu, click **Environment > Inventory**, and then click **Adapter Instances > NSX**.

2 Select the adapter instance name to display the list of objects discovered by your adapter instance.

3 Slide the display bar to the right to view the object status.

| Object Status | Description |
| --- | --- |
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

4 Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

# Azure VMware Solution

Azure VMware Solution provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

## Configuring an Azure VMware Solution Instance in VMware Aria Operations

To monitor Azure VMware Solution instances in VMware Aria Operations, you must add an Azure VMware Solution cloud account.

### Procedure

1 In the left menu, click **Data Sources > Integrations**.

2 In the Accounts tab, click **Add**.

3 On the Accounts Type page, click **Azure VMware Solution**.

4 From the **Add Cloud Account** page, enter a display name and description for the cloud account.

   - Name. Enter the name for the Azure VMware Solution instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.

5 Configure the Azure VMware Solution credentials.

| Option | Description |
| --- | --- |
| Subscription ID | Enter your subscription ID for Microsoft Azure. |
| Directory (Tenant) ID | Enter the directory (tenant) ID for your Azure Active Directory |

| Option | Description |
| --- | --- |
| Credential | Add the credentials used to access Azure VMware Solution by clicking the plus sign.<br>■ Enter an instance name for the credential values you are creating. This value is not the name of the adapter instance, but a friendly name for the secret credential.<br>■ Enter your application (Client) ID in your Azure Active Directory.<br>■ Enter the client secret that you generated for your application in the Microsoft Azure portal.<br>**Note** For steps to generate a client secret, see Generate a Client Secret for Azure VMware Solution.<br>**Note** The adapter uses the Azure public APIs to fetch the private clouds, hence internet connectivity to poll the APIs listed under Private Clouds - List in Subscription is mandatory. To view the APIs that the adapter uses internally, see Resource Manager - Overview.<br>■ Enter any required local proxy information for your network. |
| Collector/Group | Determine which VMware Aria Operations collector or collector group is used to manage the cloud account.<br>The best practice is to deploy the cloud proxy on each Private Cloud instance of Azure VMware Solution. If you use a cloud proxy, ensure it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations. |

6 Click **Validate Connection** to validate the connection.

7 Click **Save**.

The page to configure the Azure VMware Solution as a private cloud appears.

8 Click **Configure**.

9 Configure the desired vCenter adapter instance:

a Click the **Add** icon against **Credential**, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The vCenter user name.
- Password. The vCenter password configured for that vCenter user name.

b Select the required collector group.

**Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy Telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of Azure VMware Solution.

Select the cloud proxy deployed on the given vCenter Server and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

For details, see Configuring Cloud Proxies in VMware Aria Operations.

It is advised not to use the default collector groups as the Azure VMware Solution management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c Click **Next** to navigate to the **vSAN** section.

10 By default, the vSAN adapter is activated.

a By default, the vCenter Server referenced credential will be used for vSAN validation.

b Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter Server username, and password, and click **OK**.

c Select **Enable SMART data collection**, if required.

d Click **Validate Connection** to validate the connection.

e Click **Next**.

11 Configure the NSX adapter.

a By default, the NSX configuration is activated.

b Click the **Add** icon against **Credential**, and enter the required credentials.

- Credential Kind: Select the configured NSX instance.
- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The user name of the NSX instance.
- Password. The password of the NSX instance.
- Click **OK**.

c Click **Validate Connection** to validate the connection.

12 Click **Save This Private Cloud**.

> **Note** The Service Discovery adapter is optional. The steps to configure the Azure VMware Solution Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information about configuring the vCenter Service Discovery, see Configure Service and Application Discovery.

13 Determine which VMware Aria Operations collector or collector group is used to manage the cloud account.

The best practice is to deploy the cloud proxy on each private cloud instance of Azure VMware Solution. If you use a cloud proxy, ensure that it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations.

#### Results

The Azure VMware Solution account, with the configured private cloud, is added to the list.

## Generate a Client Secret for Azure VMware Solution

Create an Active Directory application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure a cloud account for Azure VMware Solution.

#### Prerequisites

- Ensure that you are using Microsoft Azure Cloud.
- Ensure that you have a valid subscription in the Microsoft Azure portal with an Active Directory integration.

#### Procedure

1 Log in to the Microsoft Azure portal.

2 Create an application and generate a secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.

Complete the following tasks:

a Create an Azure Active Directory application.

> **Note**
> - Ensure that the API Permission is 'Microsoft Graph User.Read'. Reader role for the Azure subscription with AVS private clouds deployed.
> - Custom role with read permissions on resource providers 'Microsoft.AVS', 'Microsoft.VMware' and 'microsoft.connectedvmwarevsphere'.

b Navigate to **Subscriptions** and select your subscription.

c In the left pane, click **Access Control (IAM)** and then click, **Add > Add Role Assignment**. Select the role you want to assign to the application. The minimum requirement is 'Reader' or above.

**Note** To provide access to a specific resource(s), create a Resource Group for the resource(s) and give access at the resource group level.

d Click **Select members**, and in the right pane, search for and add one or more members you want to assign to the role for the resource.

e Click **Review + assign**.

f Generate a client secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.

g Copy the subscription ID, directory (tenant) ID, application (client) ID, and client secret to use in your cloud account.

## Known Limitations

Review the following list of feature limitations of Azure VMware Solution integration.

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters might appear lower than expected and capacity remaining may appear higher than expected.
- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Cost calculation based on reference database is supported on Azure VMware Solution.
- The end-user on the vCenter Server on Azure VMware Solution has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to VMware Aria Operations using the credentials of the vCenter Server on Azure VMware Solution.
- The vCenter Server on Azure VMware Solution does not support the VMware Aria Operations plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.

# Oracle Cloud VMware Solution

Oracle Cloud VMware Solution provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

## Configuring an Oracle Cloud VMware Solution Instance in VMware Aria Operations

You must configure a dedicated cloud account for Oracle Cloud VMware Solution to manage your Oracle Cloud VMware Solution instances in VMware Aria Operations. For successful configuration, the cloud account requires the OCVS tenancy OCID, user OCID, API key fingerprint, API key's private key, and Oracle home region for the Oracle IAM user with appropriate privileges using which the cloud account authenticates, authorizes and communicates with the target Oracle cloud. SDDCs are auto-discovered after you save the cloud account for Oracle Cloud VMware Solution. You can then configure the credentials to monitor the underlying vCenter/vSAN and optionally the NSX and service discovery for each of the SDDCs.

### Prerequisites

- Create an IAM user in Oracle cloud with the required permissions, generate an API signing key, an API fingerprint, retrieve tenancy OCID, user's OCID, and the user's home region. Refer to the following Oracle documentation pages for more information: Required Keys and OCIDs and Managing Regions.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the **Accounts** tab, click **Add**.
3. On the **Accounts Types** page, click Oracle Cloud VMware Solution.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the Oracle Cloud VMware Solution instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. Enter the **Tenancy OCID** which is the Oracle cloud identifier for your tenancy. A tenancy is a secure and isolated partition of the Oracle Cloud Infrastructure to create, organize, and administer your cloud resources.
6. To add credentials for the Oracle Cloud VMware Solution instance, click the **Add** icon, and enter the required credentials.
   - Credential Name: The name by which you are identifying the configured credentials.
   - User OCID: Enter the Oracle Cloud identifier of the user calling the API.
   - API Key Fingerprint: Fingerprint for the public key that was added to this user.
   - API Key's Private Key: Enter the API key's private key in string format.

- Home Region: Enter the Oracle cloud home region. When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region.

  **Note** Enter the following details if you are using a proxy server to access the internet or public services.

- Proxy Host/IP: A remote proxy server IP.
- Proxy Port: The port that is activated on a remote proxy server.
- Proxy Username: Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
- Proxy Password: Password for the proxy server username.

7 Click **Validate Connection** to validate the connection.

8 Determine which VMware Aria Operations VMware Aria Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** It is recommended that use cloud proxy. Ensure that there is access to the Internet and it can reach the Oracle Cloud VMware Solution SDDC's vCenter, and NSX FQDNs. If the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. Ensure that you have Internet connectivity for the collectors to work. For more details see, Configuring Cloud Proxies in VMware Aria Operations.

**Note** If you have installed cloud proxy in an Oracle Cloud VMware Solution, the cloud proxy may not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps as described in the Oracle documentation in the following topic: Configuring Networking Connectivity for an SDDC.

9 Under Advanced Settings, enter the following details:

- **(Optional)** Configuration Limits File Name: The Oracle Cloud VMware Solution account uses the following default configuration maximum file: `ocvs_config_limits`. This file contains the Oracle Cloud VMware Solution configuration maximum soft and hard limits, and their configured value in VMware Aria Operations. If you have increased the limits for any of the Oracle Cloud VMware Solution configurations, you must update the this configuration file to reflect the new limit in VMware Aria Operations.

10 Click **Save**.

The page to configure the SDDCs in Oracle Cloud VMware Solution appears.

11 From the list of available SDDCs in Oracle Cloud VMware Solution, click any one of the SDDCs that you want to monitor from VMware Aria Operations.

12 Configure the vCenter adapter:

a Click the **vCenter** tab, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The vCenter Server user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter Server. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
- Password. The vCenter Server password configured for that vCenter Server user name.

b Select the required collector group.

**Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy Telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of Oracle Cloud VMware Solution.

Select the cloud proxy deployed on the given vCenter Server and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

For details, see Configuring Cloud Proxies in VMware Aria Operations.

It is advised not to use the default collector groups as the Oracle Cloud VMware Solution management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c If you have installed cloud proxy in an Oracle Cloud VMware Solution SDDC, the cloud proxy may not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps as described in the Oracle documentation in the following topic: Configuring Networking Connectivity for an SDDC.

13 Configure the vSAN Adapter.

a Click the **vSAN** tab. By default, the vSAN adapter is activated.

b Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter Server username, and password, and click **OK**.

c Select **Enable SMART data collection**, if required.

d Click **Validate Connection** to validate the connection.

e Click **Next**.

14 Configure the NSX adapter.

   a Click the NSX tab and the enter the required credentials.

   b Activate NSX configuration if it is deactivated.

   c Click the **Add** icon next to the **Credential** field and enter the required credentials.

      - Credential Kind: Select either the NSX client certificate credential option or NSX credentials.
      - Credential Name: The name by which you are identifying the configured credentials.
      - User Name: The user name of the NSX instance if you have selected NSX credentials as the credentials kind.
      - Password: The password of the NSX instance if you have selected NSX credentials as the credentials kind.
      - Client certificate data: Enter client certificate data if you have selected NSX client certificate credentials as the credentials kind.
      - Client key data: Enter client key data if you have selected NSX client certificate credentials as the credentials kind.
      - Click **OK**.

   d Click **Validate Connection** to validate the connection.

15 (Optional) Configure Service Discovery. For more information, see Configure Service and Application Discovery.

16 Click **Save This SDDC**.

## Known Limitations

Review the following list of feature limitations of Oracle Cloud VMware Solution integration.

- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Bill-based costing is not supported.
- You cannot log in to VMware Aria Operations using the credentials of the vCenter Server on Oracle Cloud.
- The vCenter Server on Oracle Cloud VMware Solution does not support the VMware Aria Operations plugin.
- vCenter Server and NSX standalone accounts configured with IPs are not recognized by the Oracle Cloud VMware Solution account and do not appear under the corresponding SDDC after Oracle Cloud VMware Solution is added. For upgrade, vCenter Server and NSX should be configured with FQDNs.

# Google Cloud VMware Engine

Google Cloud VMware Engine provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

## Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations

You must configure a dedicated cloud account for Google Cloud VMware Engine to manage your Google Cloud VMware Engine instances in VMware Aria Operations. For a successful configuration, the cloud account requires a Google Cloud Platform (GCP) project ID, service account JSON for the service account with appropriate privileges, and an optional CSP refresh token. The CSP refresh token is required if you would like to use bill-based costing and have purchased Google Cloud VMware Engine through VMware. Private clouds are auto-discovered after you save the cloud account for Google Cloud VMware Engine. You can then configure the credentials to monitor the underlying vCenter/vSAN and optionally the NSX and service discovery for each of the Private Clouds.

### Prerequisites

- Create a service account in Google Cloud Platform with at-least the viewer role privileges, note down the Google Cloud Platform project ID that you would like to manage from VMware Aria Operations. Refer to the following Google Cloud Platform documentation pages for more information: Creating and Managing Service Accounts
- Generate an optional CSP refresh token for bill-based costing in the VMware Cloud Services Portal (CSP). Navigate to **API Tokens** under **My Account** and generate a CSP API token with the billing read-only role for the Google Cloud VMware Engine service.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the **Accounts** tab, click **Add**.
3. On the **Accounts Types** page, click Google Cloud VMware Engine.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the Google Cloud VMware Engine instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.

5 Enter the Google Cloud Project ID in which Google Cloud VMware Engine service has been deployed.

   Google Cloud projects form the basis for creating, enabling, and using all Google Cloud services including managing APIs, enabling billing, adding and removing collaborators, and managing permissions for Google Cloud resources. Google Cloud projects are uniquely identified by an ID called Project ID. Refer to the following Google Cloud documentation for more information: Creating and Managing Projects.

6 To add credentials for the Google Cloud VMware Engine instance, click the **Add** icon, and enter the required credentials.

   - Credential Name: The name by which you are identifying the configured credentials.
   - Service Account JSON: Create a service account in Google cloud with at least the "viewer" role privileges and download its private key as a JSON file. Enter the contents of the JSON file in this field.

     **Note** You can create and use a single service account JSON that is common, similar to a super user account, for all the projects.

   - **(Optional)** CSP Refresh Token: Enter the API token if you want to use bill-based costing and Google Cloud VMware Engine was purchased from VMware. You can generate the CSP API refresh token from the Cloud Services Portal (CSP) with at least the billing read-only role for the Google Cloud VMware Engine service.

     **Note** Configure **all** the projects that are linked to the organisation for accurate bill based costing.

     **Note** If any project of the Google Cloud VMware Engine adapter instance is configured without the CSP token, then reference or rate card based costing will occur.

   - Proxy Host/IP: A remote proxy server IP.
   - Proxy Port: The port that is activated on a remote proxy server.
   - Proxy Username: Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
   - Proxy Password: Password for the proxy server username.

7 Click **Validate Connection** to validate the connection.

8 Determine which VMware Aria Operations VMware Aria Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** It is recommended that use cloud proxy. Ensure that there is access to the Internet and it can reach the Google Cloud VMware Engine Private Cloud's vCenter and NSX FQDNs. If the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. Ensure that you have Internet connectivity for the collectors to work. For more details see, Configuring Cloud Proxies in VMware Aria Operations.

**Note** If you have installed cloud proxy in an Google Cloud VMware Engine instance, the cloud proxy may not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps described in the Google Cloud documentation in the following topic: Configuring Internet Access for Workload VMs.

9 Under Advanced Settings, enter the following details:

- **(Optional)** Configuration Limits File Name: The Google Cloud VMware Engine account uses the following default configuration maximum file: `gcve_config_limits`. This file contains the Google Cloud VMware Engine configuration maximum soft and hard limits, and their configured value in VMware Aria Operations. If you have increased the limits for any of the Google Cloud VMware Engine configurations, you must create a new configuration file (from **Configure > Configuration Files > Solutions Configuration**) and update the name of the new configuration file in this field.
- Billing Enabled: Set the option to **true** to enable bill-based costing.

10 Click **Save**.

The page to configure the Private Clouds in Google Cloud VMware Engine appears.

11 From the list of available Private Clouds that are linked to the project configured in the Google Cloud VMware Engine instance, click any one of the Private Clouds that you want to monitor from VMware Aria Operations.

12 Configure the vCenter adapter:

a Click the **vCenter** tab, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The vCenter Server user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter Server. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
- Password. The vCenter Server password configured for that vCenter Server user name.

b Select the required collector group.

**Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy Telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each Private Cloud instance of Google Cloud VMware Engine.

Select the cloud proxy deployed on the given vCenter Server and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

For details, see Configuring Cloud Proxies in VMware Aria Operations.

It is advised not to use the default collector groups as the Google Cloud VMware Engine management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c If you have installed cloud proxy in an Google Cloud VMware Engine Private Cloud, the cloud proxy may not have outbound internet access to reach the VMware Aria Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps as described in the Google documentation in the following topic: Configuring Internet Access for Workload VMs.

13 Configure the vSAN Adapter.

a Click the **vSAN** tab. By default, the vSAN adapter is activated.

b Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter Server username, and password, and click **OK**.

c Select **Enable SMART data collection**, if required.

d Click **Validate Connection** to validate the connection.

e Click **Next**.

14 Configure the NSX adapter.

a Click the NSX tab and the enter the required credentials.

b Activate NSX configuration if it is deactivated.

c Click the **Add** icon next to the **Credential** field and enter the required credentials.

- Credential Kind: Select either the NSX client certificate credential option or NSX credentials.
- Credential Name: The name by which you are identifying the configured credentials.
- User Name: The user name of the NSX instance if you have selected NSX credentials as the credentials kind.

- Password: The password of the NSX instance if you have selected NSX credentials as the credentials kind.
- Client certificate data: Enter client certificate data if you have selected NSX client certificate credentials as the credentials kind.
- Client key data: Enter client key data if you have selected NSX client certificate credentials as the credentials kind.
- Click **OK**.

d Click **Validate Connection** to validate the connection.

15 (Optional) Configure Service Discovery. For more information, see Configure Service and Application Discovery.

16 Click **Save This Private Cloud**.

After the adapters and cloud accounts are configured, VMware Aria Operations discovers and monitors the environment that runs on Google Cloud VMware Engine.

## Known Limitations

Review the following list of feature limitations of Google Cloud VMware Engine integration.

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters may appear lower than expected and capacity remaining may appear higher than expected.
- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Bill-based costing is supported if Google Cloud VMware Engine was purchased through VMware. Customers can use rate card costing if Google Cloud VMware Engine was purchased from a partner or a vendor other than VMware.
- The end-user on the vCenter Server on Google Cloud VMware Engine has limited privileges. In-guest memory collection using VMware Tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to VMware Aria Operations using the credentials of the vCenter Server on Google Cloud VMware Engine.
- The vCenter Server on Google Cloud VMware Engine does not support the VMware Aria Operations plugin.

# VMware Cloud Foundation

In VMware Cloud Foundation, a workload domain is a policy- based resource construct with specific availability and performance attributes. It combines compute (vSphere), storage

(VMware vSAN), networking (NSX), and cloud management (VMware Aria Suite) to form a single consumable entity that creates logical resource pools across compute, storage, and networking.

A workload domain consists of one or more vSphere clusters, provisioned automatically by the Domain Manager. There are two types of workload domains - the management domain and the Virtual Infrastructure (VI) workload domains.

The management domain contains the Cloud Foundation management components which include an instance of vCenter Server and a three-node NSX Manager cluster for the management domain. It uses the vSAN storage.

The Virtual Infrastructure (VI) workload domain is created for user workloads. For each VI workload domain, you can choose the storage option (vSAN, NFS, or VMFS on FC). A VI workload domain can consist of one or more vSphere clusters. Each cluster must start with a minimum of three hosts and can scale up to a maximum of 64 hosts. The domain manager automates the creation of the workload domain and the underlying vSphere cluster(s).

For the first VI workload domain in your environment, the SDDC Manager deploys a vCenter Server and a NSX Manager cluster in the management domain. For each subsequent VI workload domain, the SDDC Manager deploys an additional vCenter Server. New VI workload domains can share the same NSX Manager cluster with the existing VI workload domain, or deploy a new NSX Manager cluster. However, VI workload domains cannot share the management domain's NSX Manager cluster.

Configure the VMware Cloud Foundation account in VMware Aria Operations to monitor these constructs of VMware Cloud Foundation.

## Configuring VMware Cloud Foundation Cloud Account in VMware Aria Operations

You must configure a VMware Cloud Foundation (VCF) cloud account to monitor a VMware Cloud Foundation environment. After configuration, all VMware Cloud Foundation domains are automatically discovered along with their connection details and credentials of underlying adapters for vCenter Server, vSAN, and NSX. You must save the domains want to monitor.

### Prerequisites

The VMware Cloud Foundation is not activated by default. You can activate it from the **Data Sources > Integrations > Repository** page or from the **Data Sources > Integrations > Add Accounts** page.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Accounts Types page, click **VMware Cloud Foundation**.

4 Enter a display name and description for the cloud account.

- Name. Enter the FQDN of the SDDC Manager for the VMware Cloud Foundation instance.
- Description. Enter any additional information that helps you manage your instances.

5 (Optional) Select the **Physical Data Center** you want to associate with the VMware Cloud Foundation cloud account.

If no physical data center is created or if you want to create a new physical data center for your cloud account, you can add a new physical data center. For more information, see Adding Physical Data Centers in VMware Aria Operations

6 Enter the FQDN for the **SDDC Manager** you are trying to connect.

7 To add credentials for the VMware Cloud Foundation instance, click the **Add** icon, and enter the required credentials.

The credentials must be associated with the admin role in VMware Cloud Foundation.

**Note** If you are using an AD or domain integrated account for authenticating your VMware Cloud Foundation instance, ensure you enter the username as username@domain or domain\username.

8 Determine which VMware Aria Operations collector or collector group or cloud proxy is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group or cloud proxy to manage the adapter processes for this instance.

9 Click **Validate Connection** to validate the connection.

The Review and Accept Certificate wizard appears. Click **Accept** to validate the continue the validation.

10 (Optional) Click **Advanced Settings** and enter the name of the **VCF Configuration Limit File Name**.

11 Click **Save**.

The page to configure the Domains in VMware Cloud Foundation appears.

12 Configure a vCenter Server Cloud account. For more information, see Configure a vCenter Server Cloud Account in VMware Aria Operations.

**Important**: When you configure the vCenter Server, the Credential is auto-generated for VMware Cloud Foundation.

13 Configure a vSAN Adapter instance. For more information, see Configure a vSAN Adapter Instance.

14 Configure the NSX adapter. For more information, see Configuring the NSX Adapter.

**Important**: When you configure the NSX adapter, the Credential is auto generated for VMware Cloud Foundation.

15 Configure Service Discovery. For more information, see Configure Service and Application Discovery.

   After the adapters and cloud accounts are configured, VMware Aria Operations discovers and monitors the environment that runs on VMware Cloud Foundation.

16 Click **Save this SDDC**.

#### Results

The VMware Cloud Foundation account, with the configured Domain, is added to the list.

#### What to do next

You can view the operations of your VMware Cloud Foundation accounts and their domains from the VMware Cloud Foundation Operations page. For more information, see Monitoring VMware Cloud Foundation (VCF) Operations

# VMware Infrastructure Health

The VMware Infrastructure Health monitors the VMware cloud management plane applications and provides metrics for their health and efficiency. With its dashboards, you can track monitor the overall health and configuration of the accounts, their applications, and services. Applications discovered by this management pack are listed in the applications section.

## VMware Infrastructure Health Monitoring in VMware Aria Operations.

The VMware Infrastructure Health in VMware Aria Operations monitors the following accounts.

1 vCenter Server

2 NSX

3 VMware Cloud Foundation

The VMware Infrastructure Health focusses mainly on the operations of the VMware Cloud Foundation and is integrated with it to provide an understanding of the availability, services, virtual machines, certificates, passwords, and active alerts through the VCF Operations page. You can view the account specific and domain specific data for VMware Cloud Foundation.

## Limitations

VMware Infrastructure Health does not monitor accounts that are in the warning state.

## Configuring the VMware Infrastructure Health

Use the VMware Infrastructure Health account to monitor the health of the VMware cloud management plane accounts including VMware Cloud Foundation.

### Prerequisites

The VMware Infrastructure Health is not activated by default. You can activate it from the **Data Sources > Integrations > Repository** page. Upon activation, VMware Infrastructure Health accounts are created for each collector and cloud proxy by default. After activation, if you add a new collector, a new VMware Infrastructure Health account gets created automatically. You can edit the VMware Infrastructure Health accounts.

### Procedure

1. From the left menu, click **Data Sources > Integrations**.
2. On the Accounts tab, expand VMware Infrastructure Health and click the **VMware Infrastructure Health Adapter** you want to edit.

   The Accounts Type page opens.
3. Enter a display name and description for the VMware Infrastructure Health account.
   - Name. Edit the name of the account.
   - Description. Enter any additional information that helps you manage your instances.
4. Click **Validate Connection** to validate the connection.
5. (Optional) Click **Advanced Settings** and enter a number for **vCenter Backup Job Objects Retention Period (in Days)** to set the period (in days) that the vCenter backup job objects are retained.
6. Click **Save**.

   The updates made to the VMware Infrastructure Health account are saved.

## Monitoring the Health of VCF Deployments

Use VMware Infrastructure Health to monitor the health of the VMware Cloud Foundation (VCF) cloud account in VMware Aria Operations. You can also monitor the VCF applications along with their services.

Before you start monitoring the VMware Cloud Foundation cloud account, you must activate and configure the VMware Cloud Foundation cloud account from the **Data Sources > Integrations > Add Accounts** page. The domains you want to monitor should have their vCenter and NSX configured in each domain. For more information, see Configuring VMware Cloud Foundation Cloud Account in VMware Aria Operations. All VMware Cloud Foundation domains are automatically discovered along with their connection details and the credentials of the underlying adapters for vCenter Server and NSX.

VMware Infrastructure Health collects data related to availability, services, virtual machines, certificates, passwords, and active alerts for each configured VMware Cloud Foundation cloud account and its domains. You can view this data from the **Environment > VCF Operations** page. For more information, see the 'Monitoring VMware Cloud Foundation (VCF) Operations' topic in the *VMware Aria Operations User Guide*.

**Note** For VMware Cloud Foundation version 4.5.1 and above, certificates and password data will be collected for the SDDC Manager, vCenter, and NSX applications only. For versions below 4.5.1, no data will be displayed for certificates and passwords.

### Monitoring the Management Domains

VMware Infrastructure Health will collect the following applications of the VMware Cloud Foundation management domain:

1. VMware Aria Operations for Logs application
2. vCenter application
3. VMware Identity Manager application
4. SDDC Manager application
5. NSX application
6. VMware Aria Automation application
7. VMware Aria Suite Lifecycle Manager application
8. application

### Monitoring the Workload Domains

VMware Infrastructure Health will collect the following applications of the VMware Cloud Foundation workload domain:

1. vCenter application
2. NSX application

## Monitoring the health of vCenter Server

Once you configure the vCenter Server cloud account, VMware Infrastructure Health starts collecting the health of the following vCenter services.

- vCenter Appliance Health Services
- vCenter Health Services
- vCenter Backup job
- vCenter NTP server

- vCenter Licensing Group

**Note**

- The list of available services differ based on the type of Operating System and the version of vCenter.
- The NTP and Backup Jobs services are available only for VCVA 6.5 version and above.

### Permissions Required to Discover vCenter Services

A user must have certain privileges to discover the services of vCenter.

The user should be a member of the 'SystemConfiguration.Administrator' group or have the administrator permissions to discover the vCenter services.

### vCenter Certificate Monitoring

Once the VMware Infrastructure Health collects the health of the vCenter health services, it monitors the ESXi host and vCenter root certificates, on a daily basis. It also alerts the user to avoid environment downtime. The certificate properties are published under the Certificate Summary group.

### vCenter Backup Job

The retention period for vCenter backup job monitoring is set by default to the last 7 days. You can configure the retention period to get more number of backup job items.

#### Backup Job Retention

To set the backup job retention period for vCenter:

1. From the left menu, click **Data Sources > Integrations**, and then click **Accounts** tab.
2. On the Account tab, expand the **VMware Infrastructure Health** and click the **Vertical Ellipses > Edit**.
3. Under **Advanced Settings**, in the **vCenter Backup Job Retention Period (in Days)** field, set the period (in days) for which vCenter backup jobs should be retained.

## Monitoring Health for NSX

You can monitor the health for NSX with the VMware Infrastructure Health.

### Resources Monitored from NSX

VMware Infrastructure Health monitors the health of the following NSX services:

- http
- nsx-ui
- ntp
- syslog

- mgmt-plane-bus
- install-upgrade
- snmp
- mgmt-plane-bus
- nsx-upgrade-agent
- migration-coordinator
- syslog
- nsx-ui
- node-mgmt
- nsx-message-bus
- migration-coordinator
- ssh
- snmp
- cm-inventory
- manager
- ssh
- liagent
- ntp
- install-upgrade
- http
- controller
- cluster-boot-manager
- search
- cm-inventory
- node-mgmt
- nsx-message-bus
- nsx-upgrade-agent
- telemetry
- telemetry
- manager
- controller
- cluster-boot-manager

- liagent
- search

### NSX Certificate Monitoring

Once the VMware Infrastructure Health collects the health of the NSX health services, it monitors the configured NSX Server and NSX manager node certificates, on a daily basis. It also alerts the user to avoid environment downtime. The certificate properties are published under the Certificate Summary group.

Alerts are raised based on the number of days set for critical, immediate, and warning category under the global setting. The NSX certificates contain the following three symptoms:

- has expired or will expire shortly - for critical
- expire shortly - for immediate
- about to expire - for warning

These symptoms are triggered if the certificate expiry date has reached the threshold condition set for critical, immediate, and warning. A notification is displayed explaining the alert information.

# VMware Cloud on Dell EMC

VMware Cloud on Dell EMC provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

## Configuring a VMware Cloud on Dell EMC Instance in VMware Aria Operations

You must configure a dedicated cloud account for VMware Cloud on Dell EMC to manage your VMware Cloud on Dell EMC instances in VMware Aria Operations. For successful configuration, the cloud account requires the CSP API token with appropriate privileges using which the cloud account authorizes and communicates with the target VMware Cloud on Dell EMC. SDDCs are auto-discovered after you save the cloud account for VMware Cloud on Dell EMC. You can then configure the credentials to monitor the underlying vCenter/vSAN and optionally the NSX and service discovery for each of the SDDCs.

#### Prerequisites

- Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs. For more information, see Generating CSP API Token for VMware Cloud on Dell EMC.
  - To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on Dell EMC service roles.
  - For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on Dell EMC service roles.

- Add the IP allowlist for accessing the vCenter Server and NSX Manager in VMware Cloud on Dell EMC service. For more information. see Add IP Allowlist for Accessing vCenter and NSX Manager.
- View the list of Known Limitations of the VMware Cloud on Dell EMC integration.

Procedure

1 From the left menu, click **Data Sources > Integrations**.

2 On the Accounts tab, click **Add**.

3 On the Accounts Types page, click VMware Cloud on Dell EMC.

4 Enter a display name and description for the cloud account.

   - Name. Enter the name for the VMware Cloud on Dell EMC instance as you want it to appear in VMware Aria Operations.
   - Description. Enter any additional information that helps you manage your instances.

5 (Optional) Select the **Physical Data Center** you want to associate with the VMware Cloud on Dell EMC cloud account.

   If no physical data center is created or if you want to create a new physical data center for your cloud account, you can add a new physical data center. For more information, see Adding Physical Data Centers in VMware Aria Operations

6 To add credentials for the VMware Cloud on Dell EMC instance, click the **Add** icon, and enter the required credentials.

   - Credential Name. The name by which you are identifying the configured credentials.
   - CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating CSP API Token for VMware Cloud on Dell EMC.

   **Note** Enter the following details if you are using a proxy server to access the internet or public services.

   - Proxy Host. A remote proxy server IP.
   - Proxy Port. The port that is activated on a remote proxy server.
   - Proxy Username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as **`username@domain name`**.
   - Proxy Password. Password for the proxy server username.
   - Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

7 Determine which VMware Aria Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**Note** The CSP token is used to access the publicly available VMware Cloud Services Portal API. It is recommended to use the **Default Collector Group** for this access. If you use a cloud proxy, ensure it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations .

8 Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings > View Organization**.

9 Click **Validate Connection** to validate the connection.

10 (Optional) Click **Advanced Settings**.

a The VMware Cloud on Dell EMC account uses the following default configuration maximum file: `vmc-d_config_limits`. Enter the name of the **VMC-D Configuration Limit File Name**.

b Set the **Billing Activated** setting to **true** to activate bill based costing and to monitor the costs of running your VMware Cloud on Dell EMC instance.

The bill based costing retrieves the actual bills and uses cost data from the bills in the cost overview homepage and other cost related functionalities. Bill based cost data retrieval depends on the invoice generation by your provider.

**Note** If the bills are not available in Cloud Services Platform (CSP), then the VMware Cloud on Dell EMC instance costs calculation will automatically switch from the bill-based calculation to the list-price based calculation.

11 Click **Save**.

The page to configure the SDDC in VMware Cloud on Dell EMC appears.

12 Configure the vCenter adapter:

a Click the **Add** icon, and enter the required credentials.

- Credential Name. The name by which you are identifying the configured credentials.
- User Name. The vCenter Server user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter Server. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.

- Password. The vCenter Server password configured for that vCenter Server user name.

b Select the required collector group.

**Note** If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on Dell EMC.

Select the cloud proxy deployed on the given vCenter Server and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

For details, see Configuring Cloud Proxies in VMware Aria Operations .

It is advised not to use the default collector groups as the VMware Cloud on Dell EMC management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Aria Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

c Click **Next**.

13 By default, the vSAN adapter is activated.

a Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter Server username, and password, and click **Ok**.

b Select **Activate SMART data collection**, if required.

c Click **Validate Connection** to validate the connection.

d Click **Next**.

14 (Optional) Configure Service Discovery. For more information, see Configure Service and Application Discovery.

After the adapters and cloud accounts are configured, VMware Aria Operations discovers and monitors the environment that runs on VMware Cloud on Dell EMC.

15 Activate the NSX adapter.

a Click **Validate Connection** to validate the connection.

b Click **Next**.

16 Click **Save This SDDC**.

### Results

The VMware Cloud on Dell EMC account, with the configured SDDC, is added to the list.

## Known Limitations

Review the following list of feature limitations of VMware Cloud on Dell EMC integration.

- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Credential less service discovery is not supported for VMware Cloud on Dell EMC.
- The end-user on the vCenter Server on VMware Cloud on Dell EMC has limited privileges. In-guest memory collection using VMware Tools is not supported with virtual machines. Active and consumed memory utilization continue to work in this case.
- The vCenter Server on VMware Cloud on Dell EMC does not support the VMware Aria Operations plugin.
- You cannot log in to VMware Aria Operations using the credentials of the vCenter Server on VMware Cloud on Dell EMC.
- Workload optimization including pDRS and host-based business intent does not work because of the cluster configurations of VMware managers.
- Workload optimization is not aware of resource pools and places the virtual machines at the cluster level and outside the resource pool. You can manually correct this in the vCenter Server user interface.
- Compliance checks for VMware management objects such as hosts and vCenter Server are not available.

## Generating CSP API Token for VMware Cloud on Dell EMC

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of VMware Cloud on Dell EMC.

### Prerequisites

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on Dell EMC service roles.
- For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on Dell EMC service roles.

### Procedure

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.

5 Copy or save the generated token.

## Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Environment** > **Inventory** > **Adapter Instances** > **NSX**, and then click the user-created instance.

Table 4-13. Object Types that NSX Discovers

| Object Type | Description |
| --- | --- |
| NSX Adapter Instance | The VMware Aria Operations Management Pack for NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

### Procedure

1 In the menu, click **Environment** > **Inventory**, and then click **Adapter Instances** > **NSX**.

2 Select the adapter instance name to display the list of objects discovered by your adapter instance.

3 Slide the display bar to the right to view the object status.

| Object Status | Description |
| --- | --- |
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

4 Deselect the adapter instance name and expand the **Object Types** tag.

Each Object Type name appears with the number of objects of that type in your environment.

# Amazon Web Services

Amazon Web Services (AWS) is an embedded adapter with diagnostic dashboards for VMware Aria Operations. The adapter collects metrics from AWS.

## Supported AWS Services

AWS supports the following services in VMware Aria Operations.

| Service | Object | Description |
|---|---|---|
| Amazon MQ* | Amazon MQ Broker<br>Amazon MQ Node<br>Amazon MQ Topic<br>Amazon MQ Queue<br>Amazon MQ Virtual Host | Provides fully managed service for open source message brokers |
| Amazon Simple Queue Service (SQS) | SQS Queue | Provides fully managed message queues for microservices, distributed systems, and serverless applications |
| Amazon EC2 | EC2 Instance<br>EC2 Instance Type | Provides a secure and resizable compute capacity to support virtually any workload |
| Amazon EC2 Auto Scaling | EC2 Auto Scaling Group | Allows you to add or remove compute capacity to meet changes in demand |
| AWS Lambda | Lambda Function | Provides a serverless compute service that lets you run code without provisioning or managing servers, creating workload-aware cluster scaling logic, maintaining event integrations, or managing runtimes |
| AWS Elastic Beanstalk | Elastic Beanstalk Application<br>Elastic Beanstalk Environment | Provides an easy-to-use service for deploying and scaling web applications and services developed with Java, .NET, PHP, Node.js, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS |
| Amazon Elastic Container Registry (ECR) | ECR Container Repository<br>ECR Container Image | Allows you to share and deploy container software, publicly or privately |
| Amazon Elastic Container Service (ECS) | ECS Cluster<br>ECS Service | Provides a highly secure, reliable, and scalable way to run containers |

| Service | Object | Description |
| --- | --- | --- |
| Amazon Elastic Kubernetes Service (EKS) | EKS Cluster<br>EKS Cluster ARN | Provides the flexibility to start, run, and scale Kubernetes applications in the AWS cloud or on-premises |
| Amazon Relational Database Service (RDS) | RDS DB Instance<br>RDS Instance Type | Allows you to set up, operate, and scale a relational database in the cloud |
| Amazon DynamoDB | DynamoDB<br>DynamoDB Accelerator Cluster<br>DynamoDB Accelerator Node | Provides fast and flexible NoSQL database service for any scale |
| Amazon ElastiCache | ElastiCache Cluster<br>ElastiCache Node | Provides fully managed in-memory data store, compatible with Redis or Memcached |
| Amazon WorkSpaces | Workspace | Provides a managed, secure Desktop-as-a-Service (DaaS) solution |
| AWS CloudFormation | CloudFormation Stack | Allows you to speed up cloud provisioning with infrastructure as code |
| Amazon Virtual Private Cloud (Amazon VPC) | VPC<br>VPC NAT Gateway<br>VPC VPN Connection | Allows you to build on a logically isolated virtual network in the AWS cloud |
| Amazon CloudFront | CloudFront Distribution | Provides a fast, highly secure, and programmable content delivery network (CDN) |
| Amazon Route 53 | Route53 Hosted Zone | A highly available and scalable cloud Domain Name System (DNS) web service |
| AWS Direct Connect | AWS Direct Connect | A cloud service solution that makes it easy to establish a dedicated network connection from your premises to AWS |
| Elastic Load Balancing | Elastic Load Balancing | Allows you to distribute network traffic to improve the scalability of your applications |
| Amazon Simple Storage Service (S3) | S3 Bucket | Provides object storage built to store and retrieve any amount of data from anywhere |
| Amazon Fargate/ECS | ECS Cluster<br>ECS Service | Provides serverless compute for containers |
| Amazon Elastic Block Store (EBS) | EBS | An easy-to-use, high-performance, block-storage service designed for use with Amazon Elastic Compute Cloud (EC2) for both throughput and transaction-intensive workloads at any scale |

| Service | Object | Description |
| --- | --- | --- |
| Amazon Redshift | Redshift<br>Redshift Cluster<br>Redshift Node | Allows you to analyze all of your data with the fastest and most widely used cloud data warehouse |
| Amazon EMR | EMR Job Flow | Provides easily run and scale Apache Spark, Hive, Presto, and other big data frameworks |
| Amazon Elastic File System (EFS) | EFS | A simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources |
| Amazon Athena | Amazon Athena WorkGroup | An interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL |
| Amazon CloudSearch | Amazon CloudSearch Domain | A managed service in the AWS Cloud that makes it simple and cost-effective to set up, manage, and scale a search solution for your website or application |
| Amazon OpenSearch | Amazon ElasticSearch Domain<br>Amazon ElasticSearch Node | Provides a managed service to easily deploy, operate, and scale OpenSearch clusters in the AWS Cloud |
| Amazon Elasticsearch Service* | Amazon ElasticSearch Domain<br>Amazon ElasticSearch Node | Provides fully managed, scalable, and secure Elasticsearch service |
| Amazon Kinesis Data Firehose | Amazon Kinesis Firehose | Allows you to prepare and load real-time data streams into data stores and analytics services |
| Amazon Kinesis Data Analytics | Amazon Kinesis Analytics | Provides actionable insights from streaming data with serverless Apache Flink |
| Amazon Kinesis Data Streams | Amazon Kinesis Stream | Allows you to easily collect, process, and analyze video and data streams in real-time |
| Amazon Kinesis Video Streams | Amazon Kinesis Video Stream | Provides a fully managed AWS service that can be used to stream live video from other devices to AWS Cloud or build applications for real-time video processing or batch-oriented video analytics. |
| Amazon Managed Streaming for Apache Kafka | Amazon MSK Cluster | Provides fully managed, highly available, and secure Apache Kafka service |
| AWS Step Functions* | Amazon Step Functions State Machine | Allows you to assemble functions into business-critical applications |

| Service | Object | Description |
| --- | --- | --- |
| Amazon EventBridge* | Amazon EventBridge Rule | Provides serverless event bus that connects application data from your apps, SaaS, and AWS services |
| Amazon Simple Notification Service (SNS)* | Amazon SNS Topic | Provides fully managed pub/sub messaging, SMS, email, and mobile push notifications |
| Amazon Simple Workflow Service (SWF)* | Amazon SWF Workflow Type<br>Amazon SWF Activity Type<br>Amazon SWF Task List | Allows developers to build, run, and scale background jobs that have parallel or sequential steps |
| Amazon WorkMail* | Amazon WorkMail Organization | A secure, managed business email and calendar service with support for existing desktop and mobile email client applications |
| Amazon Connect* | Amazon Connect Instance | An easy-to-use omnichannel cloud contact center |
| Amazon Pinpoint* | Amazon Pinpoint Application | A multichannel marketing communication service |
| Amazon Neptune | Amazon Neptune DB Instance<br>Amazon Neptune DB Cluster<br>Amazon Neptune DB Engine | Provides fast, reliable graph database built for the cloud |
| Amazon Quantum Ledger Database (QLDB) | Amazon QLDB Ledger<br>Amazon QLDB Stream | Provides a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log |
| Amazon DocumentDB | Amazon DocDB DB Instance<br>Amazon DocDB DB Cluster<br>Amazon DocDB DB Engine | Provides a fast, scalable, and highly available MongoDB-compatible database service |
| Amazon Keyspaces* | Amazon keyspace | A scalable, highly available, and managed Apache Cassandra–compatible database service |
| Amazon Timestream | Amazon Timestream Database | Provides fast, scalable, serverless time series database |
| AWS CodeBuild* | Amazon CodeBuild Project | Allows you to build and test code with continuous scaling. Pay only for the build time you use |
| Amazon AppStream 2.0* | Amazon AppStream Fleet | A fully managed non-persistent application and desktop streaming service |
| Amazon GameLift* | Amazon GameLift Fleet<br>Amazon GameLift Queue | Allows you to enhance multiplayer experiences with dedicated cloud servers |
| AWS IoT* | Amazon IoT Protocol | Provides easily and securely connect devices to the cloud |
| AWS IoT Analytics* | Amazon IoTAnalytics Data Set | Provides analytics for IoT devices |

| Service | Object | Description |
| --- | --- | --- |
| Amazon Kendra* | Amazon Kendra Index<br>Amazon Kendra DataSource | Provides highly accurate intelligent search service powered by machine learning |
| Amazon Lex* | Amazon Lex Bot | A conversational AI for Chatbots |
| AWS AppSync* | AWS AppSync GraphQL API | Allows you to accelerate application development with scalable GraphQL APIs |
| Amazon Cognito* | Amazon Cognito User Pool | A simple and secure user sign-up, sign-in, and access control |
| AWS WAF - Web Application Firewall* | AWS WAF WebACL | A web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources |
| AWS Storage Gateway* | Amazon Storage Gateway | A hybrid cloud storage service that gives you on-premises access to virtually unlimited cloud storage |
| Amazon Glue | Amazon Glue | A simple, scalable, and serverless data integration |
| Amazon Aurora | Amazon Aurora DB Cluster<br>Amazon Aurora DB Instance | A MySQL and PostgreSQL-compatible relational database built for the cloud |
| Amazon SageMaker | Amazon SageMaker Endpoint<br>Amazon SageMaker Batch Transform Jobs<br>Amazon SageMaker Training Jobs<br>Amazon SageMaker Ground Truth Labeling Jobs<br>Amazon SageMaker Ground Truth Work Team | Helps data scientists and developers to prepare, build, train, and deploy high-quality machine learning (ML) models quickly by bringing together a broad set of capabilities purpose-built for ML |
| Amazon Personalize | Amazon Personalize DataSet Import Jobs<br>Amazon Personalize Solutions<br>Amazon Personalize Campaigns<br>Amazon Personalize Event Tracker | Allows you to create real-time personalized user experiences faster at scale |
| Amazon Elastic Inference | Amazon Elastic Inference | Allows you to attach low-cost GPU-powered acceleration to Amazon EC2 and Sagemaker instances or Amazon ECS tasks, to reduce the cost of running deep learning inference |
| Amazon API Gateway | Amazon API Gateway | Allows you to create, maintain, and secure APIs at any scale |

| Service | Object | Description |
| --- | --- | --- |
| Amazon Global Accelerator | Amazon Global Accelerator | Allows you to improve global application availability and performance using the AWS global network |
| Amazon FSx for Lustre | Amazon FSx for Lustre | Provides fast and scalable shared storage to power your compute workloads |
| Amazon FSx for Windows File Server | Amazon FSx for Windows File Server | Provides fully managed file storage built on Windows Server |
| AWS App Mesh (App Mesh Metrics for EKS & ECS) | **Note** These services appear under App Mesh Metrics for EKS & ECS. | A service mesh that provides application-level networking to make it easy for your services to communicate with each other across multiple types of compute infrastructure |
| Amazon S3 Glacier (Scale object of S3) | **Note** These services appear under Scale object of S3. | A secure, durable, and extremely low-cost Amazon S3 cloud storage classes for data archiving and long-term backup |
| Amazon Kinesis (Data Analytics, Data Firehose, and Data Stream) | **Note** These services appear under Data Analytics, Data Firehose, and Data Stream. | Allows you to easily collect, process, and analyze video and data streams in real-time |
| AWS Auto Scaling (EC2 Auto Scaling) | **Note** These services appear under EC2 Auto Scaling. | Allows you to add or remove compute capacity to meet changes in demand |

**Note** Services marked with * indicate that the objects have relationships only with regions.

**Note** All services are created with the following Service Descriptors:

- Account ID
- Region
- Service Type

For more information about Amazon Web Services, go to the Amazon Web Services site at https://aws.amazon.com/products.

## Charges for AWS Metrics

Amazon charges you for the metrics you collect. You can reduce costs by selecting only the metrics that are most helpful and filtering out those that are of less interest.

By default, the Management Pack for AWS requests data every five minutes. Every collection cycle makes one CloudWatch call per metric, for an object.

For information about metric costs, see Amazon CloudWatch Pricing.

Based on the costs associated with running the adapter, you can take advantage of some of the features that limit the amount of data you collect from AWS.

- Turn off auto discovery and use manual discovery.
- Go to **Advanced Settings** and select only those services that are critical to your system.
- Subscribe only to specific critical regions.
- Use acceptlist and blocklist filtering to select object import by name.
- Use acceptlist or blocklist to selectively import specific service instances using instance names. Using regex, you can filter services by providing the partial name of the instance.

## Configuring AWS

Configure the AWS in VMware Aria Operations and optionally change its properties to customize the management pack's operation.

An Amazon Web Services account has multiple types of credentials associated with the account. Sign-in credentials are used to access the Amazon Web Services Web-based console, key pairs are used to access EC2 instances, and access keys are used in the REST API that Amazon Web Services exposes.

Because the AWS adapter is based on the REST API, you must use access keys when you set up the adapter. You generate access keys from the Amazon Web Services console. You can create credentials on a per user basis. Access keys are not a username-password pair, but a generated sequence of characters.

**Note** While it is not required, it is recommended that you create a guest type account, which has a read-only access to Amazon Web Services, and use the access keys associated with this account. When you create a guest group with default permissions, they do not include read access to the Elastic Map Reduce (EMR) service. You must use the IAM console to add the following permission:

```
elasticmapreduce:DescribeJobFlows
```

### Generate Required Access Keys

To configure AWS, you must acquire an access key and secret key from the Amazon server. You can acquire these keys as an Amazon Web Services Admin user or as an Amazon Identity and Access Management (IAM) user. For the latest instructions,

#### Prerequisites

- Ensure that you are using Amazon Web Services.
- Ensure that you have the valid permissions and roles in Amazon Web Services.

#### Procedure

1. Log in to Amazon Web Services.

2 To generate access keys, see the online documentation on the https://docs.aws.amazon.com/ site.

   Complete the following tasks:

   - Generate access keys as an Amazon Web Services Administrator.
   - Generate access keys as Amazon Web Services Identity and Access Management User.

## Configuring IAM Permissions

When you set up IAM users and groups, you can stipulate which permissions the account has for API calls. The keys you use when you set up the adapter instance must have certain permissions activated.

For each supported AWS Service, the `ReadOnlyAccess` permission is enough to collect metrics. Use the permission to create a IAM Policy for all supported services and their related services.

To use resource groups tagging API operations, see Resource Groups Tagging API Reference and Services that support the Resource Groups Tagging API.

Log in to the AWS console and create a json similar to the following to get the list of privileges for the service:

```
{
    "Version": "2012-10-17",
    "Statement": [
       {
         "Action": [
          "autoscaling:Describe*",
          "cloudwatch:Describe*",
          "cloudwatch:Get*",
          "cloudwatch:List*",
          "logs:Get*",
          "logs:List*",
          "logs:Describe*",
          "logs:TestMetricFilter",
          "logs:FilterLogEvents",
          "sns:Get*",
          "sns:List*"
        ],
        "Effect": "Allow",
        "Resource": "*"
      }
    ]
}
```

## Table 4-14. IAM Permissions

| Service | Required | Permissions |
| --- | --- | --- |
| Cloudwatch | Yes. | For the list of permissions, see Cloud Watch Read Only Access json. |
| EC2 | describeRegions is required. describeInstances and describeVolumes are only required if you subscribe to the EC2 service. | For more information, see EC2 Read Only Access json. |
| ELB (Elastic Load Balancing) | Required if subscribing to the ELB service. | For the list of permissions, see Elastic Load Balancing Read Only Access json. |
| ELB V2 | Required for application load balancer service. | `{ "Version": "2012-10-17", "Statement": [ { "Sid": "VisualEditor0", "Effect": "Allow", "Action": [ "elasticloadbalancing:DescribeLoadBalancers", "elasticloadbalancing:DescribeListeners", "elasticloadbalancing:DescribeTargetHealth", "elasticloadbalancing:DescribeTargetGroups" ], "Resource": "*" } ] }` |

### Table 4-14. IAM Permissions (continued)

| Service | Required | Permissions |
|---|---|---|
| EMR | Required if subscribing to the EMR service. | describe*<br>`{ "Effect": "Allow", "Action": [ "elasticmapreduce:Describe*", "elasticmapreduce:List*", "elasticmapreduce:ViewEventsFromAllClustersInConsole" "s3:GetObject", "s3:ListAllMyBuckets", "s3:ListBucket", "sdb:Select", "cloudwatch:GetMetricStatistics" ], "Resource": "*" }` |
| RDS | Required if subscribing to RDS service. | For the list of permissions, see RDS Read Only Access json. |
| ElasticCache | Required if subscribing to ElasticCache service. | For the list of permissions, see Elastic Cache Read Only Access json. |
| SQS | Required if subscribing to SQS service. | For the list of permissions, see SQS Read Only Access json. |
| Elastic Container Registry | | For the list of permissions, see Elastic Container Read Only Access json. |
| Elastic Container Service | | list* |
| Lambda | | For the list of permissions, see Lambda Read Only Access json and refer to the AWS Lambda policy. |
| DynamoDB | | For the list of permissions, see Dynamo DB Read Only Access json. |
| DAX | | describe*<br>list* |
| Redshift | | For the list of permissions, see Redshift Read Only Access json. |
| Virtual Private Cloud | | For the list of permissions, see VPC Read Only Access json. |
| Cloud Front Distribution | | For the list of permissions, see Cloud Front Distribution Read Only Access json. |

### Table 4-14. IAM Permissions (continued)

| Service | Required | Permissions |
| --- | --- | --- |
| Direct Connect | | For the list of permissions, see Direct Connect Read Only Access json. |
| VPN Connection | | describe* |
| VPC NAT Gateway | | describe* |
| Elastic IP | | describe* |
| CloudformationStack | | For the list of permissions, see Cloud Formation Read Only Access json. |
| S3 | | For the list of permissions, see S3 Read Only Access json. |
| Workspaces | | describe* |
| Hosted Zone | | list* |
| Health Checks | | list* |
| Neptune DB | | For the list of permissions, see Neptune Read Only Access |
| Personalzie | | list*<br>describe* |
| Sagemaker | | For the list of permissions, see SageMaker Read Only |
| Fsx | | For the list of permissions, see FSx Read Only Access |
| Global Accelerator | | For the list of permissions, see Global Accelerator Read Only Access |
| APIGateway | | get* |
| Elastic Inference | | describe* |
| Glue | | get* |
| DocumentDB | | For the list of permissions, see Doc DB Read Only Access |
| QLDB | | For the list of permissions, see QLDB Read Only |
| Aurora DB | | For the list of permissions, see RDS Read Only Access |
| Storage Gateway | Required if subscribing to the storage gateway service. | listGateways<br>describeGatewayInformation<br>listFileShares<br>describeSMBFileShares<br>describeNFSFileShares<br>listVolumes |

## Update Configuration Settings in the Properties File

The `amazonaws.properties` file provides configuration options. You can optionally change the following configuration properties if required.

### Table 4-15. Amazon Web Services Property Settings

| Property | Description |
| --- | --- |
| firstcollecthistoryhours | Determines how far in the past to collect data when the adapter starts. The default is 0, meaning no historical collection. |
| maxquerywindowminutes | The maximum query window for collections, in minutes. The default is 60. The adapter asks AWS for metrics for a maximum of this many minutes. |
| maxhoursback | The maximum number of hours back from the current time that the adapter attempts to collect. The default value is 336, or two weeks, because Cloudwatch keeps only two weeks worth of metrics. |
| includetransient | False by default. Set to true to allow the adapter to import known transient objects. Transient objects currently include any EMR job that is set to terminate on completion and all of the supporting cluster EC2 instances that belong to that job. |
| threadcount | Default is 4. Controls how many threads are active while making calls to cloudwatch to get metrics. This threadcount is per region. The total number of threads is this value times the number of regions. |
| collecttimeout | Controls how long the adapter waits for all metric collection calls to return from AWS during a collection cycle. The value is measured in seconds. The default value is 240 seconds, which is in line with the default 5 minute collection cycle. |

## Add a Cloud Account for AWS

You can add a AWS cloud account instance to your VMware Aria Operations implementation. VMware Aria Operations collects data from Amazon Web Services (AWS) and supports both commercial and AWS GovCloud accounts. It covers multiple services and regions. You can also add application metrics and other custom metrics from your AWS services and use them in dashboards, alerts and reports.

### Prerequisites

- For key-based authentication, obtain the Access Key and Secret Key values. See Generate Required Access Keys. These values are not the same as your log in credentials for the Amazon Web Services site.
- For role-based authentication, obtain the AWS IAM Role ARN, and External ID. For more information, see KB article 94820.

- Determine the services for which you collect metrics. See, Supported AWS Services. The default value * includes all services in your subscription. If you do not want to subscribe to all services, you can specify region identifiers in the Regions text box.
- Determine the regions to which you subscribe. Amazon Web Services are divided into different regions. The default value * includes all regions in your subscription. If you do not want to subscribe to all regions, you can specify region identifiers in the Regions text box.

Table 4-16. Amazon Web Services Regions for AWS Standard Account

| Region-Friendly Name | Region Identifier |
|---|---|
| US East (N. Virginia) | us-east-1 |
| US East (Ohio) | us-east-2 |
| US West (N. California) | us-west-1 |
| US West (Oregon) | us-west-2 |
| Asia Pacific (Hong Kong) | ap-east-1 |
| Asia Pacific (Tokyo) | ap-northeast-1 |
| Asia Pacific (Seoul) | ap-northeast-2 |
| Asia Pacific (Osaka-Local) | ap-northeast-3 |
| Asia Pacific (Mumbai) | ap-south-1 |
| Asia Pacific (Singapore) | ap-southeast-1 |
| Asia Pacific (Sydney) | ap-southeast-2 |
| Asia Pacific (Jakarta) | ap-southeast-3 |
| Canada (Central) | ca-central-1 |
| China (Beijing) | cn-north-1 |
| China (Ningxia) | cn-northwest-1 |
| Europe (Frankfurt) | eu-central-1 |
| Europe (Zurich) | eu-central-2 |
| Europe (Ireland) | eu-west-1 |
| Europe (London) | eu-west-2 |
| Europe (Paris) | eu-west-3 |
| Europe (Stockholm) | eu-north-1 |
| Europe (Milan) | eu-south-1 |
| Europe (Spain) | eu-south-2 |

Table 4-16. Amazon Web Services Regions for AWS Standard Account (continued)

| Region-Friendly Name | Region Identifier |
|---|---|
| South America (São Paulo) | sa-east-1 |
| Africa (Cape Town) | af-south-1 |
| Middle East (Bahrain) | me-south-1 |
| Middle East (UAE) | me-central-1 |

Table 4-17. Amazon Web Services Regions for AWS GovCloud (US) Account

| Region-Friendly Name | Region Identifier |
|---|---|
| AWS GovCloud (US-East) | us-gov-east-1 |
| AWS GovCloud (US) | us-gov-west-1 |

- Determine any blocked list or allowed list filters. These filters use regular expressions to filter in or out specific objects by name. For example, an allowed list filter of `.*indows.*` allows only objects with a name including "indows". A blocked list filter of `.*indows.*` filters out all objects with that string in their name.
- To publish custom metrics to CloudWatch, see Publishing Custom Metrics.
- To collect additional metrics for EC2 through CloudWatch agent, set up the agent. For details, see Installing CloudWatch Agent.

Procedure

1. In the left menu, click **Data Sources > Integrations**.
2. On the Accounts tab, click **ADD**.
3. On the Account Types page, click **AWS**.
4. Configure the instance settings.

| Option | Action |
|---|---|
| **Name** | Enter a name for the adapter instance. |
| **Description** | Enter a description. |
| **Account Type** | Select the AWS account type.<br>Select **AWS Standard Account** to set up a commercial AWS account.<br>Select **AWS GovCloud (US) Account** to set up a AWS GovCloud (US) account to monitor GovCloud services and regions. |

| Option | Action |
| --- | --- |
| **Services/ GovCloud Services** | Select the services from which you want to capture metrics. If you want to collect metrics for specific services, then click the drop-down icon and select one or more services. For example, **`Amazon CloudFormation, Amazon EC2`**. If you do not select any of the services, the metrics for all the services get collected.<br>The services marked with an asterix* for example, **`AWS AppSync*`** are grouped together under AWS Other Services. These services display the relationship with the regions only. For more information on supported AWS services, see Supported AWS Services.<br>**Note** If you select **AWS GovCloud (US) Account** as the account type, you can only select the services applicable for AWS GovCloud (US). |
| **Regions/ GovCloud Regions** | Select the regions you want to subscribe to. If you want to subscribe to specific regions, then click the drop-down icon and select one or more regions. For example, `US East (N. Virginia),US East (Ohio)`. If you want to subscribe to all the regions, do not select any of the regions.<br>**Note** If you select **AWS GovCloud (US) Account** as the account type, you can only select the regions applicable for AWS GovCloud (US). The regions for AWS GovCloud (US) are AWS GovCloud (US-East) and AWS GovCloud (US). |
| **Credential** | Add the credentials used to access the AWS environment by clicking the plus sign and select the **Credential Kind**.<br>**Note** For each credential kind, you must enter a **Credential name**. This name is not the name of the adapter instance, but a friendly name to identify the credential with ease.<br>■ **Instance Profile Authentication**: This authentication must be used by VMware SRE users only.<br>■ **Role-based authentication**: Enter the **Credential name**, **AWS IAM Role ARN**, and the **External ID**. For more information, see KB article 94820.<br>■ **Key-based authentication**: Use the AWS key-based authentication an application (running in an AWS Account) using access key and secret key.<br>**Note** If you are configuring an AWS GovCloud (US) Account, then the access key and secret key values must be specific to AWS GovCloud.<br>Enter the **Credential name** and the **Access Key** and **Secret Key** values.<br>Optionally, enter any required local proxy information for your network.<br>**Note** This proxy information activates VMware Aria Operations that is deployed on a local network to have a public network access for AWS to collect statistics for its accounts. |
| **Collector / Group** | Select the collector upon which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default has been selected for optimal data collecting. |

5 Click **Test Connection** to validate the connection.

6 Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

| Option | Action |
| --- | --- |
| **Collect Custom Metrics** | Set this option to true if you want to import all the custom metrics from your AWS account.<br>To publish custom metrics in VMware Aria Operations, the metrics dimension names should match the following service mappings: |

| Service Name | Dimension Name |
| --- | --- |
| dax_cluster | ClusterId |
| dax_node | NodeId |
| dynamodb | TableName |
| efs | FileSystemId |
| eks | ClusterName |
| elasticbeanstalk_env | EnvironmentName |
| redshift_node | NodeID |
| redshift_cluster | ClusterIdentifier |
| s3_bucket | BucketName |
| vpc_nat_gateway | NatGatewayId |
| vpc_vpn | VpnId |
| workspace | WorkspaceId |
| ec2_auto_scale_group | AutoScalingGroupName |
| cloudfront_distribution | DistributionId |
| direct_connect | ConnectionId |
| ec2_instance | InstanceId |
| ec2_volume | VolumeId |
| transit_gateway | TransitGateway |
| ecs_cluster | ClusterName |
| ecs_service | ServiceName |
| elasticache_cachecluster | CacheClusterId |
| elasticache_cachenode | CacheNodeId |
| ec2_load_balancer | LoadBalancerName |
| application_load_balancer | LoadBalancer |

| Option | Action |
|---|---|

| Service Name | Dimension Name |
|---|---|
| network_load_balancer | LoadBalancer |
| emr_job_flow | JobFlowId |
| lambda_function | FunctionName |
| rds_dbinstance | DBInstanceIdentifier |
| hosted_zone | HostedZoneId |
| health_check | HealthCheckId |
| sqs_queue | QueueName |
| amazon_neptune_db_instance | DBInstanceIdentifier |
| amazon_neptune_db_cluster | DBClusterIdentifier |
| amazon_personalize_data_import | DatasetimportjobArn |
| amazon_personalize_event_tracker | EventTrackerArn |
| amazon_personalize_solution | SolutionArn |
| amazon_personalize_campaign | CampaignArn |
| amazon_sagemaker_endpoint | EndpointName |
| amazon_sagemaker_batch_transform_job | Host |
| amazon_sagemaker_ground_truth_labeling_job | LabelingJobName |
| amazon_sagemaker_ground_truth_work_team | Workteam |
| amazon_global_accelerator | Accelerator |
| amazon_api_gw | ApiName |
| amazon_elastic_inference | ElasticInferenceAcceleratorId |
| amazon_glue_job | JobName |
| amazon_qldb_ledger | LedgerName |
| amazon_qldb_stream | StreamId |

| | |
|---|---|
| **Support Auto Discovery** | Set this option to true for automatic discovery of AWS services. If you set this value to false, when you create an adapter instance you must perform a manual discovery of services. |
| **Allowed List Regex** | Add regular expressions to allow only objects with names that fit the criteria you specify. |

| Option | Action |
| --- | --- |
| **Blocked List Regex** | Add regular expressions to filter out objects by name. |
| **Actions Activated** | Activate this option to run actions specific to the AWS virtual machine using VMware Aria Operations. By default, this option is set to **true**.<br>For details on running actions, see Power On, Power Off, and Reboot Actions. |
| **Account ID** | The AWS account ID is auto populated. Each AWS account uses a unique account ID to run actions. The account ID is the same one used to log in Amazon Web Services site. |

7 Click **Save Settings**.

### What to do next

Make sure that VMware Aria Operations is collecting data.

| Where to View the Information | Information to View |
| --- | --- |
| Collection Status and Collection State columns in the MP for AWS Solution Details pane on the **Accounts** tab on the **Integrations** page. | The collection status appears approximately 10 minutes after you have configured the adapter. |
| Environment Overview | The objects related to AWS are added to the inventory trees. |
| Dashboards | AWS dashboards are added to VMware Aria Operations. |

### Tagging Groups

AWS uses tagging groups. The tagging groups appear under the AWS Entity Status in the Inventory page.

Table 4-18. Tagging Groups

| Group Name | Description |
| --- | --- |
| PoweredOn | Objects with this tag are in the running state. |
| PoweredOff | Objects with this tag are in the stopped state. |
| Transient | Objects with this tag are not expected to persist for long periods of time. |
| NotExisting | Objects with this tag do not exist in the Amazon Web Services system. You can use this tag to take advantage of the periodic purge feature of VMware Aria Operations, that the `controller.properties` file on the Analytics server controls. |

## View AWS Objects

You can use the inventory tree to browse and select objects. The inventory tree shows a hierarchical arrangement of the AWS objects by region.

Procedure

1 In the left pane of VMware Aria Operations, click the **Environment** icon.
   - Click **Accounts** to list the AWS adapter instances.
   - Click **Regions** to to list the AWS regions.
2 To view the child objects, expand the regions and then expand the regions per account.

   **Note** All the account-specific objects related to a region are grouped under the region per account section.

3 To display information about the object, select an object in the inventory tree.

# Microsoft Azure

Microsoft Azure is an embedded adapter with diagnostic dashboards for VMware Aria Operations. The adapter collects metrics from Microsoft Azure.

## Supported Azure Services

Microsoft Azure supports the following services.

| Service | Object | Description |
|---|---|---|
| Azure App Service | Azure Bot Service | Allows you to quickly create powerful cloud apps for web and mobile |
| Azure Virtual Machines | Azure Linux Virtual Machines<br>Azure CycleCloud | Provides a provision Windows and Linux virtual machines in seconds |
| Azure Kubernetes Service | Azure Kubernetes Cluster | Allows you to simplify the deployment, management and operations of Kubernetes |
| Azure Virtual Machine Scale Sets | Azure Virtual Scale Set Instance | Allows you to manage and scale up to thousands of Linux and Windows virtual machines |
| Azure Cosmos DB | Azure Cosmos DB | Provides a fast NoSQL database with open APIs for any scale |
| Azure MySQL Server | Azure MySQL Server | Provides fully managed, scalable MySQL Database |
| Azure PostgreSQL Server | Azure PostgreSQL Server | Provides fully managed, intelligent and scalable PostgreSQL |
| Azure SQL Database | Azure SQL Database | Provides managed, intelligent SQL in the cloud |
| Azure Application Gateway | Azure Application Gateway | Allows you to build secure, scalable and highly available web front ends in Azure |

| Service | Object | Description |
| --- | --- | --- |
| Azure Load Balancer | Azure Load Balancer | Helps you deliver high availability and network performance to your applications |
| Azure Virtual Network | Azure Virtual Network | Provision private networks, optionally connect to on-premises datacenters |
| Azure Virtual Network Gateway | Azure Virtual Network Gateway | Helps you establish secure, cross-premises connectivity |
| Azure Disk | Azure Disk | Provides a high-performance, highly durable block storage for Azure Virtual Machines |
| Azure Storage Account | Table storage<br>Azure Storage File<br>■ Azure Storage File Share<br>■ Azure Storage File Share Directory<br>Azure Storage Containers<br>■ Azure Storage Blob<br>■ Azure Storage Blob Container<br>Azure Storage Queue | Provides durable, highly available and massively scalable cloud storage |
| Azure Event Hub | Azure Event Hub | Provides simple, secure and scalable real-time data ingestion |
| Azure Dedicated Host | Azure Dedicated Host | Provides a dedicated physical server to host Azure VMs for Windows and Linux |
| Azure Function | Azure Function | Allows you to process events with serverless code |
| Azure Batch Account | Azure Batch Account | Provides cloud-scale job scheduling and compute management |
| Azure Cloud Services (Classic) | Azure Cloud Services (Classic) | Provides platform as a service (PaaS) technology engineered to deploy web and cloud applications that are scalable, reliable, and inexpensive to operate |
| Azure OpenShift Cluster | Azure OpenShift Cluster | Provides fully managed OpenShift service |
| Azure Container Instances | Azure Container Instances | Allows you to easily run containers on Azure without managing servers |
| Azure Container Registry | Azure Container Registry | Allows you to store and manage container images across all types of Azure deployments |
| Azure Web Apps | Azure Web Apps | Allows you to quickly create and deploy mission critical web apps at scale |

| Service | Object | Description |
| --- | --- | --- |
| Azure Cache for Redis | Azure Cache for Redis | Allows you to accelerate applications with high-throughput, low-latency data caching |
| Azure MariaDB DataBase | Azure Database for MariaDB Server | A managed MariaDB database service for app developers |
| Azure SQL Managed Instance | Azure SQL Managed DataBase | Allows you to modernise your existing applications at scale and do more with a managed, secure and always up-to-date SQL instance in the cloud |
| Azure App Configuration | Azure App Configuration | Provides fast, scalable parameter storage for app configuration |
| Azure Service Bus | Azure Service Bus | Allows you to connect across private and public cloud environments |
| Azure Media Service | Azure Media Streaming EndPoint | Allows you to encode, store, and stream video and audio at scale |
| Azure Notification Hub | Azure Notification Hub | Allows you to send push notifications to any platform from any back end |
| Azure DNS | Azure DNS Zone<br>Azure Private DNS Zone | Allows you to host your DNS domain in Azure |
| Azure ExpressRoute Circuit | Azure ExpressRoute Circuit | Provides dedicated private network fiber connections to Azure |
| Azure Firewall | Azure Firewall | A native firewalling capabilities with built-in high availability, unrestricted cloud scalability and zero maintenance |
| Azure Front Door | Azure Front Door | Provides scalable, security-enhanced delivery point for global, microservice-based web applications |
| Azure Content Delivery Network (CDN) | Azure CDN Profile<br>Azure CDN Profile EndPoint | Allows you to reduce load times, save bandwidth and speed responsiveness |
| Azure Network Watcher | Azure Network Watcher | Provides network performance monitoring and diagnostics solution |
| Azure Traffic Manager Profile | Azure Traffic Manager Profile | Allows you to route incoming traffic for high performance and availability |
| Azure Virtual WAN | Azure Virtual WAN | Allows you to optimise and automate branch to branch connectivity through Azure |
| Azure Key Vault | Azure Key Vault | Allows you to safeguard and maintain control of keys and other secrets |
| Azure Data Lake Storage Gen1 | Azure Data Lake Storage Gen1 | Provides massively scalable, secure data lake functionality built on Azure Blob Storage |

| Service | Object | Description |
| --- | --- | --- |
| Azure NetApp Account | Azure NetApp Account Volume<br>Azure NetApp Account CapacityPool | Provides enterprise-grade Azure file shares, powered by NetApp |
| Azure SignalR | Azure SignalR | Provides add real-time web functionalities easily |
| Azure Cognitive Search | Azure Cognitive Search | AI-powered cloud search service for mobile and web app development |
| Azure Machine Learning | Azure Machine Learning | Helps you bring AI to everyone with an end-to-end, scalable, trusted platform with experimentation and model management |
| Azure Cognitive Service Account | Translator | Allows you to add smart API capabilities to activate contextual interactions |
| Azure Analysis Service | Azure Analysis Service | Provides an enterprise-grade analytics engine as a service |
| Azure Data Explorer Cluster | Azure Data Explorer Cluster | Provides fast and highly scalable data exploration service |
| Azure Stream Analytics | Azure Stream Analytics Job<br>Azure Stream Analytics Cluster | Provides real-time analytics on fast moving streams of data from applications and devices |
| Azure Synapse Analytics | Azure Synapse Analytics Workspace<br>Azure Synapse Analytics SQL Pool<br>Azure Synapse Analytics Spark Pool | Provides limitless analytics service with unmatched time to insight |
| Azure Data Factory | Azure Data Factory | Provides hybrid data integration at enterprise scale, made easy |
| Azure Data Lake Analytics | Azure Data Lake Analytics | Provides distributed analytics service which makes big data easy |
| Azure HDInsight | Azure HDInsight | Provides enterprise-ready, managed cluster service for open-source analytics |
| Azure Power BI Embedded | Azure Power BI Embedded | Allows you to add analytics and interactive reporting to your applications. |
| Azure Purview Account | Azure Purview Account | Allows you to maximize business value with unified data governance |
| Azure Data Box | Azure Data Box | Provides appliances and solutions for offline data transfer to Azure |
| Azure API Management | Azure API Management | Allows you to publish APIs to developers, partners, and employees securely and at scale |
| Azure Digital Twins | Azure Digital Twins | Allows you to build next-generation IoT spatial intelligence solutions |

| Service | Object | Description |
| --- | --- | --- |
| Azure DDos Protection Plan | Azure DDos Protection Plan | Protects your applications from Distributed Denial of Service (DDoS) attacks |
| Azure IoT Central | Azure IoT Central | Accelerates the creation of IoT solutions |
| Azure Spatial Anchors | Azure Spatial Anchors | Allows you to create multi-user, spatially aware mixed reality experiences |
| Azure Event Hub | Azure Event Hub | Provides a simple, secure and scalable real-time data ingestion |
| Azure IoT Hub | Azure IoT Hub | Allows you to connect, monitor and manage billions of IoT assets |
| Azure Automation | Azure Automation | Allows you to simplify cloud management with process automation |
| Azure Time Series Insights | Azure Time Series Insights | Allows you to explore and analyse time-series data from IoT devices |
| Azure Event Grid | ■ Azure Event Grid Domains<br>■ Azure Event Grid Topics<br>■ Azure Event Grid Subscriptions | Provides reliable event delivery at massive scale |

## Charges for Azure Metrics

Microsoft Azure charges you for the metrics you collect. You can reduce costs by selecting only the metrics that are most helpful and filtering out those that are of less interest.

By default, the Microsoft Azure requests data every five minutes. Every collection cycle makes one Azure Monitor call per metric, for an object.

For information about metric costs, see Azure Monitor Pricing.

Based on the costs associated with running the adapter, you can take advantage of some of the features that limit the amount of data you collect from Microsoft Azure.

- Turn off auto discovery and use manual discovery.
- Go to **Advanced Settings** and select only those services that are critical to your system.
- Subscribe only to specific critical regions.
- Use acceptlist and blocklist filtering to select object import by name.
- Use acceptlist or blocklist to selectively import specific service instances using instance names. Using regex, you can filter services by providing the partial name of the instance.

# Configuring Microsoft Azure

To configure Microsoft Azure, you must activate it in VMware Aria Operations and optionally change properties to customize it.

Microsoft Azure is a native management pack. You must activate the management pack if it is deactivated. For more information, see The Repository Tab.

After activating the management pack, you must create an application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure the management pack in VMware Aria Operations.

**Note**

- You can install and use the management pack only with an enterprise license of VMware Aria Operations.
- The management pack has a default time granularity based on the services that it monitors. You cannot configure this granularity against the metrics. You can increase the collection interval but you must not decrease it. The default interval is 10 minutes.

## Generate a Client Secret

Create an Active Directory application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure a cloud account for Microsoft Azure.

### Prerequisites

- Ensure that you are using Microsoft Azure Cloud.
- Ensure that you have a valid subscription in the Microsoft Azure portal with an Active Directory integration.

### Procedure

1. Log in to the Microsoft Azure portal.
2. Create an application and generate a secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.

   Complete the following tasks:

   a. Create an Azure Active Directory application.

      **Note** Ensure that the API Permission is 'Microsoft Graph User.Read'.

   b. Navigate to **Subscriptions** and select your subscription.

c In the left pane, click **Access Control (IAM)** and then click, **Add > Add Role Assignment**. Select the role you want to assign to the application. The minimum requirement is 'Reader' or above.

**Note** To provide access to a specific resource(s), create a Resource Group for the resource(s) and give access at the resource group level.

d Click **Select members**, and in the right pane, search for and add one or more members you want to assign to the role for the resource.

e Click **Review + assign**.

f Generate a client secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.

g Copy the subscription ID, directory (tenant) ID, application (client) ID, and client secret to use in your cloud account.

## Add a Cloud Account for Microsoft Azure

Microsoft Azure is an embedded adapter in which each adapter instance has diagnostic dashboards and collects metrics from Microsoft Azure. You can add a cloud account to configure an adapter instance in VMware Aria Operations. VMware Aria Operations collects data from Microsoft Azure and supports both standard/commercial and Azure GovCloud subscriptions. It covers multiple services and regions. You can also add application metrics and other custom metrics from your Azure services and use them in dashboards, alerts and reports.

### Prerequisites

- If Microsoft Azure is deactivated, activate it in VMware Aria Operations. For more information, see the The Repository Tab.
- Generate a client secret in the Microsoft Azure portal that you can use in this configuration. For more information, see Generate a Client Secret.

Table 4-19. Microsoft Azure Regions for Microsoft Azure Standard/Commercial Account

| Region-Friendly Name | Region Identifier |
|---|---|
| Australia Central | australiacentral |
| Australia Central 2 | australiacentral2 |
| Australia East | australiaeast |
| Australia Southeast | australiasoutheast |
| Brazil South | brazilsouth |
| Brazil Southeast | brazilsoutheast |
| Canada Central | canadacentral |
| Canada East | canadaeast |

### Table 4-19. Microsoft Azure Regions for Microsoft Azure Standard/Commercial Account (continued)

| Region-Friendly Name | Region Identifier |
|---|---|
| Central India | centralindia |
| Central US | centralus |
| Central US EUAP | centraluseuap |
| East Asia | eastasia |
| East US | eastus |
| East US 2 | eastus2 |
| East US 2 EUAP | eastus2euap |
| France Central | francecentral |
| France South | francesouth |
| Germany North | germanynorth |
| Germany West Central | germanywestcentral |
| Japan East | japaneast |
| Japan West | japanwest |
| Jio India Central | jioindiacentral |
| Jio India West | jioindiawest |
| Korea Central | koreacentral |
| Korea South | koreasouth |
| North Central US | northcentralus |
| North Europe | northeurope |
| Norway East | norwayeast |
| Norway West | norwaywest |
| South Africa North | southafricanorth |
| South Africa West | southafricawest |
| South Central India | southcentralindia |
| South Central US | southcentralus |
| South East Asia | southeastasia |
| South India | southindia |
| Sweden Central | swedencentral |

Table 4-19. Microsoft Azure Regions for Microsoft Azure Standard/Commercial Account (continued)

| Region-Friendly Name | Region Identifier |
| --- | --- |
| Sweden South | swedensouth |
| Switzerland North | switzerlandnorth |
| Switzerland West | switzerlandwest |
| UAE Central | uaecentral |
| UAE North | uaenorth |
| UK South | uksouth |
| UK West | ukwest |
| West Central US | westcentralus |
| West Europe | westeurope |
| West India | westindia |
| West US | westus |
| West US 2 | westus2 |
| West US 3 | westus3 |

Table 4-20. Microsoft Azure Regions for Azure GovCloud (US) Account

| Region-Friendly Name | Region Identifier |
| --- | --- |
| Azure GovCloud (US DoD Central) | usdodcentral |
| Azure GovCloud (US DoD East) | usdodeast |
| Azure GovCloud (US Gov Arizona) | usgovarizona |
| Azure GovCloud (US Gov Texas) | usgovtexas |
| Azure GovCloud (US Gov Virginia) | usgovvirginia |

Procedure

1 From the left menu, click **Data Sources > Integrations**.

2 On the Accounts tab, click **Add**, and on the Account Types page, select **Microsoft Azure**.

3 Enter the cloud account information.

| Option | Action |
| --- | --- |
| **Name** | Enter a name for the adapter instance. |
| **Description** | Enter a description for the adapter instance. |

4 Configure the connection.

| Option | Action |
| --- | --- |
| **Account Type** | Select the Azure account type.<br>Select **Azure Standard Account** to set up a commercial Azure account.<br>Select **Azure GovCloud (US) Account** to set up a Azure GovCloud (US) account to monitor GovCloud services and regions. |
| **Services/ GovCloud Services** | Select the services from which you want to capture metrics. If you want to collect metrics for specific services, then click the drop-down icon and select one or more services. For example, **`Azure Analysis Service`**, **`Azure Application Gateway`**. If you do not select any of the services, the metrics for all the services get collected. For more information on supported Microsoft Azure services, see Supported Azure Services.<br>**Note** If you select **Azure GovCloud (US) Account** as the account type, you can only select the services applicable for Azure GovCloud (US). |
| **Regions/ GovCloud Regions** | Select the regions you want to subscribe to. If you want to subscribe to specific regions, then click the drop-down icon and select one or more regions. For example, `Central US`, `East Asia`. If you want to subscribe to all the regions, do not select any of the regions.<br>**Note** If you select **Azure GovCloud (US) Account** as the account type, you can only select the regions applicable for Azure GovCloud (US). |
| **Subscription ID** | Enter your subscription ID for Microsoft Azure. |
| **Directory (Tenant) ID** | Enter the directory (tenant) ID for your Azure Active Directory. |
| **Credential** | Add the credentials used to access Microsoft Azure by clicking the plus sign.<br>■ Enter a name for the credential values you are creating. This value is not the name of the adapter instance, but a friendly name for the secret credential.<br>■ Enter your application (client) ID in your Azure Active Directory.<br>■ Enter the client secret that you generated for your application in the Microsoft Azure portal.<br>■ (Optional) Enter any required local proxy information for your network.<br>**Note** This proxy information activates VMware Aria Operations that is deployed on a local network to have public network access for Azure to collect statistics for its accounts. |
| **Collector/Group** | Select the collector on which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default is selected for optimal data collection. |

5 Click **Validate Connection** to test the connection.

**Note** If the test connection fails, do not add the cloud account.

If you add the cloud account with a failed test connection, VMware Aria Operations might not collect data for the adapter instance. To resolve this issue, remove the cloud account and add it again with the correct information. If you are using a proxy, ensure that the proxy connection is efficient.

6 Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

| Option | Action |
| --- | --- |
| **Collect Custom Metrics** | Set this option to **true** if you want to import all the custom metrics from your Azure account. |
| **Actions Enabled** | Activate this option to run actions specific to Azure Virtual Machine using VMware Aria Operations. By default, this option is set to **true**.<br>For details on running actions, see Power On, Power Off, and Reboot Actions.<br>**Note** If you have activated actions, it is mandatory to add the **Application (Object) ID** in the **Manage Credential** dialog box. |

7 Click **Add**.

### What to do next

Ensure that VMware Aria Operations is collecting data.

| Where to View the Information | Information to View |
| --- | --- |
| **Environment** | The objects related to the adapter instance are added to the inventory. For more information, see View Objects for Microsoft Azure. |
| **Dashboards** | The dashboards for the adapter instance are added to VMware Aria Operations. For more information, see Microsoft Azure Dashboards. |

## View Objects for Microsoft Azure

You can use the inventory tree in VMware Aria Operations to browse and select objects for an adapter instance of Microsoft Azure. The inventory tree shows a hierarchical arrangement of the objects by cloud account and by region.

### Prerequisites

Configure an adapter instance of Microsoft Azure. For more information, see Add a Cloud Account for Microsoft Azure.

**Note** When you monitor large-scaled Azure end-points (>1000 objects), change the default collection cycle to 15 minutes so that there is enough time to collect data for all the objects from a scaled end-point.

### Procedure

1. From the left menu, click **Environment > Object Browser**.
2. In the left pane, expand **Microsoft Azure** and select either of the following options:
   - To view the objects by region, click **Azure Resources By Region**.
   - To view the objects by cloud account, click **Azure Resources By Subscription**.
3. To view the object information by region, region per cloud account, subregion, cloud account, or resource group, select either of the following options:
   - If you are viewing objects by region, select a region. You can click the **Azure Region per Subscription** tab to view the object information for the region per cloud account. You can also expand the inventory tree for each region and select a subregion.
   - If you are viewing objects by cloud account, select a cloud account. You can also expand the inventory tree for each cloud account and select a resource group.
4. To view information about each object, select either of the following options:
   - If you are viewing objects by region, expand the inventory tree for a subregion and select an object.
   - If you are viewing objects by cloud account, select an object under a cloud account or expand the inventory tree for a resource group and select an object.

   You can expand the inventory tree for an SQL Server object and select an SQL Database object to view information about the database object.

# OS and Application Monitoring

You can monitor application services in VMware Aria Operations. You can use the product-managed agent to monitor physical servers and VMs that are managed or unmanaged by the vCenter Server cloud account (target machine).

For example, as an administrator, you might need to ensure that the infrastructure provided for running the application services is sufficient and that there are no problems. If you receive a complaint that a particular application service is not working properly or is slow, you can troubleshoot by looking at the infrastructure on which the application is deployed. You can view

important metrics related to the applications and share the information with the team managing the applications. You can use VMware Aria Operations to deploy the agents and send the related application data to VMware Aria Operations. You can view the data in VMware Aria Operations and share it with the team so that they can troubleshoot the application service.

You can conduct remote checks, monitor Windows services and Linux processes, monitor operating systems and applications, and run custom scripts in VMware Aria Operations.

## Product-Managed Telegraf on Different Types of Machines

The following (object types) are now supported for application monitoring based on the type of machine the product-managed Telegraf agent is running on.

- `Virtual Machine`: Refers to the vCenter Server. All the capabilities of product-managed Telegraf agents and monitoring of the vCenter Server VMs remain the same from previous releases. For more information, see Install and Uninstall a Telegraf Agent and Additional Operations from the Manage Telegraf Agents Page .
- `Endpoint`: VMware Aria Operations additionally supports the following types of machines that are categorized under an object type called `Endpoint`:
  - Physical servers
  - Machines that are not directly controlled by VMware Aria Operations. For example, VMs that are not managed by the vCenter Server cloud account.

    To install and uninstall a product-managed Telegraf agent on a remote machine or a physical server, use the helper script. For more information, see Install/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform . To manage the life cycle of agents on a remote machine or a physical server, use the VMware Aria Operations UI or suite-api. For more information, see Additional Operations from the Manage Telegraf Agents Page .

  In the application monitoring documentation for product-managed Telegraf agents, the term **target machine/s** refers to physical servers and the VMs that are managed and unmanaged by the vCenter Server cloud account.

## High Availability for Application Monitoring

You can now use application monitoring high availability activated collector groups in failover mode to monitor applications without loosing any data (metrics) from the application that is monitored even if your cloud proxy is down or if any of the cloud proxy components fail. Here are some important points:

- Cloud proxies in an application monitoring high availability activated collector group work in a primary-secondary mode.
- You cannot convert an existing Telegraf agent on all target machines, to high availability. To convert an existing target machine to high availability, you must install or re-install the agent with an application monitoring high availability activated collector group.

- You will lose all application data when you add a cloud proxy used by Telegraf agents, to the application monitoring high availability activated collector group.
- Failover or failback has a downtime of three collection cycles.
- Addition or removal of a cloud proxy from the application monitoring high availability activated collector has a downtime of three collection cycles.

## Steps to Activate Application Monitoring High Availability on Collector Groups

Here are the steps to activate application monitoring high availability on collector groups and use it for application monitoring and data collection:

1 Activate high availability for specific cloud proxies when you add a collector group. For more information, see Adding a Collector Group.

2 (Optional) You can view the cloud proxies that are configured for high availability from the **Cloud Proxies** page. For more information, see Monitoring the Health of Cloud Proxies.

3 While installing an agent, select a collector group that has application monitoring high availability activated. For more information, see Install an Agent from the UI.

4 (Optional) You can view the following agent details such as:

   a If high availability is activated for the collector group (Collector Group column).

   b The virtual IP configured on an application monitoring high availability activated collector group to which all target machines send data (Virtual IP details column).

   c The IP address of the primary cloud proxy that is collecting data (Cloud Proxy column).

   You can view these details from the **Manage Telegraf Agents** page (**Environment > Applications > Manage Telegraf Agents**). For more information, see Additional Operations from the Manage Telegraf Agents Page .

## Supported Application Services

OS and Application monitoring in helps virtual infrastructure administrators and application administrators to discover operating systems and applications running in provisioned guest operating systems at a scale and to collect run-time metrics of the operating system and application for monitoring and troubleshooting respective entities.

The following 23 application services are supported.

| Application Services |
| --- |
| Active Directory |
| Active MQ |
| Apache HTTPD |
| Cassandra |

| Application Services |
| --- |
| HyperV |
| Java Application |
| JBoss Server |
| MongoDB |
| MS Exchange |
| Microsoft IIS |
| Microsoft SQL Server |
| MySQL |
| Network Time Protocol |
| Nginx |
| Oracle DB |
| Pivotal TC Server |
| PostgreSQL |
| RabbitMQ |
| Riak KV |
| SharePoint Server |
| Tomcat Server |
| Oracle WebLogic Server |
| WebSphere |

## VeloCloud Application Services

Eight VeloCloud application services are supported in VMware Aria Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected. The following services are supported in VMware Aria Operations:

| Application Services |
| --- |
| VeloCloud Orchestrator |
| Nginx |
| Clickhouse |

| Application Services |
| --- |
| Network Time Protocol |
| MySQL |
| Redis |
| Java Application |
| VeloCloud Gateway |

## Supported Platforms

VMware Aria Operations supports monitoring for the following platforms and app combinations with API support.

### Platforms Supported by VMware Aria Operations for OS and Application Monitoring

| Platform | Version | Architecture | Application |
| --- | --- | --- | --- |
| Red Hat Enterprise Linux | 7.x<br>8.x<br>9.x | 64-bit | OS Metrics and all supported applications. |
| CentOS | 7.x<br>8.x | 64-bit | OS Metrics and all supported applications. |
| Oracle Linux | 7.x<br>8.x<br>9.x | 64-bit | OS Metrics and all supported applications. |
| Rocky Linux | 8.x<br>9.x | 64-bit | OS Metrics and all supported applications. |
| Ubuntu | 18.04 LTS<br>16.04 LTS<br>20.x | 64-bit | OS Metrics and all supported applications. |
| VMware Photon Linux | 1.0<br>2.0<br>3.0 | 64-bit | Only OS metrics and custom monitoring actions are supported for photon OS.<br>Site Recovery Manager 8.2 runs on Photon 2.0<br>vSphere- vSphere 6.7 & 6.5 runs on Photon OS 1.0<br>VMware vSAN 6.7 & VMware vSAN 6.5 runs on Photon OS 1.0<br>Unified Access Gateway 3.7 runs on Photon 3.0 & 3.6 runs on Photon 2.0. |

| Platform | Version | Architecture | Application |
|---|---|---|---|
| SUSE Linux Enterprise Server | 12.x<br>15.x | 64-bit | OS Metrics and all supported applications. |
| Windows | Windows Server 2019<br>Windows Server 2016<br>Windows 2012<br>Windows Server 2012 R2<br>Windows Server 2022 | 64-bit | OS Metrics and all supported applications. |

## Prerequisites

To monitor your application services and operating systems, complete all the prerequisites so that cloud proxy can communicate successfully with the target machines.

### Port Information and Communication with the Target Machines

Port information and communication details with the target machines to install the Telegraf agent from the UI or via the helper script is detailed on this page.

Ensure that cloud proxy's ports are reachable from target machines both via the IP and FQDN.

For the latest port information, see VMware Ports and Protocols.

Figure 4-1. Port Information and Communication with the vCenter Server, VMs, VMware Aria Operations, and Cloud Proxy (Agent Install from the UI)

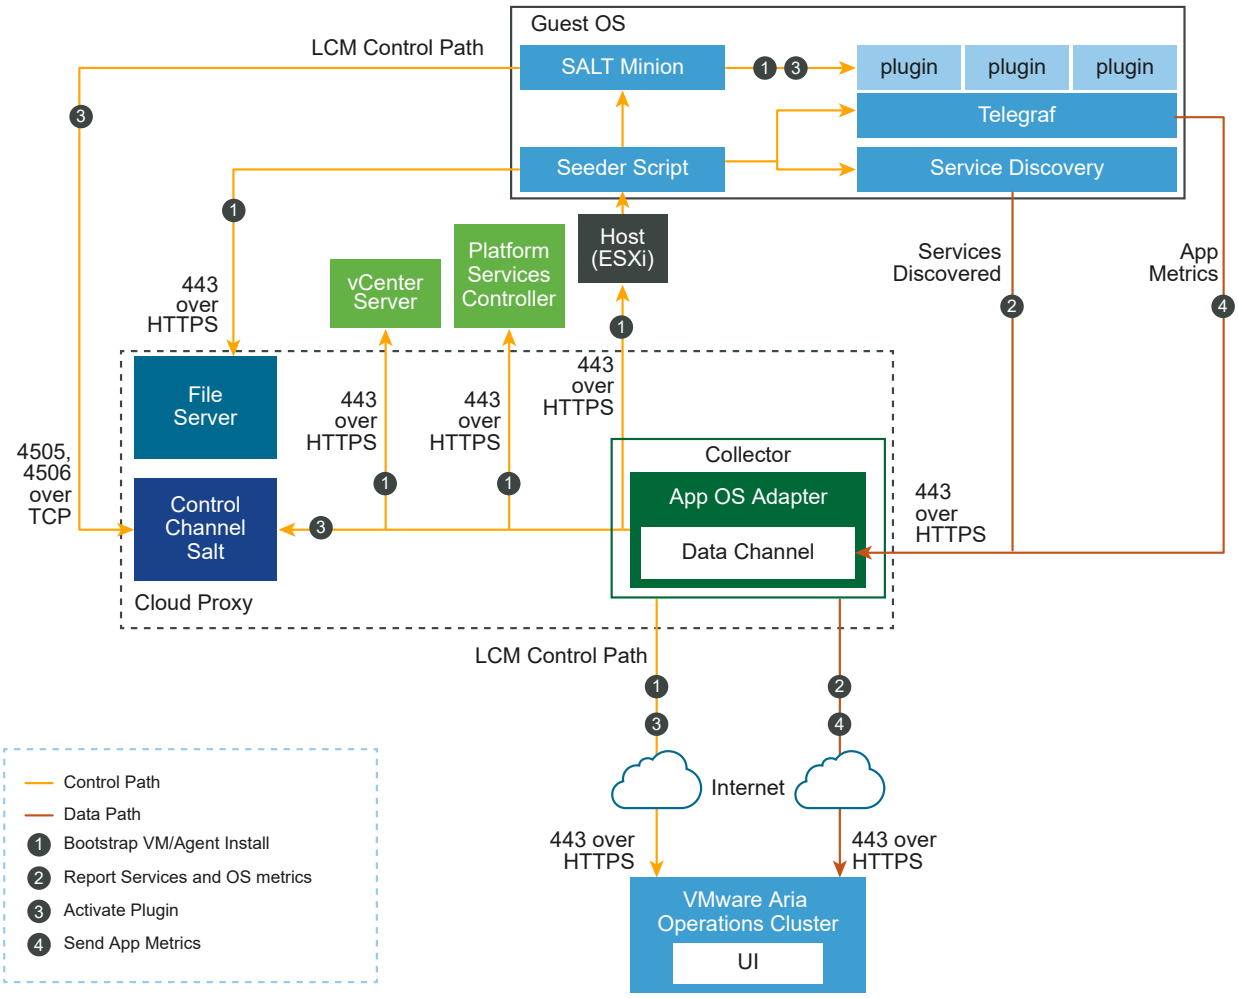

Figure 4-2. Port Information and Communication with the Target Machines and VMs for Script-Based Agent Install

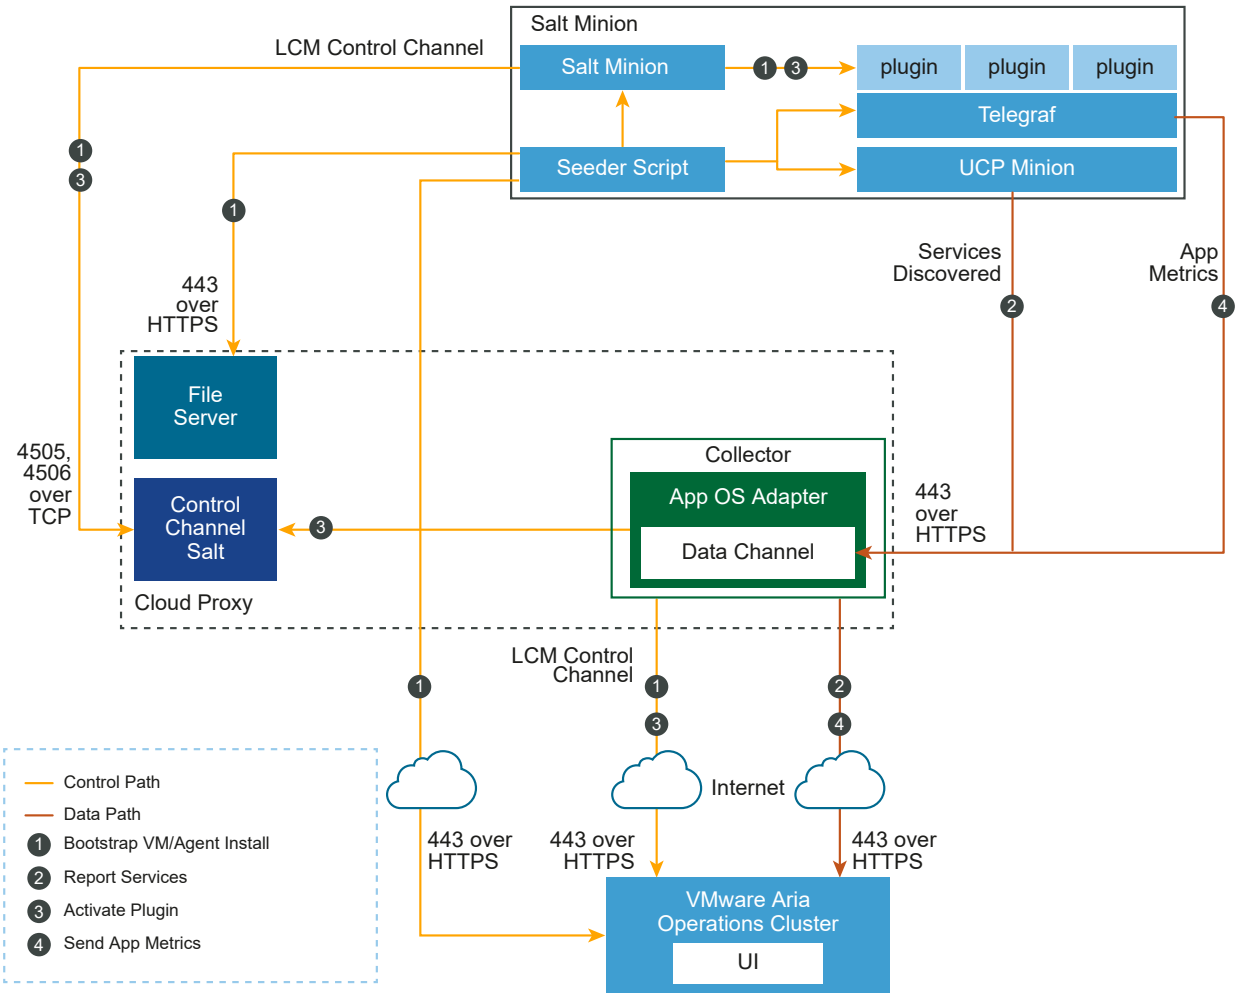


## Communication with Cloud Proxy and vCenter Server

Complete all the prerequisites required so that cloud proxy can communicate with vCenter Server. In the case of application monitoring high availability activated collector groups, these prerequisites are applicable for cloud proxies that are a part of the collector group.

**Note** Applicable to VMs managed by the vCenter Server cloud account.

- Port 443 in the ESXi where the workload VMs are deployed must be accessible to cloud proxy.
- Port 443 in Platform Services Controller and vCenter Server is accessible to cloud proxy. Open this port if vCenter Server is configured with an external Platform Services Controller.

- Configure a vCenter adapter. The vCenter Server user should have read access at the vCenter Server level and should also have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.

  **Note** For script-based agent install and uninstall, the `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries` permissions are not required.

## Communication with Cloud Proxy and Target Machines

Complete the prerequisites required during the handshake of cloud proxy with the target machines.

For the handshake of cloud proxy with the target machines, the following prerequisites are required:

- The target machine where Telegraf should run, cloud proxy, and VMware Aria Operations should be time synchronized. In case of VMs managed by vCenter Server cloud account, the ESXi instance where the VM is deployed, vCenter Server, VMs, VMware Aria Operations and cloud proxy should be time synchronized.
- The target machines must have access to ports 443, 4505, and 4506 on cloud proxy and on the virtual IP of the application monitoring high availability activated collector group. You can verify access by running the following commands on the target machine:
  - For Linux Machines:

    ```
    timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/4505"
    echo $?
    timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/4506"
    echo $?
    timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/443"
    echo $?
    ```

  - For Windows Machines:

    ```
    wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:4505
    wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:4506
    wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:443
    ```

    **Note** If you do not have `wget.exe` on the Windows machine, navigate to the `%temp%` folder or it's parent folder in File Explorer and search for `wget.exe` after you attempt installation.

    **Note** In the above commands, use virtual IP in the case of application monitoring high availability activated collector groups. For individual cloud proxy or cloud proxy which belongs to application monitoring high availability deactivated collector groups, use the cloud proxy FQDN.

- The necessary privileges for a user which are required for agent installation are mentioned in the User Account Prerequisites page.

- Target machine configuration requirements.
  - Linux requirements

    Commands: `/bin/bash, sudo, tar, awk, curl`

    Packages: `coreutils (chmod, chown, cat), shadow-utils (useradd, groupadd, userdel, groupdel), net-tools`

    Configure mount point on `/tmp` directory to allow script execution.
  - Windows requirement
    - The Visual C++ version must be higher than 14.
    - Performance Monitors on a Windows OS VM must be activated.
  - Windows 2012 R2 requirement

    The target machine must be updated with the Universal C Runtime. Refer to the following link for more information.
- VMware Tools must be installed and running on the vCenter Server VM on which you want to install the agent. For information about supported VMware Tools versions, click this link.

  **Note** Applicable only to VMs managed by the vCenter Server cloud accounts.
- Add a `tmp` folder with "exec" permission to install agents on the latest UAG Photon OS VM. To configure a mount point on the `/tmp` directory to allow script execution, run the following command: `<mount -o remount,exec /tmp>`

## User Account Prerequisites

There are certain user account prerequisites required for the install of agents.

### Windows Target Machine User Account Requirements

- To install agents,
  - The user must be either an administrator, or
  - A non-administrator who belongs to the administrator group.

### Linux Target Machine User Account Requirements

For Linux target machines, there are two user accounts for the Telegraf agent, such as the install user and the run-time user. User credentials which are provided during agent installation, are for the install user. The *arcuser* is a run-time user and needs a set of privileges which are necessary for the agent's components to run.

- `/tmp` mount point should be mounted with exec mount option.
- The following are minimal necessary permissions of the user to install agents and should be mentioned in `sudoers` file:

For example, for a user called **telegrafinstall**, you can find the `sudoers` file in the `/etc/sudoers` file or in the folder `/etc/sudoers.d/`:

```
Defaults:telegrafinstall !requiretty

Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/
cp*,/usr/bin/mkdir*,/bin/mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-
bootstrap.sh,/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh

telegrafinstall ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS
```

### Run-Time User Prerequisites

There are two ways in which a run-time user is created in Linux target machines: automatically and manually. A run-time user has a standard name and group, which is the *arcuser* and *arcgroup* respectively. If the **Create run time user on linux virtual machines, with required permissions as part of agent installation** check box is selected, the *arcuser* and *arcgroup* are created automatically. The check box is selected by default. A run-time user is also created automatically during a script-based install. If you choose to manually create the *arcuser* and *arcgroup*, here are the steps to do it manually:

Create the *arcgroup* and *arcuser* and associate the *arcgroup* as the primary group of the *arcuser*.

1 The *arcgroup* must be the primary group of the *arcuser*.

   The following commands can be used to create the *arcgroup* and *arcuser*:

   `groupadd arcgroup`

   `useradd arcuser -g arcgroup -M -s /bin/false`

2 The *arcuser* must be created with no home directory and no access to the login shell.

   For example, the `/etc/passwd` entry for the *arcuser* is as follows after adding *arcuser* and *arcgroup*.

   `arcuser:x:1001:1001::/home/arcuser:/bin/false`

3 The **arcuser** must have password-less specific set of privileges as mentioned below, which must be written in **/etc/sudoers** file or in the folder **/etc/sudoers.d/**:

```
Defaults:arcuser !requiretty
Cmnd_Alias VAPCOMMANDS=/usr/bin/systemctl * ucp-telegraf.service, !/usr/bin/systemctl
* * ucp-telegraf.service, /bin/systemctl * ucp-telegraf.service, !/bin/systemctl *
* ucp-telegraf.service, /usr/bin/systemctl * ucp-minion.service, !/usr/bin/systemctl
* * ucp-minion.service, /bin/systemctl * ucp-minion.service, !/bin/systemctl *
* ucp-minion.service, /usr/bin/systemctl * salt-minion.service, !/usr/bin/systemctl
* * salt-minion.service, /bin/systemctl * salt-minion.service, !/bin/systemctl
* * salt-minion.service, /usr/bin/systemctl * ucp-salt-minion.service, !/usr/bin/
systemctl * * ucp-salt-minion.service, /bin/systemctl * ucp-salt-
minion.service, !/bin/systemctl * * ucp-salt-minion.service, /usr/bin/netstat, /bin/
```

```
netstat, /opt/vmware/ucp/tmp/telegraf_post_install_linux.sh, /opt/vmware/ucp/bootstrap/
uaf-bootstrap.sh, /opt/vmware/ucp/content/runscript.sh, /opt/vmware/ucp/ucp-minion/bin/ucp-
minion.sh, /usr/bin/systemd-run, /bin/systemd-run
arcuser ALL=(ALL) NOPASSWD: VAPCOMMANDS
```

## Steps to Monitor Applications

You can monitor and collect metrics for your application services and operating systems.

The following flowchart describes how you can set up VMware Aria Operations for application monitoring in the case of VMs managed by the vCenter Server cloud account.

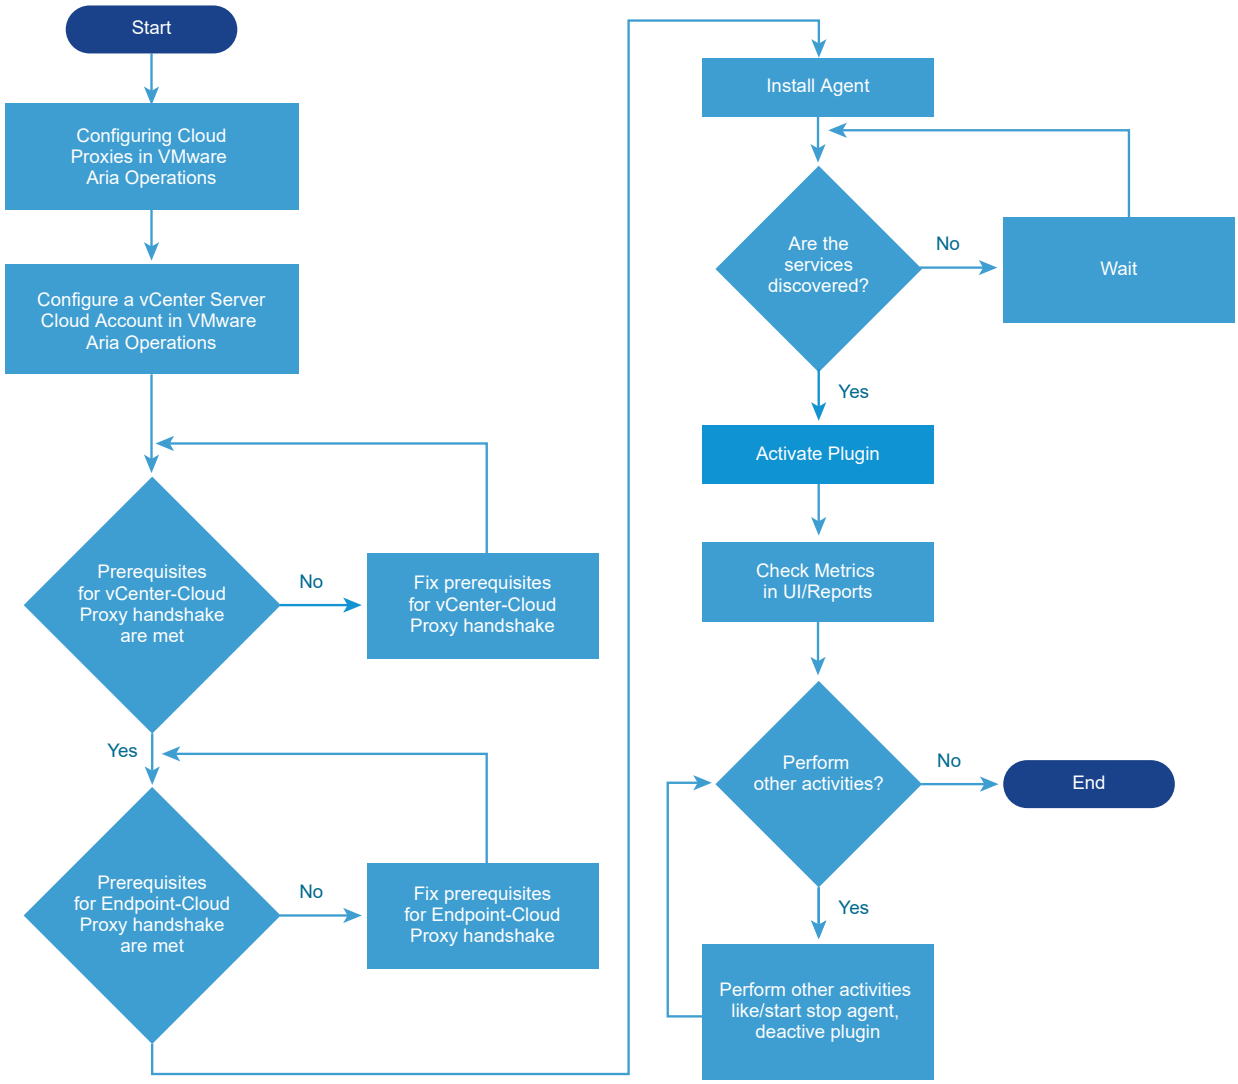


Follow these steps to monitor applications.

1 Configure cloud proxy on which the AppOS adapter instance is created.

   For more information, see Configuring Cloud Proxies in VMware Aria Operations.

2 Configure a vCenter Server cloud account that uses the cloud proxy configured in step 1.

For more information, see Configure a vCenter Server Cloud Account in VMware Aria Operations.

**Note** This step is not applicable for target machines.

3 Complete all the prerequisites.

For more information, see Prerequisites.

4 Install agents on selected target machines.

For more information, see Install an Agent from the UI or Install/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform .

5 Activate an application service.

For more information, see Activate an Application Service.

6 View the summary of application services and operating systems discovered in VMware Aria Operations.

For more information about monitoring your applications in VMware Aria Operations, see Summary of Discovered and Supported Operating Systems and Application Services .

## Install and Uninstall a Telegraf Agent

You can install and uninstall Telegraf agents on VMs managed by the vCenter Server cloud account from the user interface of VMware Aria Operations or by running a script on physical servers and on machines not managed by VMware Aria Operations.

### Install an Agent from the UI

You must select the vCenter Server VMs on which you want to install the agent. All the vCenter Server VMs of the vCenter Server cloud account are listed in the **Manage Telegraf Agents** page.

#### Prerequisites

**Note** You cannot install an agent from the VMware Aria Operations user interface, on physical servers and target machines (Discovered in VMware Aria Operations as the `Endpoint` object type).

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

#### Procedure

1 From the left menu, select **Environment > Applications**. From the **Applications** panel, select **Manage Telegraf Agents**. You see the **Manage Telegraf Agents** page.

2 From the **Manage Telegraf Agents** page, select the VMs managed by the vCenter Server cloud account on which you want to install the agent.

3 From the **Manage Telegraf Agents** page, click **Agent Actions**, and then click **Install**.

4 If the VMware Aria Operations Application Monitoring Adapter instance has not been created for the specific vCenter Server, you see the **Installing Telegraf Agent** dialog box. You can configure high availability for application monitoring via collector groups:

- Select **Cloud Proxy Groups, with High-Availability** if you want to associate an application monitoring high availability activated collector group to the vCenter Server. You can choose the collector group from the **HA Enabled Group** drop down list, or
- Select **Cloud Proxies, without High-Availability** if you want to associate a cloud proxy to the vCenter Server. You can choose a cloud proxy from the **Cloud Proxy** drop down list.
- Click **Done**.

**Note** If the vCenter Server is on the default collector, the cloud proxy options are blank. If the vCenter Server is on a specific cloud proxy, the cloud proxy option is automatically populated. To deploy the VMware Aria Operations Application Monitoring Adapter instance on a different cloud proxy, select another cloud proxy or application monitoring high availability activated collector group.

5 You see the **Manage Agent** dialog box.

6 From the **How do you want to provide VM Credentials** page, complete the following steps:

a If you have a common user name and password for all the VMs, select the **Common username and password** option.

b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.

c Click **Next**.

7 From the **Provide Credentials** page, depending on whether you have a common credential for all VMs managed by the vCenter Server cloud account or different credentials for all VMs managed by the vCenter Server cloud account, enter the following details:

a If the selected VMs have a common user name and password, enter the common user name and password.

b For different user names and passwords for each VM, download the CSV template and add the required details such as the user name, password for each VM managed by the vCenter Server cloud account. Use the **Browse** button to select the template.

c The **Create run time user on Linux virtual machines, with required permissions as part of agent installation** check box is selected by default. For more information, see User Account Prerequisites.

d Click **Next**.

8 From the **Summary** page, you can view the list of VMs managed by the vCenter Server cloud account on which the agent is to be deployed.

9 Click **Install Agent**. Refresh the UI to view the agents that are installed.

On User Account Control (UAC) deactivated Windows VMs managed by the vCenter Server cloud account, the Telegraf agent installation completes with the **Agent running** agent status and **Install Success** Last operation status, if all prerequisites are performed.

**UAC Activated on Windows VMs Managed by the vCenter Server Cloud Account**

In the case of UAC activated VMs, for a non-administrator user who belongs to the administrator group, Telegraf agent installation completes with the **Not Started** agent status and **Download Success** last operation status, if all the prerequisites are performed. The bits are downloaded to the VM. You must manually install the bits. From `C:\VMware\UCP\downloads`, run a bootstrap launcher, and then the agent status becomes **Agent Running** and the Last operation status becomes **Install Success**.

- Open PowerShell with administrator privileges.
- Go to `$SYSTEMDRIVE\VMware\UCP\downloads` folder: `cd $SYSTEMDRIVE\VMware\UCP\downloads`
- Run the `cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1` command.
- View the results from `uaf_bootstrap.log`.
- Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Telegraf Agents** page.

10 The agent discovers the application services that are installed on the target machines. You can view the application services by clicking the drop down arrow against the VMs managed by the vCenter Server cloud account on which the agent is installed in the **Manage Telegraf Agents** page. You can view the status of agent installation from the **Agent Status** column in the **Manage Telegraf Agents** page.

11 After installation, an OS object will be created under VMs managed by the vCenter Server cloud account topology. To view the object, select a VM from **Environment > Applications > Manage Telegraf Agents** page. Click the **Go details** button, and click on the **Metrics** tab. Expand the **Object relationship** widget, and select the OS object.

12 If agent installation fails, a detailed error message is displayed when you hover over the red exclamation icon before the Install Failed text in the **Last operation status** column. You can take appropriate action based on the error message to resolve the agent installation issues.

#### What to do next

You can manage the services on each agent.

For information about uninstalling an agent, see Uninstall an Agent.

### Install/Uninstall an Agent Using a Script on Linux Platforms

You can install or uninstall an agent on a target machine using a script.

Prerequisites

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- The cloud account must be configured for the vCenter Server to which the VM belongs. The application monitoring adapter that is mapped to the vCenter Server is created if it does not exist.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account. For more information about object types, see OS and Application Monitoring.

- The VM managed by the vCenter Server cloud account must be available in VMware Aria Operations.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account.

- The guest IP must be properly configured and should be unique across vCenter Servers. If more than one VM is managed by the vCenter Server cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account. For more information about object types, see OS and Application Monitoring.

- Ensure that the Internet is enabled.
- The user must have the following permissions:
  - Administration > REST-APIs > All other Read, Write APIs
  - Administration > REST-APIs > Read access to APIs
  - Environment > Applications > Manage Telegraf Agent
  - Environment > Actions > Bootstrap virtual machines
  - Environment > Actions > Download bootstrap

  **Note** To check permissions, navigate to **Administration > Access Control > Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version must be >=10.2.

  **Note** Applicable only to VMs managed by the vCenter Server cloud account.

### Procedure

1 Log in to the target machine on which you want to install/uninstall the agent and download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/ telegraf-utils.sh`. Run one of the following commands:

```
wget --no-check-certificate “https://<CloudProxy>/downloads/salt/telegraf-utils.sh”
curl -k “https://<CloudProxy>/downloads/salt/telegraf-utils.sh” --output telegraf-utils.sh
```

**Note** Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2 Make the script executable by running the following command:

```
chmod +x telegraf-utils.sh
```

3 Create the user's refresh token associated with the current organization from the following location in the Cloud Service portal: `User/Organization Settings > My Account > API Tokens > Generate a New API Token`

4 To execute the script and install/uninstall the agent, run the following command:

```
telegraf-utils.sh  product-managed -c cloud_proxy_ip_or_collector_group_name -t token (-v
vmwareariaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d download_tmp_dir
-s sleep_seconds -i list_of_IPs -o operation]

Description of arguments:
      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector
Group Name
      -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem
vmwareariaoperations_auth_token.
            saas: CSP Refresh Token of the user/account. For getting new token, follow -
User/
                  Organization Settings >> My Account >> API Tokens >> Generate a New API
Token.
                  Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq

      -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation.
Default: current directory

      -g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware
Aria Operations SaaS gateway URL
      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL
      -o : [OPTIONAL] operation - The type of operation to be performed
{install,uninstall} Default: install
      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script execution.
This optional parameter will be helpful when this script is used in vRA to have agent
installed on deploying VM. Recommended time 600 seconds.
      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated IPs
that are set as properties of a VM in VMware Aria Operations (VM -> Properties -> Network
-> <Integer> -> IP Address). If single adapter has multiple IP Addresses, then delimit
```

```
them using '_'
                    ex:- If your VM has 3 adapters with IP Address as follows: Adapter1:
10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4
                    Then this parameter should be given as
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"

Example: /bin/bash telegraf-utils.sh product-managed -t 8dab02cc-277c-4383-b910-
bd2e89c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**Note** The `-c` argument is mandatory to run the helper script.

**Note** The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**Note** When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Install/Uninstall an Agent Using a Script on a Windows Platform

You can install or uninstall an agent on a target machine using a script.

#### Prerequisites

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- The cloud account must be configured for the vCenter Server to which the VM belongs. The application monitoring adapter that is mapped to the vCenter Server is created if it does not exist.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account. For more information about object types, see OS and Application Monitoring.

- The VM managed by the vCenter Server cloud account must be available in VMware Aria Operations.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account.

- The guest IP must be properly configured and should be unique across vCenter Servers. If more than one VM is managed by the vCenter Server cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.

  **Note** This prerequisite is applicable only to the VMs managed by the vCenter Server cloud account. For more information about object types, see OS and Application Monitoring.

- Ensure that the Internet is enabled.
- The VMware Aria Operations user must have the following permissions:
  - Administration > REST-APIs > All other Read, Write APIs
  - Administration > REST-APIs > Read access to APIs
  - Environment > Applications > Manage Telegraf Agent
  - Environment > Actions > Bootstrap virtual machines
  - Environment > Actions > Download bootstrap

  **Note** To check permissions navigate to **Administration > Access Control > Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version must be >=10.2.

  **Note** Applicable only to VMs managed by the vCenter Server cloud account.

Procedure

1 Log in to the target machine on which you want to install/uninstall the agent, download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/telegraf-utils.ps1` .

  If the script download fails with the following message: `The request was aborted: Could not create SSL/TLS secure channel`, follow the steps mentioned in Script Download Fails on a Windows Platform.

  Run one of the following commands:

```
PowerShell command:
Invoke-WebRequest "https://<CloudProxy>/downloads/salt/telegraf-utils.ps1" -OutFile
telegraf-utils.ps1

Or if you have the wget tool:
wget --no-check-certificate https://<CloudProxy>/downloads/salt/telegraf-utils.ps1
```

  **Note** Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2 Create the user's refresh token associated with the current organization from the following location in the Cloud Service portal: `User/Organization Settings > My Account > API Tokens >> Generate a New API Token`

3 To execute the script and install/uninstall the agent, run the following command:

```
telegraf-utils.ps1  product-managed -c cloud_proxy_ip_or_collector_group_name -t token (-v
vmwarearіaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d download_tmp_dir
-s sleep_seconds -i list_of_IPs -o operation]

Description of arguments:
      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector
Group Name
      -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem
vmwarearіaoperations_auth_token.
            saas: CSP Refresh Token of the user/account. For getting new token, follow -
User/
                  Organization Settings >> My Account >> API Tokens >> Generate a New API
Token.
                  Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq

      -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation.
Default: current directory

      -g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware
Aria Operations SaaS gateway URL
      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL
      -o : [OPTIONAL] operation - The type of operation to be performed
{install,uninstall} Default: install
      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script execution.
This optional parameter will be helpful when this script is used in vRA to have agent
installed on deploying VM. Recommended time 600 seconds.
      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated IPs
that are set as properties of a VM in VMware Aria Operations (VM -> Properties -> Network
-> <Integer> -> IP Address). If single adapter has multiple IP Addresses, then delimit
them using '_'
                    ex:- If your VM has 3 adapters with IP Address as follows: Adapter1:
10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4
                    Then this parameter should be given as
```

```
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"

Example: .\telegraf-utils.ps1 product-managed -t 8dab02cc-265c-4392-b910-
bd2e89c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**Note** The `-c` argument is mandatory to run the helper script.

**Note** The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

**Note** Do not use a space in the configuration path. Paths with spaces can be passed as a short name notation, such as `c:\PROGRA~1` for `c:\Program Files`.

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**Note** When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Uninstall an Agent

You must select the VMs managed by the vCenter Server cloud account on which you want to uninstall the agent.

#### Prerequisites

**Note** You cannot uninstall an agent from the UI for physical servers and VMs unmanaged by the vCenter Server cloud account (`Endpoint` object type in VMware Aria Operations).

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

#### Procedure

1. From the left menu, select **Environment > Applications**. From the **Applications** panel, select **Manage Telegraf Agents**. You see the **Manage Telegraf Agents** page.
2. From the **Manage Telegraf Agents** page, select the VMs that are managed by the vCenter Server cloud account on which you want to uninstall the agent.
3. From the **Manage Telegraf Agents** page, click the horizontal ellipsis, and then click **Uninstall**. You see the **Manage Agent** dialog box.

4 From the **How do you want to provide VM Credentials** page, complete the following steps:

   a If you have a common user name and password for all the VMs, select the **Common username and password** option.

   b If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.

   c Click **Next**.

5 From the **Provide Credentials** page, depending on whether you have a common credential for all the VMs or different credentials for all VMs, enter the following details:

   a If your VM has a single user name and password, enter the common user name and password.

   b For multiple user names and passwords for each VM, download the CSV template and add the details. Use the **Browse** button to select the template.

   c Click **Next**.

6 From the **Summary** page, you can view the list of VMs on which the agent is uninstalled.

7 Click **Uninstall Agent**. Refresh the UI to view the progress of agent uninstallation.

   The **Agent Status** column and the missing drop down arrow against the VM in the workspace, indicate that uninstallation is complete and that there are no application services discovered on each agent.

   **UAC Activated on Windows Machine**

   The bits are downloaded to the VM. You have to manually uninstall the bits.

   - Open PowerShell with administrator privileges.
   - Go to `$SYSTEMDRIVE\VMware\UCP\downloads` folder: `cd $SYSTEMDRIVE\VMware\UCP\downloads`
   - Run the `cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1` command.
   - View the results from `uaf_bootstrap.log`.
   - Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Telegraf Agents** page.

## Activate an Application Service

To monitor application services running on the target machines, you must configure Telegraf agents in VMware Aria Operations.

**Note** The maximum allowed application services count per target machine is 200.

After you have installed the agent, you can activate plugins to monitor application services. Additionally, you can monitor a discovered service on VMs managed and unmanaged by the vCenter Server cloud account and physical servers.

You can configure discovered services and carry out custom monitoring actions. For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.

**Prerequisite**

- If plugin activation requires the location of a file (for example, client certificates for SSL Trust) on the target machine, the location and the files should have appropriate read permissions for the *arcuser* to access those files. Configuration of application services and custom monitoring using Telegraf is referred to as plugins.

  **Note** If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

- Linux process activation for Pid files works only if the Pid file and its parent directories have read permission for **Others**.

**Activate an Application Service on a Single Target Machine**

To monitor an application service on a single target machine, complete the following steps:

1. From the left menu, click **Environment > Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.
2. Select the target machine on which the agent is running. You can use the filter functionality too. For example, filter by **Agent Status > Agent Running**.
3. Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Services Discovered** section.
4. From the **Services Discovered** section, select a discovered service or custom monitored service, click the vertical ellipsis and then click **Add**.
5. Activate the application service from dialog box that is displayed on the right side.
6. Enter the details for each instance that you add and click **Save**. For configuration details of each application, see Configuring Supported Application Services.

   Fields with a star are mandatory.

   For more information about the status details that appear against the application services, see the table called Status Details in Additional Operations from the Manage Telegraf Agents Page .

   To edit or delete instances of application services, click the **Edit** or **Delete** options from the vertical ellipsis against application service you added. After the services have been added and saved, click the drop down arrow against the application service to view the list of services and their status.

**Activate an Application Service on Multiple Target Machines**

To monitor an application service on multiple target machines, complete the following steps:

1 From the left menu, click **Environment > Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.

2 Select the target machines on which the agent is running where the agent type is `product-managed`. You can use the filter functionality too. For example, filter by **Agent Status > Agent Running** and **Agent Type > Product-managed agent**.

**Note** You can only activate an application service for all object types (Endpoint and Virtual Machine) with a product-managed Telegraf agent installed and running.

3 Click the **Horizontal Ellipsis** and then click **Activate Service**. You see the **Activate Service** dialog box.

4 From the **Activate Service** dialog box, select a discovered service or custom monitoring service by clicking the drop-down arrow in the **Select a Service** field and click **Confirm**.

5 Activate the application service from dialog box that is displayed on the right side.

6 Enter the details for each instance that you add and click **Save**. For configuration details of each application, see Configuring Supported Application Services.

   Fields with a star are mandatory.

   For more information about the status details that appear against the application services, see the table called Status Details in Additional Operations from the Manage Telegraf Agents Page .

   To edit or delete instances of application services, click the **Edit** or **Delete** options from the vertical ellipsis against application service you added. After the services have been added and saved, click the drop down arrow against the application service to view the list of services and their status.

   For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.

The following special characters are permitted in the DB user field: `'[]{} (),.<> ?:!|/~@#$%^&*-_ +=`

You can provide DB name lists in the following format `['DBNAME_1', 'DBNAME_2', 'DBNAME_3']` where DBNAME_1, DBNAME_2, DBNAME_3 must not contain quotes such as `'` and `"`.

### Application Availability

When an application service is activated, the **Application Availability** metric is collected and displays if the application service is running on the target machine or if it is down. **1** indicates that the application service is running on the target machine and **0** indicates that the application service is down. This metric is available for all supported application services except JAVA application service.

For information about deactivating a service, see Deactivate an Application Service.

## Configuring Supported Application Services

Twenty-three application services are supported in VMware Aria Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

### Active Directory

Active Directory is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

### Active MQ

ActiveMQ is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Server URL | Yes | http://localhost:8161 |
| User name | Yes | User name for Active MQ. Example: admin |
| Password | Yes | Password |
| Installed Path | Yes | The path on the target machine where Active MQ is installed.<br>Example:<br>For Linux target machines: `/opt/apache-activemq`<br>For Windows target machines: `C:\apache-activemq-5.15.2` |

### Apache HTTPD

Apache HTTPD is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Status Page URL | Yes | http://localhost/server-status?auto |
| User name | No | User name for Apache HTTPD service. **`Example:root`** |
| Password | No | Password |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Response String Match | No | Optional substring or regex match in the body of the response (case sensitive). For example, **`ok`**. You must escape regex special characters **`?`**, **`(`**, **`)`**, **`.`** with a backslash. For example, for matching against **`?doodles`**, escape the **`?`** with **`\?`**, such as, **`\?doodles`**. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

## Cassandra Database

Cassandra database is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the application instance. |
| Installed Path | Yes | Valid file path. |
| URL | Yes | http://localhost:8778 |

## Hyper-V

Hyper-V is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the application service. |

## Java

Java is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |

| Name | Mandatory? | Comment |
|---|---|---|
| Installed Path | Yes | The path on the target machine where Java is installed. Example: For Linux target machines: /opt/vmware/ucp ; For Windows target machines: C:\VMware\UCP |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

### JBoss

JBoss is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where JBoss is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

### MongoDB

MongoDB is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MongoDB is running. Example: 27017 |
| Hostname | No | Optional hostname for the MongoDB Service. |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Username | No | User name for MongoDB. Example: Root |
| Password | No | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

## MS Exchange

MS Exchange is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

## MS IIS

MS IIS is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

## MS SQL

MS SQL is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Instance | Yes | Instance name of the MS SQL server. |
| Port | No | The port where MS SQL is running. Example: 1433 |
| Hostname | No | Optional hostname for the MS SQL Service. |
| Username | Yes | User name for MS SQL. Example: Root |
| Password | Yes | Password |

## MySQL

MySQL is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MySQL is running. Example: 3306 |
| User name | Yes | User name for MySQL service. Example: Root |
| password | Yes | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Hostname | No | Optional hostname for the MySQL Service |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma separated. For example, 'database1','database2','database3'. |
| TLS Connection | No | Allowed values are true, false, and skip-verify. |

## Nginx

Nginx is supported in VMware Aria Operations.

### Table 4-21.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the application instance. |
| Status Page URL | Yes | http://127.0.0.1:8081/nginx_status |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |

Table 4-21. (continued)

| Name | Mandatory? | Comment |
| --- | --- | --- |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain and host verification. Expected: True/False. |

### NTPD

NTPD is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

### Oracle Database

Oracle database is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the application instance. |
| OracleDB Username | Yes | User name for the Oracle database instance. |
| OracleDB Password | Yes | Password for the Oracle database instance. |
| OracleDB SID | Yes | SID of the Oracle database instance. |

### Pivotal Server

Pivotal Server is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where Pivotal server is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

## Postgres

Postgres is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where PostgreSQL is running. Example: 5432 |
| User name | Yes | User name for PostgreSQL service. Example: Root |
| Password | Yes | Password |
| SSL Connection | No | Allowed values are deactivate, verify-ca, verify-full. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: true/false. |
| Hostname | No | Optional hostname for the PostgreSQL Service. |
| Default Database | No | The database for initiating connection with the server |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated, for example , 'database1','database2','database3'. |
| Ignored Databases | No | Comma-separated list of databases that need not be monitored. Each of the database names to be excluded from monitoring must be enclosed in single quotes and the databases themselves should be comma-separated for example, 'database1','database2','database3'. |

## RabbitMQ

RabbitMQ is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Management Plugin URL | Yes | http://localhost:15672 |
| User name | No | User name for RabbitMQ. Example: Guest |
| Password | No | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |
| Nodes | No | Each of the RabbitMQ data collection nodes should be in single quotes and the nodes themselves should be comma-separated. The list of nodes must be enclosed in square brackets. For example ['rabbit@node1','rabbit@node2',.....] |

## Riak

Riak is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Server URL | Yes | http://localhost:8098 |

### Sharepoint

Sharepoint is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

### Tomcat

Tomcat is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where Tomcat is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on thetarget machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

### Weblogic

Weblogic is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:7001 |
| Installed Path | Yes | The path on the target machine where WebLogic is installed. |
| User name | Yes | User name for WebLogic. Example: admin |
| Password | Yes | Password |

| Name | Mandatory? | Comment |
|---|---|---|
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

### Websphere

Websphere is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| IBM Websphere Server URL | Yes | Example : http://localhost:9081 |
| Websphere Authorization Token | Yes | To generate the token, follow the below steps:<br>■ Go to https://www.base64encode.org.<br>■ Type in the user and password created in the format: user:password<br>■ Click the Encode button.<br>■ Copy the resulting Base64 encoded string. Example: d2F2ZWZyb250OndhdmVmcm9udA== |

### Remote Checks

### HTTP Remote Check

HTTP is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the remote check instance. |
| URL | Yes | http://localhost |
| Method | Yes | GET/POST/PUT |
| Proxy | No | Proxy URL: http://localhost |
| Response Timeout | No | Timeout for the connection in seconds. For example, 10. |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Follow Redirects | No | True/False if redirects from the server. For example, true/false (all small values). |
| Body | No | HTTP request body. |
| Response String Match | No | Substring or regex match in the response body. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificates | No | Path to the SSL certificate file on the target machine. |
| SSL Key | No | Path to the SSL key file on the target machine. |
| Skip Host & chain verification | No | Use SSL but skip chain and host verification. Expected: True/False. |
| Headers | No | HTTP Request Headers, which can look like the following. For example: `accept = "application/json"` `Authorization = "5609fe9d-cddd-4654-898b-26d9b67f137a::4ca86c16-ac01-4e17-90da-fb55eb36571a"` |

## ICMP Remote Check

ICMP is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the remote check instance. |
| FQDN/IP | Yes | Host name to send the packets. Example: *example.org* |
| Count | No | Number of ping packets to send per interval. For example, 1. |
| Ping Interval | No | Time to wait between ping packets in seconds. For example, 10.0.<br>**Note** Follow the decimals as mentioned in the example. |
| Timeout | No | Timeout to wait for ping response in seconds. For example, 10.0.<br>**Note** Follow the decimals as mentioned in the example. |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Deadline | No | The total ping deadline in seconds. For example, 30. |
| Interface | No | Interface or source from which to send a ping. |

### TCP Remote Check

TCP is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the remote check instance. |
| Address | Yes | <hostname>:port |
| Send | No | The given string is sent to the TCP. It can be any string of your choice. |
| Expect | No | The given string is expected from the TCP. It can be any string of your choice. |
| Timeout | No | Timeout for the connection to the TCP server. For example, 10. |
| Read Timeout | No | Timeout for the response from the TCP server. For example, 10. |

### UDP Remote Check

UDP is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display name of the remote check instance. |
| Address | Yes | <hostname>:port |
| Send | Yes | The given string is sent to the UDP. |
| Expect | Yes | The given string is expected from the UDP. |
| Timeout | No | Timeout for the connection to the UDP server. For example, 10. |
| Read Timeout | No | Timeout for the response from the UDP server. For example, 10. |

### Configuring Supported VeloCloud Services

Eight VeloCloud application services are supported in VMware Aria Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

### VeloCloud Orchestrator

VeloCloud Orchestrator and the following services are supported in VMware Aria Operations.

- VeloCloud Orchestrator
- Nginx

  **Note** To activate the plugin for ngnix service you must use the loopback address in the url `http://127.0.0.1/nginx_status`.

- Clickhouse
- Network Time Protocol
- MySQL
- Redis
- Java Application

  **Note** Java application gets discovered after bootstrapping a VeloCloud Orchestrator virtual machine, but you must ignore it, as we do not monitor the Java application.

In VeloCloud Orchestrator, we monitor the following services. For each of these services we display a metric which indicates the service status:

- Backend
- Portal
- Upload

VeloCloud Orchestrator details.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the VeloCloud Orchestrator instance. |

### Nginx

Nginx is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Status Page URL | Yes | http://127.0.0.1/nginx_status |
| SSL CA | No | Path to the SSL CA file on the end point VM. |
| SSL Certificate | No | Path to the SSL Certificate file on the end point VM. |

| Name | Mandatory? | Comment |
| --- | --- | --- |
| SSL Key | No | Path to the SSL Key file on the end point VM. |
| Skip SSL Verification. | No | Use SSL but skip chain & host verification. Expected: True/False. |

### ClickHouse

ClickHouse is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Servers URL | Yes | http://127.0.0.1:8123 |
| User name | No | User name for the ClickHouse service. |
| Password | No | Password |

### NTPD

NTPD is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

### MySQL

MySQL is supported in VMware Aria Operations.

To activate the MySQL plug-in and fetch the credentials, refer to the article Steps to fetch password for telegraf user of MySQL, while activating plugin (81153) at the VMware Support Knowledge Base.

Use the port number 3306 to run MySQL and the telegraf credentials and activate the plug-in.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MySQL is running. Example: 3306 |
| User name | Yes | User name for the MySQL service. Example: Root |
| password | Yes | Password |
| SSL CA | No | Path to the SSL CA file on the end point VM. |

| Name | Mandatory? | Comment |
|---|---|---|
| SSL Certificate | No | Path to the SSL Certificate file on the end point VM. |
| SSL Key | No | Path to the SSL Key file on the end point VM. |
| Hostname | No | Optional hostname for the MySQL Service |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated. For example, 'database1','database2','database3'. |
| TLS Connection | No | Accepted values are true, false, and skip-verify. |

### Redis

Redis is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Redis URL | Yes | servers = ["tcp://localhost:6379"] |
| SSL CA | No | Secure Socket Layer Certification Authority. |
| SSL Certificate | No | Secure Socket Layer Certificate. |
| SSL Key | No | Secure Socket Layer Key |
| Skip SSL Verification. | No | Skips verification for SSL. |

### VeloCloud Gateway

VeloCloud Gateway and the following services are supported in VMware Aria Operations

- Network Time Protocol
- VeloCloud Gateway

In VeloCloud Gateway, we monitor the following processes. For each of these processes, we display a metric which indicates the process status.

- bgpd
- watchquagga
- gwd

- mgd
- natd
- ssh
- vc procmon
- vcsyscmd

VeloCloud Gateway details.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the VeloCloud Gateway instance. |

NTPD

NTPD is supported in VMware Aria Operations.

| Name | Mandatory? | Comment |
| --- | --- | --- |
| Display Name | Yes | Display Name of the application instance. |

### Pre-Requirements for Application Services

For the Telegraf agent to collect metrics for some of the application services, you must make modifications in the target machines. After you make these modifications, the agent starts collecting metrics. You must SSH to the virtual machine where you have deployed the agent and modify the configuration files.

#### Apache HTTPD

For the agent to collect metrics, modify the conf file available in `/etc/httpd/conf.modules.d/status.conf` and activate the mod_status for the HTTPD plugin.

```
<IfModule mod_status.c>

<Location /server-status>

    SetHandler server-status

</Location>

ExtendedStatus On

</IfModule>
```

If the conf file is not available, you must create one. Restart the HTTPD service after modifying the conf file with the following command:

```
systemctl restart httpd
```

### Java Application

Support for Java application monitoring is currently limited to applications that support jolokia based JMX monitoring. Currently, as part of Java plugin activation, both the jolokia.jar and jolokia.war files are placed in the "Installed Path". If the Java application to be monitored is a servlet container, then the jolokia.war file is automatically deployed to the "Installed Path" that you provided when you configured the Java application in VMware Aria Operations. If auto deployment is not supported or not activated, restart the application after activating the plugin or refer to the application's product documentation for the deployment of manual .war.

For all other java applications, first activate the plugin then include the path to the jolokia.jar file in the `"-javaagent:" JVM flag as a command line parameter`. For example, `"-javaagent:/<path_to_jolokia>/jolokia.jar"`, `<path_to_jolokia>` is the "Installed Path" which was provided during application configuration. Restart the application.

**Note** For monitoring a customized third-party java application, ensure that:

- The jolokia based JMX monitoring is supported.
- In case auto deploy is turned off, manual .war deployment is supported for servlet container application.
- `"-javaagent:" JVM flag` is added as part of the command line parameter (JVM_OPTS)

### Nginx

Add the following lines to the conf file available in `/etc/nginx/nginx.conf`:

```
http {
    server {
        location /status {
            stub_status on;
   access_log off;
   allow all;
  }
    }
}
```

Restart the Nginx service with the following command:

```
systemctl restart nginx
```

### Postgres

In the configuration file available in the `/var/lib/pgsql/data/pg_hba.conf`, change the value of `local all postgres peer` to `local all postgres md5` and restart the service with the following command:

```
sudo service postgresql restart
```

### Cassandra

To monitor the Cassandra database application, the Jolokia jar must be included as a JVM input to the Cassandra database application. Complete the following steps:

1 Modify `/etc/default/cassandra`.

```
echo "export JVM_EXTRA_OPTS=\"-javaagent:/usr/share/java/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=localhost\"" | sudo tee -a /etc/default/cassandra
```

**Note** You can download the latest version of the Jolokia agent from https://jolokia.org/download.html. For a .JAR file deployment, you have to restart the application service after including the full file path of the JAR in the JMX argument of the JAVA process which you are monitoring.

2 Alternatively, you can activate the agent by modifying `cassandra-env.sh`. Include the following line at the end of the `cassandra-env.sh`:

```
JVM_OPTS="$JVM_OPTS -javaagent:/usr/share/java/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=localhost"
```

After you see the JVM inputs, restart the Cassandra service.

### Oracle Database

To monitor the Oracle database, complete these steps:

1 Download the instant client library from: `https://www.oracle.com/database/technologies/instant-client/downloads.html`.

You must download the Oracle instant library and included it in the PATH.

2 Install Python 3.6 or later. Install the cx_Oracle module.

```
python3 -m pip install cx_Oracle --upgrade
```

3 Create a User.

```
CREATE USER <UserName> IDENTIFIED BY <yourpassword>;
 GRANT select_catalog_role TO <UserName>;
 GRANT CREATE SESSION TO <UserName>;
```

4 **For a Windows VM:**

a Append the Oracle instant client library path and the Python path to the PATH environment variable.

b Install the ARC agent.

5 **For a Linux VM:**

a Install the ARC agent.

b Get the file `exec_oracle_python.sh(wget --no-check-certificate https://<CP_IP>/downloads/salt/content-accessories/exec_oracle_python.sh)` from `/opt/vmware/ucp/`.

c Fill in `/opt/vmware/ucp/exec_oracle_python.sh` with absolute values for `LD_LIBRARY_PATH, ORACLE_HOME, PYTHON_BIN`, and `TNS_ADMIN`.

`LD_LIBRARY_PATH`: Path to the Oracle instant client library. For example, `/opt/vmware/ucp/instantclient_21_4`.

`ORACLE_HOME`: Directory in the file system where the Oracle software is installed. For example, `/u01/app/oracle/product/19.3.0/dbhome_1/`.

`PYTHON_BIN`: Path to the Python executable. For example, `/usr/bin/python`.

`TNS_ADMIN`: Environment variable that points to the directory where the SQL*Net configuration files are located. For example, `/u01/app/oracle/product/19.3.0/dbhome_1/network/admin`.

d Activate the Oracle DB plugin.

### Active MQ 5.16 and Later Versions

To activate Active MQ 5.16 and later versions, complete the following steps:

- Navigate to `/opt/activemq/apache-activemq-5.16.0/webapps/api/WEB-INF/classes/jolokia-access.xml`
- Remove or comment out the following lines:

```
<cors>
  <strict-checking/>
</cors>
```

- Restart the Active MQ service.

### Microsoft SQL Server

The user account must have the following permissions to monitor the Microsoft SQL Server application service with Telegraf.

```
USE master;
GO
CREATE LOGIN [telegraf] WITH PASSWORD = N'mystrongpassword';
GO
GRANT VIEW SERVER STATE TO [telegraf];
GO
GRANT VIEW ANY DEFINITION TO [telegraf];
GO
```

**Note** If you want to monitor named instances of the Microsoft SQL Server application service, then the plugin activation input should not have the Port field configured. Use Port only when there is no instance in the server (default 1433).

## Additional Operations from the Manage Telegraf Agents Page

After you have configured cloud proxy and installed a product-managed Telegraf agent, configure a vCenter Server Cloud account (relevant only to VMs managed by the vCenter Server cloud account). You can manage the agents on all the target machines from the **Manage Telegraf Agents** page.

From the **Manage Telegraf Agents** page, you can:

- Use the filter to easily find physical servers or VMs managed or unmanaged by the vCenter Server cloud accounts.
- View data collected on all target machines.
- Start, stop, and update agents on all target machines where the product-managed agent is installed. You can also discover and manage the services on each product-managed Telegraf agent that you install.

  **Note** The maximum allowed application services count per target machine is 200.

Additionally, all the types of supported monitored object types such as `Virtual Machine` and `Endpoint` are consolidated in the **Manage Telegraf Agents** page so that you have visibility of the entire environment you are monitoring and can manage agents and monitor the content from one single page, depending on the agent you are using and type of the objects available. For more information about object types, see OS and Application Monitoring.

### Telegraf Agents in VMware Aria Operations

For more information about Telegraf agents, see What Are The Different Types of Telegraf Agents in VMware Aria Operations?.

**Note** Data collected from product-managed agents, open-source agents, and user-managed agents can be viewed from the **Manage Telegraf Agents** page. For objects that have open-source agents and user-managed Telegraf agents installed and running, you can only view data that is collected, you cannot install, uninstall, or manage the lifecycle of the agents.

### Where You Manage the Agents

To manage the agents and application services, from the left menu, click **Environment > Applications**. From the **Applications** panel, click the **Manage Telegraf Agents** tab.

**Note** In the table below, **Vertical Ellipsis** is applicable to single VM operations and **Agent Actions** is applicable to group actions.

## Table 4-22. Options

| Options | Description |
| --- | --- |
| **Agent Actions/Vertical Ellipsis > Install** | Installs the agents on the selected VM/s managed by the vCenter Server cloud account. Select the VM/s managed by the vCenter Server cloud account on which you want to install the agent, click Agent Actions/vertical ellipsis and then click **Install**. For more information, see Install an Agent from the UI.<br>**Note**<br>■ **Install** from the UI is available only for VMs managed by the vCenter Server cloud account using the product-managed agent.<br>■ For `Endpoint` object types, product-managed agent install is not supported.<br>■ **Install** from the UI is not available for open-source agents and user-managed agents. |
| **Agent Actions/Vertical Ellipsis > Uninstall** | Uninstalls the agent. Select the VM/s managed by the vCenter Server cloud account on which you want to uninstall the agent, click Agent Actions/vertical ellipsis and then click **Uninstall**. For more information, see Uninstall an Agent.<br>**Note**<br>■ Uninstall from the UI is available only for VMs managed by the vCenter Server cloud account using the product-managed agent.<br>■ For `Endpoint` objects project-managed agent uninstall is not supported.<br>■ **Uninstall** from the UI is not available for open-source agents and user-managed agents. |
| **Agent Actions Ellipsis/Vertical Ellipsis > Update** | Updates agents that are at a lower version. Select the VMs managed and unmanaged by the vCenter Server cloud account and physical servers on which you want to update the agent, click Agent Actions/vertical ellipsis and then click **Update**. After the agents are updated, the last operation status changes to **Content Upgrade Success**.<br>**Note** You can use the **Update** option only for objects with a product-managed Telegraf agent installed and running.<br>**Note** After you upgrade VMware Aria Operations and then update product-managed Telegraf agents that are running at an earlier version, there could be a data collection gap of a maximum of 2-3 collection cycles. |
| **Agent Actions/Vertical Ellipsis > Start** | If you have temporarily stopped sending metrics to VMware Aria Operations, you can use the **Start** option to start data collection for the application service. Select the VMs managed and unmanaged by the vCenter Server cloud account and physical servers on which you want to start the agent and click Agent Actions/vertical ellipsis and then click **Start**.<br>**Note** You can use the **Start** option only for objects with a product-managed Telegraf agent installed and running. |
| **Agent Actions/Vertical Ellipsis > Stop** | During a maintenance period, you can temporarily stop sending application service metrics to VMware Aria Operations. Select the VMs managed and unmanaged by the vCenter Server cloud account and physical servers on which you want to stop the agent and click Agent Actions/vertical ellipsis and then click **Stop**.<br>**Note** You can use the **Stop** option only for objects with a product-managed Telegraf agent installed and running. |

### Table 4-22. Options (continued)

| Options | Description |
| --- | --- |
| **Vertical Ellipsis > Go To Details** | Displays the **Summary** tab of the selected VMs managed and unmanaged by the vCenter Server cloud account and physical servers.<br>**Note** You can view the summary and details for objects with product-managed Telegraf agents, open-source Telegraf agents, and user-managed Telegraf agents installed and running. |
| **Agent Actions > Activate Services** | Allows you to activate a service.<br>**Note** You can only activate a service for objects with a product- managed Telegraf agent installed and running.<br>**Note** Applicable for bulk actions. |
| Filter | Filters the VMs managed and unmanaged by the vCenter Server cloud account and physical servers based on the name of the VM or the host name in the case of `Endpoint` type objects, the operating system it runs on, object type, agent type, agent status, the application service discovered, the last operation status, the power status of the VMs managed by the vCenter Server cloud account, and the agent version. |

You can also view specific details from the options in the data grid.

### Table 4-23. Data Grid Options

| Option | Description |
| --- | --- |
| Name | Name of the virtual machine or host name in case of physical servers and VMs unmanaged by the vCenter Server cloud account. |
| Object Type | There are four types of object types:<br>■ Virtual Machine: Includes only VMs managed by the vCenter Server cloud account.<br>■ Endpoint: Includes physical servers and VMs unmanaged by the vCenter Server cloud account.<br>**Note** The agent types supported by the `Endpoint` object type are: `product-managed agent`, `open-source agent`, and `user-managed agent`.<br>■ EC2 Instance<br>■ Azure Virtual Machine<br>**Note** EC2 Instance and Azure Virtual Machine object types are listed for open-source Telegraf. |
| Operating System | Operating system running on the target machine. |

### Table 4-23. Data Grid Options (continued)

| Option | Description |
| --- | --- |
| Agent Status | Displays the status of the agent at the target machine.<br>■ Agent Running. Indicates that the agent is running.<br>**Note** Applicable only for product managed Telegraf agents.<br>■ Agent Stopped. Indicates that the agent has stopped.<br>■ Agent Unhealthy: Indicates that a component of the product-managed agent is not healthy.<br>■ Not Installed. Indicates that the agent is not installed. |
| Agent Type | Types of agents:<br>■ Product-managed agent: Telegraf agent managed by VMware Aria Operations.<br>■ Open-source agent: Open source Telegraf agent.<br>■ User-Managed agent: Telegraf agent used to monitor physical servers.<br>**Note** The `User-managed agent` option is available only if you have upgraded from a previous version of VMware Aria Operations. |
| Last Operation Status | Status of the last operation.<br>**Note** Applicable only for product-managed agents.<br>The possible values are:<br>■ No Operation<br>■ Install Success<br>■ Install Failed<br>■ Install In Progress<br>■ Start Success<br>■ Start Failed<br>■ Start In Progress<br>■ Stop Success<br>■ Stop Failed<br>■ Stop In Progress<br>■ Update Success<br>■ Update Failed<br>■ Update In Progress<br>■ Uninstall Success<br>■ Uninstall Failed<br>■ Uninstall In Progress<br>■ Download Success |

### Table 4-23. Data Grid Options (continued)

| Option | Description |
| --- | --- |
| State | Power status of the VMs managed by the vCenter Server cloud account. The possible values are:<br>■ Powered On<br>■ Powered Off<br>**Note** Powered on and Powered Off status is applicable only for the `Virtual Machine` object type.<br>■ Unknown |
| Virtual IP Details | Virtual IP configured on an application monitoring high availability activated collector group to which target machines will send data. For collector groups that are not activated for application monitoring high availability, the column will be empty. |
| Collector Group | Displays the name of the collector group to which the cloud proxy mentioned in the **Cloud Proxy** column belongs. |
| Cloud Proxy | You can see either:<br>■ The cloud proxy IP through which Telegraf agent management is performed and data is collected, or<br>■ For high availability activated collector groups, the IP address of the primary cloud proxy that is collecting data is displayed. |
| Agent Version | Version of the agent. A gray dot is displayed if the target machine requires an update for product managed agents.<br>■ For product-managed agents: The version displayed is the same as the VMware Aria Operations version.<br>■ For open-source agents: The version displayed is the open source Telegraf agent version that you use. |
| vCenter Name | Name of the vCenter Adapter cloud account to which that vCenter Server VMs resource belongs.<br>**Note** The field is empty for `Endpoint`, `EC2 Instance`, and `Azure VM` object types. |
| Collection State | VMware Aria Operations Application Management cloud account collection state which monitors the VM. |
| Collection Status | VMware Aria Operations Application Management cloud account collection status which monitors the VM. |

Table 4-24. Status of Application Services

| Icon | Description |
| --- | --- |
| Green tick icon against the application service | Indicates that the application service is activated and the application service instances are receiving data. |
| Red exclamation icon against the application service | Indicates that the application service has been activated but there is a problem with data collection.<br>When there is more than one agent plugin of the same kind, and one of them is activated, but the other is not collecting data, a red exclamation icon is still displayed against the application service. |
| Red icon against the application service instance | Indicates that there is an error in receiving data with an error message displayed after the name of the application service instance. |
| Gray question icon against the application service | Indicates that the application services requires reactivation. The application service must be reactivated. For reactivation, see Activate an Application Service for more information |
| Gray pause icon against the application service instance | Indicates that the application service instance has been stopped. |
| Blue icon with three horizontal dots against the agent plugin | Indicates data is being received. |
| Progress status icon | After you have added the parameters and activated the application service, the progress status is displayed until data collection starts. |

To manage the agent, follow these steps:

1 Install the agent.

For more information, see Install an Agent from the UI.

**Note**

- You can install the agent from the UI only for VMs managed by the vCenter Server cloud account.
- To install agents on physical servers and VMs unmanaged by the vCenter Server cloud account, use the script. For more information to install product-managed Telegraf agents, see Install/Uninstall an Agent Using a Script on Linux Platformsand Install/Uninstall an Agent Using a Script on a Windows Platform . For information to install open-source agents, see Monitoring Applications using Open Source Telegraf on a Linux Platform and Monitoring Applications using Open Source Telegraf on a Windows Platform.

2 Manage the application services on each agent.

For more information, see Activate an Application Service.

**Note** You can manage application services only for objects with a product-managed Telegraf agent installed and running.

For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.

3 Stop and start the agents on the target machines.

When you stop an agent, you cannot activate or deactivate a plugin. If the target machine is powered off or if you lose connection with cloud proxy, you cannot configure or activate a plugin.

**Note** You can stop and start agents only for objects with a product-managed Telegraf agent installed and running.

4 Uninstall the agent.

For more information, see Uninstall an Agent.

**Note** You can uninstall the agent from the UI only for managed vCenter Server VMs.

5 Update agents that are at a lower version.

### What Are The Different Types of Telegraf Agents in VMware Aria Operations?

The three types of Telegraf agents are product-managed agent, open source agent, and user-managed agent in VMware Aria Operations.

- **Product-managed agent**: Managed by VMware Aria Operations. You can install and uninstall (only for VMs managed by the vCenter Server cloud account), activate application services, and manage the life cycle of the Telegraf agents such as upgrade the agent, start the agent, stop the agent, and so on. For more details, see Additional Operations from the Manage Telegraf Agents Page .
- **Open source agent**: Refers to the open source Telegraf agent that you use to monitor curated and non-curated application services completely by yourself. You can install the open source Telegraf agent using a script and also configure it using the helper script to send data to VMware Aria Operations. For more information, see Monitoring Applications using Open Source Telegraf on a Linux Platform and Monitoring Applications using Open Source Telegraf on a Windows Platform. You cannot manage the life-cycle of open source telegraf agents in VMware Aria Operations, such as activate an application service, and stop/start upgrade. You can only view configured application instances and their data collection. Agent management is carried out on the target machine by the user.

- **User-managed agent**: User-managed agents were used to monitor physical servers in prior releases. Please note the following:
  - Available only if you upgrade from previous versions of VMware Aria Operations. If you want to continue with the existing user-managed agents you can only view the metrics collected and you do not have to make any changes in VMware Aria Operations.
  - Deprecated for fresh install. You will not be able to install new user-managed agents.
  - As a first time user if you want to monitor physical servers, use the product-managed Telegraf agent and use the helper script for installation of the Telegraf agent. For more information, see Install/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform

### Custom Script

You can run custom scripts in the target machine and collect custom data which can then be consumed as a metric.

#### Prerequisites

- All the scripts that you run using the custom script, must output a single integer value. If the output is not a single integer value, an error is displayed in the user interface.
- The custom script uses Telegraf's `exec` plugin to run scripts on a target machine's operating system. In Linux operating systems, a special user called *arcuser* with specific privileges, is created for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that *arcuser* user. Ensure that the *arcuser* can run the scripts that use the custom script (the *arcuser* must have permissions to run the script). For example, the *arcuser* created automatically by cloud proxy, does not have privileges to run scripts which are stored under the `/root` directory.
- In Windows operating systems, a system user is used for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that system user. Ensure that the system user has privileges and can run the custom script.
- The script must be placed in the `/opt/vmware` folder.
- The environment variable PATH must be set for PowerShell scripts to be executed on the Windows operating system if the user used the *powershell -File* prefix. For example, *"%SYSTEMROOT%\System32\WindowsPowerShell\v1.0\"* in the environment variable PATH variable.

#### How to Run Custom Scripts

- From the **Manage Telegraf Agents** tab, filter by **Agent Status > Agent Running**.
- Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Custom Script** option, click the vertical ellipsis and then click **Add**.

- From the **Manage Custom Service** dialog box, you can add and configure the Windows services to be monitored.

### Instance Settings

| Option | Description |
|---|---|
| Status | Activate the custom script execution. |
| Display Name | Add a suitable name for the script. The `*` is an invalid character and must not be used in the name. |
| Filepath | Enter the path to the script file on the target machine.<br>Example:<br>For Linux machines:<br>`/opt/vmware/scripts/customscript.py`<br>For Windows machines:<br>`c:\scripts\customscript.ps1` |
| Prefix | Enter a prefix if necessary.<br>Example:<br>For Linux machines: python2.7, /bin/bash, or perl, etc<br>For Windows machines: *powershell -File* |
| Args | List the arguments in the script. |
| Timeout | Enter a script execution timeout on the VM. |

After you save the script, it appears under **Custom Script**. You can edit or delete scripts by clicking the **Edit** or **Delete** options from the vertical ellipsis against the custom script you added. After the scripts have been added and saved, click the drop down arrow against **Custom Script**, to view the list of scripts and their status.

**Note**

- The custom script must throw all errors in the format `ERROR|<Error_message>` for the error propagation to work. If the script does not throw an error in the given format, VMware Aria Operations displays an error message `Unable to parse the error message. Please check the endpoint` in the user interface. This is by design, until cloud proxy propagates the exact error message.
- The bash script must start with `shebang (#!/bin/bash)`.

### All Metrics Tab

When data is collected successfully, you can view the script as a metric for the target machine, in the **All Metrics** tab. The script metrics are created under an object called `Custom Script` which is a single object per target machine. All the metrics from the scripts for the target machine are placed under that `Custom Script` object that contains all the custom scripts you have created.

You can view the output for the specific metric. The metric name under the `Scripts` folder is the display name that the user specifies while creating the script configuration. For example, if you set the display name as **`Python script`**, then a metric is created with the name **Python script** if data is collected successfully.

### Activate Remote Checks

You can activate remote checks such as ICMP Check, UDP Check, TCP Check, and HTTP Check.

- ICMP check sends a ping from the target machine to a specified resource.
- UDP check performs a UDP based connection check from the target machine to a specified resource.
- TCP check performs a connection check from the target machine to the port of a specified resource.
- HTTP check performs REST API calls from the target machine to a specified URL.

#### Procedure

1. From the **Manage Telegraf Agents** tab, filter by **Agent Status > Agent Running**.
2. Expand the drop-down arrow against the relevant target machine on which the agent is installed. You see the **Custom Monitoring** section.
3. Against the **ICMP Check**, **UDP Check**, **TCP Check**, or **HTTP Check** options, click the vertical ellipsis and then click **Add**.
4. From the dialog box that appears on the right side, you can activate and configure the remote checks to be monitored.

   For configuration information, see Configuring Supported Application Services.
5. Click **Save**.

### Monitor Windows Services

After you install an agent on a target machine, you can monitor existing or custom Windows services that run on the target machine.

#### How to Monitor Windows Services

- From the **Manage Telegraf Agents** tab, filter by **Agent Status > Agent Running**.
- From the data grid, expand the drop-down arrow in front of the relevant Windows target machine name on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Services** option, click the vertical ellipsis and then click **Add**.

- From the **Manage Service Activation** dialog box, you can add and configure the Windows services to be monitored.

  Table 4-25. Instance Settings and Other Options

| Option | Descriptions |
| --- | --- |
| Status | Activate the monitoring of the Windows service. |
| Display Name | Add a suitable name for the Windows service.<br>For new plugin activations, the VM name or the end point host name in the format <on *VM name*>, is automatically appended to the display name. For example, if the display name you enter is <system>, the VM name is automatically appended and the name is displayed as <system on *VM name*>.<br>If the display name was <system on abcd>, after an upgrade, <abcd> is replaced with the VM name or end point host name.<br>If the display name did not end with <on *text1*>, after an upgrade, <on *VM name*> is automatically appended to the existing display name.<br>The following are invalid characters and must not be used in the name: `<`, `"`, `>`, and `\|`. |
| Service Name | Enter a name of the Windows service you want to monitor. |

Save the settings to add the Windows service. To edit or delete Windows services, click the **Edit** or **Delete** options from the vertical ellipsis against Windows service you added. After the services have been added and saved, click the drop down arrow against **Services**, to view the list of Windows services and their status.

### Metrics Tab

When data is collected successfully, you can view the metric from the **Manage Telegraf Agents** page, click the **Vertical Ellipsis > Go To Details > Metrics** tab. The metrics for the Windows service are created under an object called `Services` which is a single object per target machine.

## Monitor Linux Processes

After you install an agent on a target machine, you can monitor existing or custom Linux processes that run on the machine.

### How to Monitor Linux Services

- From the **Manage Telegraf Agents** tab, filter by **Agent Status > Agent Running**.
- Expand the drop-down arrow against the relevant Linux machine on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Services** option, click the vertical ellipsis and then click **Add**.
- From the **Manage Service Activation** dialog box, you can add and configure the Linux services to be monitored.

Table 4-26. Instance Settings and Other Options

| Option | Description |
| --- | --- |
| Status | Activate or deactivate the monitoring of the Linux process. |
| Display Name | Add a suitable name for the Linux process you want to monitor.<br>For new plugin activations, the VM name or end point host name in the format <on *VM name*>, is automatically appended to the display name. For example, if the display name you enter is <system>, the VM name is automatically appended and the name is displayed as <system on *VM name*>.<br>If the display name was <system on abcd>, after an upgrade, <abcd> is replaced with the VM name or end point host name.<br>If the display name did not end with <on *text1*>, after an upgrade, <on *VM name*> is automatically appended to the existing display name.<br>The following are invalid characters and must not be used in the name: `<`, `"`, `>`, and `\|`. |
| Filter Type | Select either **Executable Name**, **Regex Pattern**, or **Pid File** as the filter type from the drop-down menu. |
| Filter Value | The filter value could be a process executable name, a regex pattern, or a pid file absolute path. |

Save the settings to add the Linux service. To edit or delete Linux services, click the **Edit** or **Delete** options from the vertical ellipsis against Linux service you added. After the services have been added and saved, click the drop down arrow against **Services**, to view the list of Linux services and their status.

### Metrics Tab

When data is collected successfully, you can view the metric from the **Manage Telegraf Agents** page, click the **Vertical Ellipsis > Go To Details > Metrics** tab. The metrics for the Linux process are created under an object called `Processes` which is a single object per target machine.

## Deactivate an Application Service

You can deactivate an application service to stop monitoring the application service that is sending data to VMware Aria Operations.

### Prerequisite

- If plugin deactivation requires the location of a file (for example, client certificates for SSL Trust) on the target machine, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

  **Note** If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

### Deactivate an Application Service

To deactivate a plugin to stop monitoring the application service that is sending data to VMware Aria Operations, complete the following steps:

1 From the left menu, click **Environment > Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.

2 Filter by **Agent Status > Agent Running**.

3 Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Services Discovered** section.

4 From the **Services Discovered** section, select a service that has been activated, click the vertical ellipsis and then click **Edit**.

5 Deactivate the application service from dialog box that is displayed on the right side.

6 Click **Save**.

For information on activating an application service, see Activate an Application Service.

## Summary of Discovered and Supported Operating Systems and Application Services

You can monitor application services and operating systems from VMware Aria Operations and get insights into the services, processes and infrastructure. You can view a summary of discovered operating systems and services, supported operating systems, and supported services.

To monitor and view applications and operating systems, from the left menu, click **Configure > Applications and Services**, and then click **Application Monitoring: Telegraf** from the right panel.

### Discovered Operating Systems and Services

You see the application services that are discovered on the end point virtual machines where the agents are installed. From the **Discovered Operating Systems and Services** section in the **Monitor Applications** page, click the <discovered> link under the name of the application service to perform lifecycle management actions for the agent and application services. For more information, see Additional Operations from the Manage Telegraf Agents Page .

### Supported Operating Systems

You see a list of supported operating systems for which VMware Aria Operations collects metrics. For more information, see Supported Platforms.

### Supported Services

You see a list of supported services for which VMware Aria Operations collects metrics. For more information, see Supported Application Services .

## Metrics Collected

Metrics are collected for operating systems, application services, remote checks, Linux processes, and Windows services.

### Operating System Metrics

Metrics are collected for Linux and Windows operating systems.

#### Linux Platforms

The following metrics are collected for Linux operating systems:

Table 4-27. Metrics for Linux

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\| Usage Idle | CPU | False |
| <Instance name>\| Usage IO-Wait | CPU | False |
| <Instance name>\|Time Active | CPU | True |
| <Instance name>\|Time Guest | CPU | False |
| <Instance name>\|Time Guest Nice | CPU | False |
| <Instance name>\|Time Idle | CPU | False |
| <Instance name>\|Time IO-Wait | CPU | False |
| <Instance name>\|Time IRQ | CPU | True |
| <Instance name>\|Time Nice | CPU | False |
| <Instance name>\|Time Soft IRQ | CPU | True |
| <Instance name>\|Time Steal | CPU | False |
| <Instance name>\|Time System | CPU | False |
| <Instance name>\|Time User | CPU | True |
| <Instance name>\|Usage Active (%) | CPU | True |
| <Instance name>\|Usage Guest (%) | CPU | False |
| <Instance name>\|Usage Guest Nice (%) | CPU | False |
| <Instance name>\|Usage IRQ (%) | CPU | True |
| <Instance name>\|Usage Nice (%) | CPU | False |
| <Instance name>\|Usage Soft IRQ (%) | CPU | True |
| <Instance name>\|Usage Steal (%) | CPU | False |
| <Instance name>\|Usage System (%) | CPU | True |

### Table 4-27. Metrics for Linux (continued)

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\|Usage User (%) | CPU | True |
| CPU Load1 (%) | CPU Load | False |
| CPU Load15 (%) | CPU Load | False |
| CPU Load5 (%) | CPU Load | False |
| <Instance name>\|IO Time | Disk IO | False |
| <Instance name>\|Read Time | Disk IO | False |
| <Instance name>\|Reads | Disk IO | False |
| <Instance name>\|Write Time | Disk IO | False |
| <Instance name>\|Writes | Disk IO | False |
| <Instance name>\|Disk Free | Disk | False |
| <Instance name>\|Disk Total | Disk | False |
| <Instance name>\|Disk Used (%) | Disk | False |
| Cached | Memory | False |
| Free | Memory | False |
| Inactive | Memory | False |
| Total | Memory | True |
| Used | Memory | True |
| Used Percent | Memory | True |
| Blocked | Processes | True |
| Dead | Processes | False |
| Running | Processes | False |
| Sleeping | Processes | False |
| Stopped | Processes | False |
| Zombies | Processes | False |
| Free | Swap | False |
| In | Swap | False |
| Out | Swap | False |
| Total | Swap | True |

### Table 4-27. Metrics for Linux (continued)

| Metric | Metric Category | KPI |
|---|---|---|
| Used | Swap | True |
| Used Percent | Swap | True |
| Telegraf Availability | None | False |

## Windows Platforms

The following metrics are collected for Windows operating systems:

### Table 4-28. Metrics for Windows

| Metric | Metric Category | KPI |
|---|---|---|
| Idle Time | CPU | False |
| Interrupt Time | CPU | False |
| Interrupts persec | CPU | True |
| Privileged Time | CPU | False |
| Processor Time | CPU | False |
| User Time | CPU | False |
| DPC Time (%) | CPU | False |
| Usage Guest (%) | CPU | False |
| Usage System (%) | CPU | False |
| Usage User (%) | CPU | False |
| Avg. Disk Bytes Read | Disk | False |
| Avg. Disk sec Read | Disk | False |
| Avg. Disk sec Write | Disk | False |
| Avg. Disk Write Queue Length | Disk | False |
| Avg. Disk Read Queue Length | Disk | False |
| Disk Read Time | Disk | False |
| Disk Write Time | Disk | False |
| Free Megabytes | Disk | False |
| Free Space | Disk | False |
| Idle Time | Disk | False |
| Split IO persec | Disk | False |

Table 4-28. Metrics for Windows (continued)

| Metric | Metric Category | KPI |
| --- | --- | --- |
| Available Bytes | Memory | True |
| Cache Bytes | Memory | False |
| Cache Faults persec | Memory | False |
| Committed Bytes | Memory | True |
| Demand Zero Faults persec | Memory | False |
| Page Faults persec | Memory | True |
| Pages persec | Memory | False |
| Pool Nonpaged Bytes | Memory | True |
| Pool Paged Bytes | Memory | False |
| Transition Faults persec | Memory | False |
| Total (bytes) | Memory | False |
| Used (bytes) | Memory | False |
| Used Percent(%) | Memory | False |
| Bytes Received persec | Network | False |
| Bytes Sent persec | Network | False |
| Packets Outbound Discarded | Network | False |
| Packets Outbound Errors | Network | False |
| Packets Received Discarded | Network | False |
| Packets Received Errors | Network | False |
| Packets Received persec | Network | False |
| Packets Sent persec | Network | False |
| Elapsed Time | Process | False |
| Handle Count | Process | False |
| IO Read Bytes persec | Process | False |
| IO Read Operations persec | Process | False |
| IO Write Bytes persec | Process | False |
| IO Write Operations persec | Process | False |
| Privileged Time | Process | False |

### Table 4-28. Metrics for Windows (continued)

| Metric | Metric Category | KPI |
|---|---|---|
| Processor Time | Process | False |
| Thread Count | Process | False |
| User Time | Process | False |
| Context Switches persec | System | False |
| Processes | System | False |
| Processor Queue Length | System | False |
| System Calls persec | System | False |
| System Up Time | System | False |
| Threads | System | False |
| Used Percent (%) | Swap | False |
| Total (bytes) | Swap | False |
| Telegraf Availability | None | False |

## Application Service Metrics

Metrics are collected for 23 application services.

### Active Directory Metrics

Metrics are collected for the Active Directory application service.

### Table 4-29. Active Directory Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Database Cache % Hit (%) | Active Directory Database | True |
| Database Cache Page Faults/sec | Active Directory Database | True |
| Database Cache Size | Active Directory Database | False |
| Data Lookups | Active Directory DFS Replication | False |
| Database Commits | Active Directory DFS Replication | True |
| Avg Response Time | Active Directory DFSN | True |
| Requests Failed | Active Directory DFSN | False |
| Requests Processed | Active Directory DFSN | False |
| Dynamic Update Received | Active Directory DNS | False |
| Dynamic Update Rejected | Active Directory DNS | False |

## Table 4-29. Active Directory Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Recursive Queries | Active Directory DNS | False |
| Recursive Queries Failure | Active Directory DNS | False |
| Secure Update Failure | Active Directory DNS | False |
| Total Query Received | Active Directory DNS | True |
| Total Response Sent | Active Directory DNS | True |
| Digest Authentications | Active Directory Security System-Wide Statistics | True |
| Kerberos Authentications | Active Directory Security System-Wide Statistics | True |
| NTLM Authentications | Active Directory Security System-Wide Statistics | True |
| Directory Services:<InstanceName>\|Base Searches persec | Active Directory Services | False |
| Directory Services:<InstanceName>\|Database adds persec | Active Directory Services | False |
| Directory Services:<InstanceName>\|Database deletes persec | Active Directory Services | False |
| Directory Services<InstanceName>\|Database modifys/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|Database recycles/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Inbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Inbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Outbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Outbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Pending Replication Operations | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Pending Replication Synchronizations | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Sync Requests Made | Active Directory Services | False |
| Directory Services<InstanceName>\|DRA Sync Requests Successful | Active Directory Services | False |

### Table 4-29. Active Directory Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Directory Services<InstanceName>\|DS Client Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\|DS Directory Reads/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|DS Directory Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\|DS Server Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\|DS Threads in Use | Active Directory Services | True |
| Directory Services:<InstanceName>\|LDAP Active Threads | Active Directory Services | False |
| Directory Services:<InstanceName>\|LDAP Client Sessions | Active Directory Services | True |
| Directory Services<InstanceName>\|LDAP Closed Connections/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|LDAP New Connections/sec | Active Directory Services | True |
| Directory Services<InstanceName>\|LDAP Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\|LDAP Successful Binds/sec | Active Directory Services | False |
| Directory Services<InstanceName>\|LDAP UDP operations/sec | Active Directory Services | False |
| Directory Services:<InstanceName>\|LDAP Writes/sec | Active Directory Services | False |
| Application Availability | Active Directory | False |

## Apache HTTPD

Metrics are collected for the Apache HTTPD application service.

### Table 4-30. Apache Tomcat

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Tomcat Server | False |
| Class Loading\|Loaded Class Count | Tomcat Server | False |
| Class Loading\|Total Loaded Class Count | Tomcat Server | False |

## Table 4-30. Apache Tomcat (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Class Loading\|Unloaded Class Count | Tomcat Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Tomcat Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Time | Tomcat Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Number of Object Pending Finalization Count | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Tomcat Server | False |

## Table 4-30. Apache Tomcat (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| JVM Memory\|Pool:<InstanceName>\|Usage\|Used Memory | Tomcat Server | False |
| Process CPU Usage (%) | Tomcat Server | True |
| System CPU Usage (%) | Tomcat Server | True |
| System Load Average (%) | Tomcat Server | True |
| Threading\|Thread Count | Tomcat Server | False |
| Uptime | Tomcat Server | True |
| Application Availability | Tomcat Server | False |
| JSP Count | Tomcat Server Web Module | False |
| JSP Reload Count | Tomcat Server Web Module | False |
| JSP Unload Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Error Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Processing Time | Tomcat Server Web Module | False |
| Cache : Hit Count | Tomcat Server Web Module | False |
| Cache : Lookup Count | Tomcat Server Web Module | False |
| Current Thread Count | Tomcat Server Global Request Processor | True |
| Current Threads Busy | Tomcat Server Global Request Processor | True |
| errorRate | Tomcat Server Global Request Processor | False |
| Total Request Bytes Received | Tomcat Server Global Request Processor | False |
| Total Request Bytes Sent | Tomcat Server Global Request Processor | False |
| Total Request Count | Tomcat Server Global Request Processor | True |
| Total Request Error Count | Tomcat Server Global Request Processor | True |
| Total Request Processing Time | Tomcat Server Global Request Processor | False |

## Microsoft SQL Server Metrics

Metrics are collected for the Microsoft SQL Server application service.

### Table 4-31. MS SQL Metrics

| Metric Name | Category | KPI |
|---|---|---|
| CPU<InstanceName>\|CPU Usage (%) | Microsoft SQL Server | False |
| Database IO\|Rows Reads Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Reads/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes/Sec | Microsoft SQL Server | False |
| Performance\|Access Methods\|Full Scans per second | Microsoft SQL Server | False |
| Performance\|Access Methods\|Index Searches | Microsoft SQL Server | False |
| Performance\|Access Methods\|Page Splits per second | Microsoft SQL Server | False |
| Performance\|Broker Activation\| Stored Procedures Invoked per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Buffer cache hit ratio (%) | Microsoft SQL Server | True |
| Performance\|Buffer Manager\| Checkpoint Pages/sec | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Lazy writes per second | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page life expectancy | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page lookups per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page reads per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page writes per second | Microsoft SQL Server | False |
| Performance\|Databases\|Active Transactions | Microsoft SQL Server | True |
| Performance\|Databases\|Data File(s) Size | Microsoft SQL Server | True |
| Performance\|Databases\|Log Bytes Flushed/Sec | Microsoft SQL Server | False |
| Performance\|Databases\|Log File(s) Size | Microsoft SQL Server | False |

## Table 4-31. MS SQL Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Performance\|Databases\|Log File(s) Used Size | Microsoft SQL Server | False |
| Performance\|Databases\|Log Flush Wait Time | Microsoft SQL Server | False |
| Performance\|Databases\|Log Flushes per second | Microsoft SQL Server | False |
| Performance\|Databases\|Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|Write Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|XTP Memory Used | Microsoft SQL Server | False |
| Performance\|General Statistics\|Active temp Tables | Microsoft SQL Server | False |
| Performance\|General Statistics\|Logins per second | Microsoft SQL Server | False |
| Performance\|General Statistics\| Logouts per second | Microsoft SQL Server | False |
| Performance\|General Statistics\| Processes Blocked | Microsoft SQL Server | False |
| Performance\|General Statistics\|Temp Tables Creation Rate | Microsoft SQL Server | False |
| Performance\|General Statistics\|User Connections | Microsoft SQL Server | False |
| Performance\|Locks\|Average Wait Time | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Requests per second | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Wait Time | Microsoft SQL Server | True |
| Performance\|Locks\|Lock Waits per second | Microsoft SQL Server | True |
| Performance\|Locks\|Number of Deadlocks per second | Microsoft SQL Server | True |
| Performance\|Memory Manager\| Connection Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Lock Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Log Pool Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Memory Grants Pending | Microsoft SQL Server | True |

## Table 4-31. MS SQL Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Performance\|Memory Manager\|SQL Cache Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Target Server Memory | Microsoft SQL Server | True |
| Performance\|Memory Manager\|Total Server Memory | Microsoft SQL Server | True |
| Performance\|Resource Pool Stats\| internal\|Active memory grant amount | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|CPU Usage Percentage (%) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read Bytes per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO | Microsoft SQL Server | False |
| Wait Stats:<InstanceName>\|Wait Time (ms) | Microsoft SQL Server | False |
| Wait Stats<InstanceName>\|Number of Waiting tasks (ms) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO Throttled Per Second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Write Bytes per second (Bps) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Write IO Throttled per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Used Memory | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| Batch Requests Per Second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Compilations per second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Re-Compilations per second | Microsoft SQL Server | False |
| Performance\|Transactions \| Free space in tempdb (KB) | Microsoft SQL Server | False |
| Performance\|Transactions \| Transactions | Microsoft SQL Server | False |
| Performance\|Transactions \| Version Store Size (KB) | Microsoft SQL Server | False |

### Table 4-31. MS SQL Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Performance\|User Settable Counter \| User Counter 0 to 10 | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Active Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Blocked Tasks | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|CpU Usage (%) | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Queued Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Request Completed/sec | Microsoft SQL Server | False |
| Application Availability | Microsoft SQL Server | False |

There are no metrics collected for Microsoft SQL Server Database.

**PostgreSQL**

Metrics are collected for the PostgreSQL application service.

### Table 4-32. PostgreSQL

| Metric Name | Category | KPI |
| --- | --- | --- |
| Buffers\|Buffers Allocated | PostgreSQL | False |
| Buffers\|Buffers Written by Backend | PostgreSQL | True |
| Buffers\|Buffers Written by Background Writer | PostgreSQL | True |
| Buffers\|Buffers Written During Checkpoints | PostgreSQL | True |
| Buffers\|fsync Call Executed by Backend | PostgreSQL | False |
| Checkpoints\|Checkpoints sync time | PostgreSQL | False |
| Checkpoints\|Checkpoints write time | PostgreSQL | False |
| Checkpoints\|Requested checkpoints performed count | PostgreSQL | False |
| Checkpoints\|Scheduled checkpoints performed count | PostgreSQL | False |
| Clean scan stopped count | PostgreSQL | False |
| Application Availability | PostgreSQL | False |
| Disk Blocks\|Blocks Cache Hits | PostgreSQL Database | False |

### Table 4-32. PostgreSQL (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Disk Blocks\|Blocks Read | PostgreSQL Database | False |
| Disk Blocks\|Blocks Read Time | PostgreSQL Database | False |
| Disk Blocks\|Blocks Write Time | PostgreSQL Database | False |
| Statistics\|Backends Connected | PostgreSQL Database | False |
| Statistics\|Data Written by Queries | PostgreSQL Database | True |
| Statistics\|Deadlocks Detected | PostgreSQL Database | True |
| Statistics\|Queries Cancelled | PostgreSQL Database | True |
| Statistics\|Temp Files Created by Queries | PostgreSQL Database | False |
| Transactions\|Transactions Committed | PostgreSQL Database | True |
| Transactions\|Transactions Rolled Back | PostgreSQL Database | True |
| Tuples\|Tuples Deleted | PostgreSQL Database | True |
| Tuples\|Tuples Fetched | PostgreSQL Database | True |
| Tuples\|Tuples Inserted | PostgreSQL Database | True |
| Tuples\|Tuples Returned | PostgreSQL Database | True |
| Tuples\|Tuples Updated | PostgreSQL Database | True |

## Microsoft IIS Metrics

Metrics are collected for the Microsoft IIS application service.

### Table 4-33. IIS Metrics

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Service Request Queues<InstanceName>AppPool\|CurrentQueueSize | IIS HTTP Service Request Queues | True |
| HTTP Service Request Queues<InstanceName>AppPool\|RejectedRequests | IIS HTTP Service Request Queues | False |
| Web Services<InstanceName> Web Site\|Bytes Received | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Sent/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Total/sec | IIS Web Services | False |

## Table 4-33. IIS Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Web Services<InstanceName> Web Site\|Connection Attempts/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Current Connections | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Get Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Locked Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Not Found Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Post Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Service Uptime | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Bytes Sent | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Get Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Post Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Put Requests | IIS Web Services | False |
| Current File Cache Memory Usage (bytes) | IIS Web Services Cache | False |
| File Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Misses | IIS Web Services Cache | False |
| Total Flushed URIs | IIS Web Services Cache | False |
| URI Cache Hits | IIS Web Services Cache | False |
| URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| URI Cache Misses | IIS Web Services Cache | False |
| ASP.NET<InstanceName>\|Application Restarts | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Request Wait Time | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Current | IIS ASP.NET | True |

## Table 4-33. IIS Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| ASP.NET<InstanceName>\|Requests Queued | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Rejected | IIS ASP.NET | True |
| MS.NET<InstanceName>\|Allocated Bytes/sec | MS.NET | True |
| MS.NET<InstanceName>\|Current Queue Length | MS.NET | False |
| MS.NET<InstanceName>\|Finalization Survivors | MS.NET | False |
| MS.NET<InstanceName>\|Gen 0 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 0 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 heap size | MS.NET | False |
| MS.NET<InstanceName>\|IL Bytes Jitted / sec | MS.NET | False |
| MS.NET<InstanceName>\|Induced GC | MS.NET | False |
| MS.NET<InstanceName>\|Large Object Heap size | MS.NET | False |
| MS.NET<InstanceName>\|No of current logical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current physical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|No of Exceps Thrown / sec | MS.NET | True |
| MS.NET<InstanceName>\|No of total recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|Percent Time in Jit | MS.NET | False |
| MS.NET<InstanceName>\|Pinned Objects | MS.NET | False |

### Table 4-33. IIS Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| MS.NET<InstanceName>\|Stack Walk Depth | MS.NET | False |
| MS.NET<InstanceName>\|Time in RT checks | MS.NET | False |
| MS.NET<InstanceName>\|Time Loading | MS.NET | True |
| MS.NET<InstanceName>\|Total No of Contentions | MS.NET | False |
| MS.NET<InstanceName>\|Total Runtime Checks | MS.NET | True |
| Application Availability | Microsoft IIS | False |

## MS Exchange Metrics

Metrics are collected for the MS Exchange application service.

### Table 4-34. MS Exchange Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Active Manager Server\|Active Manager Role | MS Exchange | False |
| Active Manager Server\|Database State Info Writes per second | MS Exchange | False |
| Active Manager Server\| GetServerForDatabase Server-Side Calls | MS Exchange | False |
| Active Manager Server\|Server-Side Calls per second | MS Exchange | True |
| Active Manager Server\|Total Number of Databases | MS Exchange | True |
| ActiveSync\|Average Request Time | MS Exchange | True |
| ActiveSync\|Current Requests | MS Exchange | False |
| ActiveSync\|Mailbox Search Total | MS Exchange | False |
| ActiveSync\|Ping Commands Pending | MS Exchange | False |
| ActiveSync\|Requests per second | MS Exchange | True |
| ActiveSync\|Sync Commands per second | MS Exchange | True |
| ASP.NET\|Application Restarts | MS Exchange | False |
| ASP.NET\|Request Wait Time | MS Exchange | True |
| ASP.NET\|Worker Process Restarts | MS Exchange | False |

## Table 4-34. MS Exchange Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Autodiscover Service\|Requests per second | MS Exchange | True |
| Availability Service\|Average Time to Process a Free Busy Request | MS Exchange | True |
| Outlook Web Access\|Average Search Time | MS Exchange | True |
| Outlook Web Access\|Requests per second | MS Exchange | False |
| Outlook Web Access\|Current Unique Users | MS Exchange | False |
| Application Availability | MS Exchange | False |
| Performance\|Database Cache Hit (%) | MS Exchange Database | False |
| Performance\|Database Page Fault Stalls per second | MS Exchange Database | True |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database | True |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database | True |
| Performance\|I/O Log Reads Average Latency | MS Exchange Database | False |
| Performance\|I/O Log Writes Average Latency | MS Exchange Database | False |
| Performance\|Log Record Stalls per second | MS Exchange Database | False |
| Performance\|Log Threads Waiting | MS Exchange Database | False |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database Instance | False |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database Instance | False |
| Performance\|Log Record Stalls per second | MS Exchange Database Instance | False |
| Performance\|Log Threads Waiting | MS Exchange Database Instance | False |
| Performance\|LDAP Read Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Search Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Searches Timed Out per minute | MS Exchange Domain Controller | False |
| Performance\|Long Running LDAP Operations per minute | MS Exchange Domain Controller | False |

### Table 4-34. MS Exchange Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Connection Attempts per second | MS Exchange Web Server | True |
| Performance\|Current Connections | MS Exchange Web Server | False |
| Performance\|Other Request Methods per second | MS Exchange Web Server | False |
| Process\|Handle Count | MS Exchange Windows Service | False |
| Process\|Memory Allocated | MS Exchange Windows Service | False |
| Process\|Processor Time (%) | MS Exchange Windows Service | True |
| Process\|Thread Count | MS Exchange Windows Service | False |
| Process\|Virtual Memory Used | MS Exchange Windows Service | False |
| Process\|Working Set | MS Exchange Windows Service | False |

## JBoss Server Metrics

Metrics are collected for the JBoss Server application service.

### Table 4-35. JBoss Server Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Jboss Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Jboss Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Jboss Server | False |
| Class Loading\|Loaded Class Count | Jboss Server | False |
| Class Loading\|Total Loaded Class Count | Jboss Server | False |
| Class Loading\|Unloaded Class Count | Jboss Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Jboss Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Http Listener<InstanceName>\|Error Count | Jboss Server | False |

## Table 4-35. JBoss Server Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Http Listener<InstanceName>\| Request Count | Jboss Server | False |
| Https Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Https Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Https Listener<InstanceName>\|Error Count | Jboss Server | False |
| Https Listener<InstanceName>\| Request Count | Jboss Server | False |
| Process CPU Usage (%) | Jboss Server | False |
| System CPU Usage (%) | Jboss Server | False |
| System Load Average (%) | Jboss Server | False |
| Threading\|Daemon Thread Count | Jboss Server | False |
| Threading\|Peak Thread Count | Jboss Server | False |
| Threading\|Thread Count | Jboss Server | False |
| Threading\|Total Started Thread Count | Jboss Server | False |
| Uptime | Jboss Server | False |
| UTILIZATION\|Heap Memory Usage | Jboss Server | False |
| Application Availability | Jboss Server | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Jboss JVM Garbage Collector | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Jboss JVM Garbage Collector | False |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Jboss JVM Memory | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Jboss JVM Memory | False |

### Table 4-35. JBoss Server Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Jboss JVM Memory | False |
| JVM Memory\|Object Pending Finalization Count | Jboss JVM Memory | True |
| UTILIZATION\|Active Count | Jboss Datasource Pool | False |
| UTILIZATION\|Available Count | Jboss Datasource Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Used Memory | Jboss JVM Memory Pool | False |

## RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

### Table 4-36. RabbitMQ Metrics

| Metric Name | Category | KPI |
|---|---|---|
| CPU\|Limit | RabbitMQ | False |
| CPU\|Used | RabbitMQ | True |

## Table 4-36. RabbitMQ Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Disk\|Free | RabbitMQ | False |
| Disk\|Free limit | RabbitMQ | False |
| FileDescriptor\|Total | RabbitMQ | False |
| FileDescriptor\|Used | RabbitMQ | False |
| Memory\|Limit | RabbitMQ | False |
| Memory\|Used | RabbitMQ | True |
| Messages\|Acked | RabbitMQ | False |
| Messages\|Delivered | RabbitMQ | False |
| Messages\|Delivered get | RabbitMQ | False |
| Messages\|Published | RabbitMQ | False |
| Messages\|Ready | RabbitMQ | False |
| Messages\|Unacked | RabbitMQ | False |
| Socket\|Limit | RabbitMQ | False |
| Socket\|Used | RabbitMQ | True |
| UTILIZATION\|Channels | RabbitMQ | True |
| UTILIZATION\|Connections | RabbitMQ | True |
| UTILIZATION\|Consumers | RabbitMQ | True |
| UTILIZATION\|Exchanges | RabbitMQ | True |
| UTILIZATION\|Messages | RabbitMQ | True |
| UTILIZATION\|Queues | RabbitMQ | True |
| Application Availability | RabbitMQ | False |
| Messages\|Publish in | RabbitMQ Exchange | False |
| Messages\|Publish out | RabbitMQ Exchange | False |
| Consumer Utilisation | RabbitMQ Queue | False |
| Consumers | RabbitMQ Queue | False |
| Memory | RabbitMQ Queue | False |
| Messages\|Ack | RabbitMQ Queue | False |
| Messages\|Ack rate | RabbitMQ Queue | False |

### Table 4-36. RabbitMQ Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Messages\|Deliver | RabbitMQ Queue | False |
| Messages\|Deliver get | RabbitMQ Queue | False |
| Messages\|Persist | RabbitMQ Queue | False |
| Messages\|Publish | RabbitMQ Queue | False |
| Messages\|Publish rate | RabbitMQ Queue | False |
| Messages\|Ram | RabbitMQ Queue | False |
| Messages\|Ready | RabbitMQ Queue | False |
| Messages\|Redeliver | RabbitMQ Queue | False |
| Messages\|Redeliver rate | RabbitMQ Queue | False |
| Messages\|Space | RabbitMQ Queue | False |
| Messages\|Unack | RabbitMQ Queue | False |
| Messages\|Unacked | RabbitMQ Queue | False |
| Messages | RabbitMQ Queue | False |

There are no metrics collected for RabbitMQ Virtual Host.

## MySQL Metrics

Metrics are collected for the MySQL application service.

### Table 4-37. MySQL Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Aborted connection count | MySQL | True |
| Connection count | MySQL | True |
| Event wait average time | MySQL | False |
| Event wait count | MySQL | False |
| Binary Files\|Binary Files Count | MySQL | False |
| Binary Files\|Binary Size Bytes | MySQL | False |
| Global Status\|Aborted Clients | MySQL | False |
| Global Status\|Binlog Cache Disk Use | MySQL | False |
| Global Status\|Bytes Received | MySQL | False |
| Global Status\|Bytes Sent | MySQL | False |

## Table 4-37. MySQL Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Global Status\|Connection Errors Accept | MySQL | False |
| Global Status\|Connection Errors Internal | MySQL | False |
| Global Status\|Connection Errors Max Connections | MySQL | False |
| Global Status\|Queries | MySQL | False |
| Global Status\|Threads Cached | MySQL | False |
| Global Status\|Threads Connected | MySQL | False |
| Global Status\|Threads Running | MySQL | False |
| Global Status\|Uptime | MySQL | False |
| Global Variables\|Delayed Insert Limit | MySQL | False |
| Global Variables\|Delayed Insert Timeout | MySQL | False |
| Global Variables\|Delayed Queue Size | MySQL | False |
| Global Variables\|Max Connect Errors | MySQL | False |
| Global Variables\|Max Connections | MySQL | False |
| Global Variables\|Max Delayed Threads | MySQL | False |
| Global Variables\|Max Error Count | MySQL | False |
| InnoDB\|All deadlock count | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Dirty | MySQL | False |
| InnoDB\|Buffer Pool Dump Status | MySQL | False |
| InnoDB\|Buffer Pool Load Status | MySQL | False |
| InnoDB\|Buffer Pool Pages Data | MySQL | False |
| InnoDB\|Buffer Pool Pages Dirty | MySQL | False |
| InnoDB\|Buffer Pool Pages Flushed | MySQL | False |
| InnoDB\|Buffer pool size | MySQL | True |
| InnoDB\|Checksums | MySQL | False |

### Table 4-37. MySQL Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| InnoDB\|Open file count | MySQL | False |
| InnoDB\|Row lock average time | MySQL | False |
| InnoDB\|Row lock current waits | MySQL | False |
| InnoDB\|Row lock maximum time | MySQL | False |
| InnoDB\|Row lock time | MySQL | False |
| InnoDB\|Row lock waits | MySQL | True |
| InnoDB\|Table lock count | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Delete | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Fetch | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Insert | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Update | MySQL | False |
| Process List\|Connections | MySQL | False |
| Application Availability | MySQL | False |
| IO waits average time | MySQL Database | False |
| IO waits count | MySQL Database | True |
| Read high priority average time | MySQL Database | False |
| Read high priority count | MySQL Database | False |
| Write concurrent insert average time | MySQL Database | False |
| Write concurrent insert count | MySQL Database | False |

## NGINX Metrics

Metrics are collected for the NGINX application service.

### Table 4-38. NGINX Metrics

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Status Info\|Accepts | Nginx | True |
| HTTP Status Info\|Active connections | Nginx | False |
| HTTP Status Info\|Handled | Nginx | True |
| HTTP Status Info\|Reading | Nginx | False |

### Table 4-38. NGINX Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Status Info\|Requests | Nginx | False |
| HTTP Status Info\|Waiting | Nginx | True |
| HTTP Status Info\|Writing | Nginx | False |
| Application Availability | Nginx | False |

## SharePoint Metrics

Metrics are collected for the SharePoint Server application service.

### Table 4-39. SharePoint Server Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Sharepoint Foundation\|Active Threads | SharePoint Server | True |
| Sharepoint Foundation\|Current Page Requests | SharePoint Server | False |
| Sharepoint Foundation\|Executing SQL Queries | SharePoint Server | False |
| Sharepoint Foundation\|Executing Time/Page Request | SharePoint Server | True |
| Sharepoint Foundation\|Incoming Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Object Cache Hit Count | SharePoint Server | False |
| Sharepoint Foundation\|Reject Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Responded Page Requests Rate | SharePoint Server | True |
| SQL query executing time | SharePoint Server | False |
| Application Availability | SharePoint Server | False |
| Network\|Received Data Rate | SharePoint Web Server | True |
| Network\|Sent Data Rate | SharePoint Web Server | True |
| Process\|Processor Time (%) | SharePoint Windows Service | False |
| Process\|Threads | SharePoint Windows Service | False |

## Oracle WebLogic Server Metrics

Metrics are collected for the Oracle WebLogic Server application service.

### Table 4-40. Oracle WebLogic Server Metrics

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Process Cpu Load | Oracle WebLogic Server | True |
| UTILIZATION\|System Cpu Load | Oracle WebLogic Server | False |
| UTILIZATION\|System Load Average | Oracle WebLogic Server | False |
| Application Availability | Oracle WebLogic Server | False |
| UTILIZATION\|Collection Time | Weblogic Garbage Collector | True |
| UTILIZATION\|Connections HighCount | Weblogic JMS Runtime | True |
| UTILIZATION\|JMS Servers TotalCount | Weblogic JMS Runtime | False |
| UTILIZATION\|Active Total Count Used | Weblogic JTA Runtime | False |
| UTILIZATION\|Active Transactions TotalCount | Weblogic JTA Runtime | False |
| UTILIZATION\|Transaction Abandoned TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Transaction RolledBack App TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Heap Memory Usage | Weblogic JVM Memory | True |
| UTILIZATION\|Non Heap Memory Usage | Weblogic JVM Memory | False |
| UTILIZATION\|Peak Usage | Weblogic JVM Memory Pool | True |
| UTILIZATION\|Usage | Weblogic JVM Memory Pool | False |
| UTILIZATION\|UpTime | Weblogic JVM Runtime | False |

## Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

### Table 4-41. Pivotal TC Server Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Pivotal TC Server | False |
| Class Loading\|Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Total Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Unloaded Class Count | Pivotal TC Server | False |

## Table 4-41. Pivotal TC Server Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| File Descriptor Usage\|Max File Descriptor Count | Pivotal TC Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Time | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Pivotal TC Server | True |
| JVM Memory\|Number of Object Pending Finalization Count | Pivotal TC Server | True |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Pivotal TC Server | False |

### Table 4-41. Pivotal TC Server Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Pool:<InstanceName>\|Usage\|Used Memory | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| System CPU Usage (%) | Pivotal TC Server | True |
| Uptime | Pivotal TC Server | True |
| Threading\|Thread Count | Pivotal TC Server | False |
| System Load Average | Pivotal TC Server | False |
| Application Availability | Pivotal TC Server | False |
| Current Thread Count | Pivotal TC Server Thread Pool | False |
| Current Threads Busy | Pivotal TC Server Thread Pool | True |
| Total Request Bytes Received | Pivotal TC Server Thread Pool | False |
| Total Request Bytes Sent | Pivotal TC Server Thread Pool | False |
| Total Request Count | Pivotal TC Server Thread Pool | True |
| Total Request Error Count | Pivotal TC Server Thread Pool | True |
| Total Request Processing Time | Pivotal TC Server Thread Pool | True |
| JSP Count | Pivotal TC Server Web Module | False |
| JSP Reload Count | Pivotal TC Server Web Module | False |
| JSP Unload Count | Pivotal TC Server Web Module | False |

## ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

### Table 4-42. ActiveMQ Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Active MQ | False |
| Buffer Pool<InstanceName>\|Memory Used | Active MQ | False |
| Buffer Pool<InstanceName>\|Total Capacity | Active MQ | False |
| Class Loading\|Loaded Class Count | Active MQ | False |

## Table 4-42. ActiveMQ Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Class Loading\|Unloaded Class Count | Active MQ | False |
| Class Loading\|Total Loaded Class Count | Active MQ | False |
| File Descriptor Usage\|Max File Descriptor Count | Active MQ | False |
| File Descriptor Usage\|Open File Descriptor Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Committed Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Used Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Committed Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Used Memory | Active MQ | False |
| Application Availability | Active MQ | False |
| Threading\|Thread Count | Active MQ | False |
| Uptime | Active MQ | False |

## Table 4-42. ActiveMQ Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| UTILIZATION\|Process CpuLoad | Active MQ | False |
| UTILIZATION\|Memory Limit | ActiveMQ Broker | True |
| UTILIZATION\|Memory Percent Usage (%) | ActiveMQ Broker | True |
| UTILIZATION\|Store Limit | ActiveMQ Broker | False |
| UTILIZATION\|Store Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Temp Limit | ActiveMQ Broker | False |
| UTILIZATION\|Temp Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Total Consumer Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Dequeue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Enqueue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Message Count | ActiveMQ Broker | True |
| JVM Memory\|Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |

### Table 4-42. ActiveMQ Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Object Pending FinalizationCount | ActiveMQ JVM Memory Usage | False |
| UTILIZATION\|Process CpuLoad | ActiveMQ OS | False |
| UTILIZATION\|System Cpu Load | ActiveMQ OS | False |
| UTILIZATION\|Consumer Count | ActiveMQ Topic | True |
| UTILIZATION\|Dequeue Count | ActiveMQ Topic | True |
| UTILIZATION\|Enqueue Count | ActiveMQ Topic | True |
| UTILIZATION\|Queue Size | ActiveMQ Topic | True |
| UTILIZATION\|Producer Count | ActiveMQ Topic | False |

## Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

**Note** Metrics are collected for the Events MPM. Metrics are not collected for the other MPMs.

### Table 4-43. Apache HTTPD Metrics

| Metric Name | Category | KPI |
| --- | --- | --- |
| UTILIZATION\|Busy Workers | Apache HTTPD | True |
| UTILIZATION\|Bytes Per Req | Apache HTTPD | False |
| UTILIZATION\|Bytes Per Sec | Apache HTTPD | False |
| UTILIZATION\|CPU Load | Apache HTTPD | True |
| UTILIZATION\|CPU User | Apache HTTPD | False |
| UTILIZATION\|Idle Workers | Apache HTTPD | True |
| UTILIZATION\|Request Per Sec | Apache HTTPD | True |
| UTILIZATION\|SCBoard Closing | Apache HTTPD | False |
| UTILIZATION\|SCBoard DNS Lookup | Apache HTTPD | False |
| UTILIZATION\|SCBoard Finishing | Apache HTTPD | False |
| UTILIZATION\|SCBoard Idle Cleanup | Apache HTTPD | False |

### Table 4-43. Apache HTTPD Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| UTILIZATION\|SCBoard Keep Alive | Apache HTTPD | False |
| UTILIZATION\|SCBoard Logging | Apache HTTPD | False |
| UTILIZATION\|SCBoard Open | Apache HTTPD | False |
| UTILIZATION\|SCBoard Reading | Apache HTTPD | False |
| UTILIZATION\|SCBoard Sending | Apache HTTPD | False |
| UTILIZATION\|SCBoard Starting | Apache HTTPD | False |
| UTILIZATION\|SCBoard Waiting | Apache HTTPD | False |
| UTILIZATION\|Total Accesses | Apache HTTPD | False |
| UTILIZATION\|Total Bytes | Apache HTTPD | True |
| UTILIZATION\|Total Connections | Apache HTTPD | False |
| UTILIZATION\|Uptime | Apache HTTPD | True |
| UTILIZATION\|Asynchronous Closing Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Keep Alive Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Writing Connections | Apache HTTPD | False |
| UTILIZATION\|ServerUptimeSeconds | Apache HTTPD | False |
| UTILIZATION\|Load1 | Apache HTTPD | False |
| UTILIZATION\|Load5 | Apache HTTPD | False |
| UTILIZATION\| ParentServerConfigGeneration | Apache HTTPD | False |
| UTILIZATION\| ParentServerMPMGeneration | Apache HTTPD | False |
| Application Availability | Apache HTTPD | False |

## Oracle DB Metrics

Metrics are collected for the Oracle DB application service.

Oracle DB cannot be activated on Linux platforms.

### Table 4-44. Oracle DB Metrics

| Metric Name | Category | KPI |
| --- | --- | --- |
| Utilization\|Active Sessions | OracleDB | True |
| Utilization\|Buffer CacheHit Ratio | OracleDB | False |

## Table 4-44. Oracle DB Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Utilization\|Cursor CacheHit Ratio | OracleDB | False |
| Utilization\|Database Wait Time | OracleDB | False |
| Utilization\|Disk Sort persec | OracleDB | False |
| Utilization\|Enqueue Timeouts Persec | OracleDB | False |
| Utilization\|Global Cache Blocks Corrupted | OracleDB | False |
| Utilization\|Global Cache Blocks Lost | OracleDB | False |
| Utilization\|Library CacheHit Ratio | OracleDB | False |
| Utilization\|Logon persec | OracleDB | True |
| Utilization\|Memory Sorts Ratio | OracleDB | True |
| Utilization\|Rows persort | OracleDB | False |
| Utilization\|Service Response Time | OracleDB | False |
| Utilization\|Session Count | OracleDB | True |
| Utilization\|Session Limit | OracleDB | False |
| Utilization\|Shared Pool Free | OracleDB | False |
| Utilization\|Temp Space Used | OracleDB | False |
| Utilization\|Total Sorts persec | OracleDB | False |
| Utilization\|Physical Read Bytes Persc | OracleDB | False |
| Utilization\|Physical Read IO Requests Persc | OracleDB | False |
| Utilization\|Physical Read Total Bytes Persec | OracleDB | False |
| Utilization\|Physical Reads Persec | OracleDB | True |
| Utilization\|Physical Reads Per Txn | OracleDB | False |
| Utilization\|Physical Write Bytes Persc | OracleDB | False |
| Utilization\|Physical Write IO Requests Persc | OracleDB | False |
| Utilization\|Physical Write Total Bytes Persc | OracleDB | False |
| Utilization\|Physical Writes Persc | OracleDB | True |
| Utilization\|Physical Writes Per Txn | OracleDB | False |

### Table 4-44. Oracle DB Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Utilization\|User Commits Percentage | OracleDB | False |
| Utilization\|User Commits Persc | OracleDB | False |
| Utilization\|User Rollbacks Percentage | OracleDB | False |
| Utilization\|User Rollbacks persec | OracleDB | True |
| Utilization\|User Transaction Persec | OracleDB | False |
| Utilization\|Database Time Persc | OracleDB | False |
| Application Availability | Oracle DB | False |

## Cassandra Metrics

Metrics are collected for the Cassandra application service.

### Table 4-45. Cassandra Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Cache<InstanceName>\|Capacity | Cassandra | False |
| Cache<InstanceName>\|Entries | Cassandra | True |
| Cache<InstanceName>\|HitRate | Cassandra | True |
| Cache<InstanceName>\|Requests | Cassandra | True |
| Cache<InstanceName>\|Size | Cassandra | False |
| ClientRequest<InstanceName>\|Failures | Cassandra | False |
| ClientRequest<InstanceName>\|Latency | Cassandra | False |
| ClientRequest<InstanceName>\|Timeouts | Cassandra | False |
| ClientRequest<InstanceName>\|Total Latency | Cassandra | False |
| ClientRequest<InstanceName>\|Unavailables | Cassandra | False |
| CommitLog\|Pending Tasks | Cassandra | False |
| CommitLog\|Total Commit Log Size | Cassandra | False |
| Compaction\|Bytes Compacted | Cassandra | False |
| Compaction\|Completed Tasks | Cassandra | False |
| Compaction\|Pending Tasks | Cassandra | False |

## Table 4-45. Cassandra Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Compaction\|Total Compactions Completed | Cassandra | False |
| Connected Native Clients | Cassandra | False |
| HeapMemoryUsage\|committed | Cassandra | False |
| HeapMemoryUsage\|init | Cassandra | False |
| HeapMemoryUsage\|max | Cassandra | False |
| HeapMemoryUsage\|used | Cassandra | False |
| NonHeapMemoryUsage\|committed | Cassandra | False |
| NonHeapMemoryUsage\|init | Cassandra | False |
| NonHeapMemoryUsage\|max | Cassandra | False |
| NonHeapMemoryUsage\|used | Cassandra | False |
| ObjectPendingFinalizationCount | Cassandra | False |
| Storage\|Exceptions Count | Cassandra | False |
| Storage\|Load Count | Cassandra | False |
| Table<InstanceName>\|Coordinator Read Latency | Cassandra | False |
| Table<InstanceName>\|Live Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Total Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Write Latency | Cassandra | False |
| Table<InstanceName>\|Write Latency | Cassandra | False |
| ThreadPools<InstanceName>\|Active Tasks | Cassandra | False |
| ThreadPools<InstanceName>\| Currently Blocked Tasks | Cassandra | False |
| ThreadPools<InstanceName>\| Pending Tasks | Cassandra | False |
| Application Availability | Cassandra | False |

## HyperV Metrics

Metrics are collected for the HyperV application service.

### Table 4-46. HyperV Metrics

| Metric Name | Category | KPI |
|---|---|---|
| VM:Hyper-V Virtual Machine Health Summary\|Health Critical | HyperV | False |
| VM<instanceName>\|Physical Memory | HyperV | False |
| VM<instanceName>Hv VP 0\|Total Run Time | HyperV | False |
| VM<instanceName>\|Bytes Received | HyperV | False |
| VM<instanceName>\|Bytes Sent | HyperV | False |
| VM<instanceName>\|Error Count | HyperV | False |
| VM<instanceName>\|Latency | HyperV | False |
| VM<instanceName>\|Queue Length | HyperV | False |
| VM<instanceName>\|Throughput | HyperV | False |
| CPU<instanceName>\|Idle Time | HyperV | True |
| CPU<instanceName>\|Processor Time | HyperV | True |
| CPU<instanceName>\|User Time | HyperV | True |
| Disk<instanceName>\|Avg Disk Queue Length | HyperV | False |
| Disk<instanceName>\|Idle Time | HyperV | False |
| Disk<instanceName>\|Read Time | HyperV | True |
| Disk<instanceName>\|Write Time | HyperV | True |
| Process<instanceName>\|Private Bytes | HyperV | False |
| Process<instanceName>\|Processor Time | HyperV | False |
| Process<instanceName>\|Thread Count | HyperV | False |
| Process<instanceName>\|User Time | HyperV | False |
| System\|Processes | HyperV | False |
| System\|Processor Queue Length | HyperV | False |
| System\|System UpTime | HyperV | False |
| Memory\|Available Bytes | HyperV | False |

### Table 4-46. HyperV Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| Memory\|Cache Bytes | HyperV | False |
| Memory\|Cache Faults | HyperV | False |
| Memory\|Pages | HyperV | False |
| Network<instanceName>\|Packets Outbound Error | HyperV | False |
| Network<instanceName>\|Packets Received Error | HyperV | False |
| Application Availability | HyperV | False |

## MongoDB Metrics

Metrics are collected for the MongoDB application service.

### Table 4-47. MongoDB Metrics

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Active Reads | MongoDB | True |
| UTILIZATION\|Active Writes | MongoDB | True |
| UTILIZATION\|Connections Available | MongoDB | False |
| UTILIZATION\|Connections Total Created | MongoDB | False |
| UTILIZATION\|Current Connections | MongoDB | True |
| UTILIZATION\|Cursor Timed Out | MongoDB | True |
| UTILIZATION\|Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Document Inserted | MongoDB | False |
| UTILIZATION\|Document Deleted | MongoDB | False |
| UTILIZATION\|Flushes Per Sec | MongoDB | False |
| UTILIZATION\|Inserts Per Sec | MongoDB | False |
| UTILIZATION\|Net Input Bytes | MongoDB | False |
| UTILIZATION\|Open Connections | MongoDB | True |
| UTILIZATION\|Page Faults Per Second | MongoDB | False |
| UTILIZATION\|Net Output Bytes | MongoDB | False |
| UTILIZATION\|Queries Per Sec | MongoDB | False |
| UTILIZATION\|Queued Reads | MongoDB | True |

### Table 4-47. MongoDB Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Queued Writes | MongoDB | True |
| UTILIZATION\|Total Available | MongoDB | False |
| UTILIZATION\|Total Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Total Passes Per Sec | MongoDB | False |
| UTILIZATION\|Total Refreshing | MongoDB | False |
| UTILIZATION\|Updates Per Sec | MongoDB | False |
| UTILIZATION\|Volume Size MB | MongoDB | False |
| Application Availability | MongoDB | False |
| UTILIZATION\|Collection Stats | MongoDB DataBases | False |
| UTILIZATION\|Data Index Stats | MongoDB DataBases | True |
| UTILIZATION\|Data Indexes | MongoDB DataBases | False |
| UTILIZATION\|Data Size Stats | MongoDB DataBases | True |
| UTILIZATION\|Average Object Size stats | MongoDB DataBases | False |
| UTILIZATION\|Num Extents Stats | MongoDB DataBases | False |

## Riak KV Metrics

Metrics are collected for the Riak KV application service.

### Table 4-48. Riak KV Metrics

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|CPU Average | Riak KV | False |
| UTILIZATION\|Memory Processes | Riak KV | False |
| UTILIZATION\|Memory Total | Riak KV | False |
| UTILIZATION\|Node GETs | Riak KV | True |
| UTILIZATION\|Node GETs Total | Riak KV | False |
| UTILIZATION\|Node PUTs | Riak KV | True |
| UTILIZATION\|Node PUTs Total | Riak KV | False |
| UTILIZATION\|PBC Active | Riak KV | True |
| UTILIZATION\|PBC Connects | Riak KV | True |
| UTILIZATION\|Read Repairs | Riak KV | True |

### Table 4-48. Riak KV Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|vNODE Index Reads | Riak KV | True |
| UTILIZATION\|vNODE Index Writes | Riak KV | True |
| Application Availability | Riak KV | False |

## Network Time Protocol Metrics

Metrics are collected for the Network Time Protocol application service.

### Table 4-49. Network Time Protocol Metrics

| Metric Name | Category | KPI |
|---|---|---|
| NTPD \|delay | Network Time Protocol | True |
| NTPD \| jitter | Network Time Protocol | True |
| NTPD \| offset | Network Time Protocol | True |
| NTPD \| poll | Network Time Protocol | False |
| NTPD \| reach | Network Time Protocol | True |
| NTPD \| when | Network Time Protocol | False |
| Application Availability | Network Time Protocol | False |

## WebSphere Metrics

Metrics are collected for the WebSphere application service.

### Table 4-50. WebSphere Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Thread Pool\|Active Count\| Current | Thread Pool | False |
| Thread Pool\|Active Count\| High | Thread Pool | False |
| Thread Pool\|Active Count\| Low | Thread Pool | False |
| Thread Pool\|Active Count\| Lower | Thread Pool | False |
| Thread Pool\|Active Count\| Upper | Thread Pool | False |
| JDBC\|Close Count | JDBC | False |
| JDBC\|Create Count | JDBC | False |
| JDBC\|JDBC Pool Size\| Average | JDBC | False |

## Table 4-50. WebSphere Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| JDBC\|JDBC Pool Size\| Current | JDBC | False |
| JDBC\|JDBC Pool Size\| Lower | JDBC | False |
| JDBC\|JDBC Pool Size\| Upper | JDBC | False |
| Garbage Collection<InstanceName>\| Total Collection Count | WebSphere | False |
| Garbage Collection<InstanceName>\| Total Collection Time | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Number of Object Pending Finalization Count | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Committed Memory | WebSphere | False |

### Table 4-50. WebSphere Metrics (continued)

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Pool<InstanceName>\|Peak Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Peak Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Peak Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Used Memory | WebSphere | False |
| Process Cpu Load | WebSphere | False |
| System Cpu Load | WebSphere | False |
| System Load Average | WebSphere | False |
| Application Availability | WebSphere | False |

## Java Application Metrics

Metrics are collected for the Java application service.

### Table 4-51. Java Application Metrics

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Java Application | False |
| Buffer Pool<InstanceName>\|Memory Used | Java Application | False |
| Buffer Pool<InstanceName>\|Total Capacity | Java Application | False |
| Class Loading\|Loaded Class Count | Java Application | True |
| Class Loading\|Total Loaded Class Count | Java Application | False |
| Class Loading\|Unloaded Class Count | Java Application | False |

## Table 4-51. Java Application Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| Garbage Collection<InstanceName>\|Total Collection Count | Java Application | False |
| Garbage Collection<InstanceName>\|Total Collection Time | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Initial Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\|Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\|Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\|Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\|Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Used Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Java Application | False |

### Table 4-51. Java Application Metrics (continued)

| Metric Name | Category | KPI |
| --- | --- | --- |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Java Application | False |
| JVM Memory\|Object Pending Finalization Count | Java Application | False |
| Uptime | Java Application | True |
| Threading\|Thread Count | Java Application | True |
| Process CPU Usage % | Java Application | False |
| System CPU Usage % | Java Application | False |
| System Load Average % | Java Application | False |

## Remote Check Metrics

Metrics are collected for object types such as HTTP, ICMP, TCP, and UDP.

### HTTP Metrics

VMware Aria Operations discovers metrics for HTTP remote checks.

HTTP Metrics

### Table 4-52. HTTP Metrics

| Metric Name | KPI |
| --- | --- |
| Availability | False |
| Content Length | False |
| Response Code | False |
| Response Time | True |
| Result Code | False |

### ICMP Metrics

VMware Aria Operations discovers metrics for the ICMP object type.

### Table 4-53. ICMP Metrics

| Metric Name | KPI |
| --- | --- |
| Availability | False |
| Average Response Time | True |
| Packet Loss (%) | False |
| Packets Received | False |

### Table 4-53. ICMP Metrics (continued)

| Metric Name | KPI |
| --- | --- |
| Packets Transmitted | False |
| Result Code | False |

## TCP Metrics

VMware Aria Operations discovers metrics for the TCP object type.

### Table 4-54. TCP Metrics

| Metric Name | KPI |
| --- | --- |
| Availability | False |
| Response Time | True |
| Result Code | False |

## UDP Metrics

VMware Aria Operations discovers metrics for the UDP object type.

### Table 4-55. UDP Metrics

| Metric Name | KPI |
| --- | --- |
| Availability | False |
| Response Time | True |
| Result Code | False |

## Linux Process Metrics

Metrics are collected for Linux services.

### Table 4-56. Linux Process Metrics

| Metric Name | Category | KPI |
| --- | --- | --- |
| AVAILABILITY\|Resource Availability | Processes | False |
| UTILIZATION\|Memory Usage (%) | Processes | False |
| UTILIZATION\|CPU Usage (%) | Processes | False |
| UTILIZATION\|Number of Processes | Processes | False |

## Windows Service Metrics

Metrics are collected for Windows services.

Table 4-57. Windows Service Metrics

| Metric Name | Category | KPI |
| --- | --- | --- |
| AVAILABILITY\|Resource Availability | Services | False |
| UTILIZATION\|Memory Usage(%) | Services | False |
| UTILIZATION\|CPU Usage(%) | Services | False |

## OS and Application Monitoring Properties

Properties are collected for operating systems, application services, remote checks, Linux processes, and Windows services which can be used to create reports, views, and dashboards.

### Guest Information Properties

VMware Aria Operations displays the following guest information properties for all objects created by the OS and Application Monitoring management pack.

- Guest Info
  - Hostname
  - IP
  - OS Name
  - OS Version
  - Telegraf Version

Other properties of operating systems and application services are available under **Properties > Tags**.

# Troubleshooting

## Troubleshooting Agent Installation

### Agent Install Failure Because of the vCenter Server User Permissions

Guest operation privileges are required to install agents on end point virtual machines.

#### Problem

Agent installation fails with the following error message if there are no guest operation privileges:

```
vCenter adapter user is missing either of the following guest operations privileges –
execute, modify, query
```

#### Solution

1 Verify that you have configured a vCenter adapter.

2 The vCenter Server user account with which the vCenter adapter is configured in VMware Aria Operations, should have the following permissions: **Guest operation modifications**, **Guest operation program execution**, and **Guest operation queries**.

### Agent Install Failure Because NTP is Not in Sync

After you install or upgrade to the latest version of cloud proxy, you must set up accurate timekeeping as part of the deployment. If the time settings between cloud proxy and VMware Aria Operations are not synchronized, you face agent installation and metric collection issues. Ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts, cloud proxy and VMware Aria Operations using the Network Time Protocol (NTP).

#### Problem

- Agent installation fails

#### Solution

- Log in to cloud proxy and run the following command to stop the NTP daemon:

```
systemctl stop ntpd
```

- Run the following command to update the time immediately from an NTP server:

```
ntpdate time.vmware.com
```

**Note** Replace `time.vmware.com` with a suitable time server setting. You can use the FQDN or IP of the time server.

- Enter the following command to start the NTP daemon:

```
systemctl start ntpd
```

**Note** The system time takes about five minutes to sync with the NTP server time.

### Agent Install Fails on a Linux End Point

Install of an agent on a Linux end point fails for a non-root user with a specific set of privileges.

#### Problem

Agent installation fails with the following error if the `tty` command is not added:

```
Bootstrap Failed for VM <VM ID> with error message:{ "status":"FAILED", "data":
[ { "status":"FAILED", "message":"Failed - install - passwordless sudo access is required for
the user <Install Username> on the command mkdir. [sudo: sorry, you must have a tty to run
sudo]", "stage":"0" } ], "currentstage":"0", "totalstages":"0" }
```

The following are minimal necessary permissions of the user to install agents and should be mentioned in the `sudoers` file:

#### Solution

- The following are minimal necessary permissions of the user to install agents and should be mentioned in the `sudoers` file. For example, for a user called **telegrafinstall**, you can find the sudoers file in the **/etc/sudoers** file or in the folder **/etc/sudoers.d/:**. Add these lines to `/etc/sudoers`, if you have not added them.

```
Defaults:telegrafinstall !requiretty

Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/
cp*,/usr/bin/mkdir*,/bin/mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-
bootstrap.sh,/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh

telegrafinstall ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS
```

### Agent Install Fails on Windows When UAC is Deactivated

#### Problem

Install of the agent fails even when UAC is deactivated.

#### Solution

- Make sure UAC (previously known as LUA) was deactivated on Windows in the following way:
  a In the registry path `HKLM:\Software\Microsoft\Windows\CurrentVersion\Policies\System`, set the value for the key `EnableLUA` to `0`.
  b You must reboot the machine for the changes to take effect.

### Agent Install Fails on Windows with a Permission Denied Error

In Windows, during bootstrap, when the Telegraf folder is renamed to `ucp-telegraf`, it can result in a failure because of a permission error.

#### Problem

Sometimes, there are certain antiviruses running, which prevent the application from renaming or modifying the directory or files. In such a situation, the following error message is displayed:

```
Install telegraf [unable to install telegraf due to system error : [WinError
5] Access is denied: 'C:\\VMware\\UCP\\ucp-telegraf']"}].
```

#### Solution

- Deactivate the antivirus and then proceed with bootstrapping.

### Agent Install Does Not Progress

#### Problem

During agent install, new tasks do not progress beyond the Starting phase, from recent tasks. Adapter logs are not written.

#### Solution

Verify that the adapter instance in cloud proxy is in a Data Collecting state. If not, restart the adapter instance from the VMware Aria Operations user interface. From the left menu, navigate to **Environment > Inventory > Adapter Instances > VMware Aria Operations Application Management Adapter Instance**. From the **Objects** tab in the right pane, select the adapter instance and click **Stop Collecting** and then **Start Collecting**.

### Agent Install Fails Without an Error Message

#### Problem

Agent install fails without any error message, in the user interface

The `uaf_bootstrap.log` at the endpoint VM displays the following log message: `'findstr' is not recognized as an internal or external command`.

#### Cause

This happens because `C:\Windows\System32` is not available in the environment variable *PATH*.

#### Solution

Add `C:\Windows\System32` to the environment variable *PATH*.

### Script Download Fails on a Windows Platform

When a file is downloaded from cloud proxy to a Windows end point, it could fail due to security protocols.

#### Problem

Script download fails on a Windows platform with the following message:

```
The request was aborted: Could not create SSL/TLS secure channel.
```

There are three kinds of PowerShell scripts hosted in cloud proxy that can be downloaded and executed at the Windows end point VMs for different purposes:

- To install custom Telegraf using a script (download.ps1).
- To install custom Telegraf on a physical server (unmanagedagent_setup_sample.ps1).
- To configure open source Telegraf on managed or unmanaged VMs (open_source_telegraf_monitor.ps1).

### Solution

Ignore the `ServerCertificateValidationCallback` using the following command.

```
if (-not
([System.Management.Automation.PSTypeName]'ServerCertificateValidationCallback').Type)
{
$certCallback = @"
    using System;
    using System.Net;
    using System.Net.Security;
    using System.Security.Cryptography.X509Certificates;
    public class ServerCertificateValidationCallback
    {
        public static void Ignore()
        {
            if(ServicePointManager.ServerCertificateValidationCallback ==null)
            {
                ServicePointManager.ServerCertificateValidationCallback +=
                    delegate
                    (
                        Object obj,
                        X509Certificate certificate,
                        X509Chain chain,
                        SslPolicyErrors errors
                    )
                    {
                        return true;
                    };
            }
        }
    }
"@
    Add-Type $certCallback
 }
[ServerCertificateValidationCallback]::Ignore()
[Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12
```

After downloading and executing the required script, `ServerCertificateValidationCallback` can be activated.

## Troubleshooting Plugin Related Failures

### Unable to Activate a Plugin

Unable to activate a plugin with the same fields until the plugin configuration is deleted.

#### Problem

An error message is displayed in the user interface of VMware Aria Operations that states the following:

```
Failed to update resource: Resource with same key already exists
```

#### Solution

- Manually delete the existing plugin configuration and then continue with the activation of the plugin. If the problem persists, delete the corresponding resource from the inventory.

### RabbitMQ Plugin Error

You might receive an error while monitoring the RabbitMQ plugin.

#### Problem

While monitoring the RabbitMQ plugin, an error might occur if you upgrade to VMware Aria Operations, which has a Telegraf version of 1.19. The following message is displayed:

```
getting "/api/federation-links" failed: 404 Not Found
```

#### Cause

The Federation plugins in the RabbitMQ setup are not activated.

#### Solution

- Activate the Federation plugins in the RabbitMQ setup by running the following commands in the RabbitMQ VM:

  ```
  rabbitmq-plugins enable rabbitmq_federation

  rabbitmq-plugins enable rabbitmq_federation_management
  ```

## Troubleshooting Metric Collection

### Troubleshoot Agent Installation and Metric Collection Issues

If the time settings between cloud proxy and VMware Aria Operations are not synchronized, you might face agent installation and metric collection issues. Eventually, you might not see any metrics in the VMware Aria Operations dashboards.

#### Problem

You might notice the following issues in VMware Aria Operations:

- You cannot install an agent in the Windows and Linux target VMs.

#### Cause

Time synchronization is a prerequisite of the TLS/SSO communication between client and server.

If the VMware Aria Operations and cloud proxy are not time synchronized, the test connection fails while configuring cloud proxy in VMware Aria Operations.

If the Windows and Linux target VMs are not time synchronized with VMware Aria Operations, communication between cloud proxy and agents will break after installing the agents. Hence monitored metrics are not sent to VMware Aria Operations. Alternatively, stop and restart the agent to resolve this issue.

#### Solution

1 Check the VMware Aria Operations support bundle in the following path: `COLLECTOR/adapters/APPOSUCPAdapter/` for errors.

2 Check the cloud proxy support bundle, *ucpapi.log*, for errors.

3 Ensure time synchronization between cloud proxy, VMware Aria Operations and the Windows and Linux target VMs.

4 To start and restart the agent, see Additional Operations from the Manage Telegraf Agents Page .

### Troubleshooting Content Upgrade for an End Point

#### Problem

Content upgrade for an end point fails with the following error:

`Timeout Error. Please retry the action after some time.`

#### Cause

Sometimes content upgrade for an end point fails because of a timeout in the cloud proxy.

#### Solution

- Retrigger content upgrade for the end point to resolve the issue.

### Telegraf Agent Related Actions Fail after Cloud Proxy is Restarted or Upgraded

#### Problem

After an upgrade of vRealize Operations/vRealize Operations Cloud and cloud proxy from 8.4 to a later releases or if you restart cloud proxy, Telegraf agent management related actions could fail with the error message `Connect to Salt Master`.

#### Cause

The saltmaster docker container inside cloud proxy does not run properly causing all Telegraf related actions to fail.

#### Solution

1 SSH to the cloud proxy VM.

2 Verify if the salt process count by command is: `ps -ef | grep salt`.

3 If the salt process count is less than 15, run the following command: `/rpm-content/ucp/subsequentboot.sh`.

   You can view the log from the following location: `/opt/vmware/var/log/ucp-subsequentboot`.

## Troubleshooting Using Support Bundles

Support bundles are required to troubleshoot problems related to application monitoring. For Linux and Windows end point VMs, run the specified command and access the support bundle.

### For End Point VMs

1 Log in to the end point.

2 Run the following commands based on the end point VM's operating system type:

**For Linux End Point VMs**

```
/opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh --config /opt/vmware/ucp/salt-minion/etc/salt/
grains --action gen_support_bundle --log_level INFO
```

The support bundle is generated and placed as a ZIP file in the `/opt/vmware/ucp/support-bundle-endpoints/` directory.

**For Windows End Point VMs**

```
C:\VMware\UCP\ucp-minion\bin\ucp-minion.bat --config C:\VMware\UCP\salt\conf\grains --
action gen_support_bundle --log_level INFO
```

The support bundle is generated and placed as a ZIP file in the `%SystemDrive%\VMware\UCP\support-bundle-endpoints\` directory.

## Objects Do Not Receive Data from the Adapter Instance

### Problem

In case of a HA failover, you see the following alert: `Objects are not receiving data from the adapter instance`. This is an adapter-specific alert that does not determine whether the object is receiving data from any of the adapters. It is recommended that you deactivate this alert.

### Solution

1 Create a new policy that inherits the AppOS adapter instance's pre-applied policy. From the left menu, click **Configure > Policies**, and then in the left pane, click **Policy Definition**.

2 Click **Add** to add a policy, or select the policy and click **Edit Policy** to edit an existing policy.

You can add and edit policies and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.

3 In the **Create Policies** workspace, assign a name to the policy and enter the description.

a Give the policy a meaningful name and description so that all users know the purpose of the policy.

b From the **Inherit From** drop-down option, select the policy to use as a baseline to define the settings for your new local policy.

4 Click **Create Policy**. The **Create Policies** workspace provides options to customise your policy.

5 Click the **Alerts and Symptoms** tile and search for the alert called `Objects are not receiving data from adapter instances` from the policy created in step 1.

6 Click on the alert and deactivate it.

7 Apply the policy created in step 1 to the AppOS adapter instances.

# Monitoring Application Services and Operating Systems using Open Source Telegraf

You can use open source Telegraf that runs on a vCenter Server VM, an EC2 AWS instance, or an Azure VM to send metrics to VMware Aria Operations using a helper script or by providing specific configurations required on the end point to post metrics to cloud proxy. The helper script adds necessary configurations in the Telegraf configuration that is directly identified by VMware Aria Operations.

You can monitor application services and operating systems that are supported by VMware Aria Operations and you can also monitor unsupported application services. For a list of supported application services, see Supported Application Services .

Open source Telegraf is supported on Linux and Windows platforms.

**Notice:** VMware supports only the installation of open source Telegraf because there is a configuration required to point it to the VMware Aria Operations cloud proxy. After the initial configuration is complete, and we have made sure that the metrics are received consistently for a couple of end points, and there are no common patterns of failures observed on VMware Aria Operations, no further support is offered from VMware. Any other queries on the agent has to go through the Telegraf open source community. Issues related to Telegraf can be raised through: https://github.com/influxdata/telegraf/issues.

## Points to Note

As you can monitor both supported and unsupported application services using open source Telegraf, keep in mind the following points.

- There are no alerts generated for non-supported application services.
- You cannot add metrics for supported application services.
- Application service objects that are not supported are named as follows: *<application service name_Generic>*
- Metrics for unsupported application services are displayed without proper categorization and appears as is from Telegraf.
- Metrics for unsupported application services are not localized and are provided only in English.

- If there are multiple instances of unsupported application services, to distinguish between them in the user interface, add the tag 'identifier' and provide a unique tag value in the `telegraf.conf` file. Example, [inputs.mongodb.tags] identifier="1".
- If multiple instances of the same supported application service is monitored by open source Telegraf, in the user interface, the display name is the same for all instances of the application service. To differentiate between the instances of the same application service, edit the display name from the user interface and provide a specific name.
- Inactivated application services with a managed agent are displayed as activated in the **Manage Telegraf Agents** page. Uninstall managed agents before you use open source Telegraf.
- Telegraf version 1.19.0 is not supported when you use open source Telegraf.

## Monitoring Application Using Open Source Telegraf

You can use open source Telegraf to monitor application services and operating systems on a or Linux platform or on a Windows platform.

### Install and Configure Open Source Telegraf

Install and configure open source Telegraf to monitor your applications.

Telegraf uses input plugins (where the metrics come from) and output plugins (where the metrics go) in the configuration files. You can see all the supported plugins at Plugin Directory. Input and output plugins should be written in the Telegraf configuration file and configuration directory. The Telegraf configuration file (telegraf.conf) lists all available Telegraf plugins.

Telegraf uses the `--config` flag to specify the configuration file location and `--config-directory` flag to include files ending with `.conf` in the specified directory. On most Linux systems, the default locations are `/etc/telegraf/telegraf.conf` for the main configuration file and `/etc/telegraf/telegraf.d` for the directory of configuration files. For Windows platforms, it will be in the location where the `telegraf zip` is extracted.

For more information, see the Telegraf documentation, Get Started and Configuration Options. After you have downloaded and installed Telegraf, you are ready to begin collecting and sending data. To collect and send data Telegraf should be configured. Follow the steps below:

#### Procedure

1. Install open source Telegraf on the end point. If you have an instance installed, you can skip this step. To download and install a new instance of Telegraf, see the official documentation and search for the corresponding OS version from https://www.influxdata.com/time-series-platform/telegraf/ and https://portal.influxdata.com/downloads/.

   On a Windows platform only, after downloading and extracting Telegraf files, besides `telegraf.exe` and `telegraf.conf` files, create a folder with the name `telegraf.d` that you use to run the helper script and monitor applications.

2 Run the helper script to configure Telegraf to send data to cloud proxy.

   For more information, see Monitoring Applications Using Open Source Telegraf on a Linux Platform and Monitoring Applications Using Open Source Telegraf on a Windows Platform.

   After you run the helper script, a `cloudproxy-http.conf` file is created and the output plugin is added to the file with the properties necessary for sending data and the input plugins are updated for OS metrics. For OS metrics configuration template, see Telegraf Configuration Details for Operating Systems. For information about the `cloudproxy-http.conf` configuration template, see Sample Configurations.

3 Using open source Telegraf, you can collect metrics from different types of application services.

   - If an application service (curated plugin) that is supported by Application Monitoring in VMware Aria Operations is running on the end point and you want to monitor it, update the Telegraf configuration file or directory with necessary inputs for Telegraf.

     For a list of supported application services (curated plugins), see Supported Application Services . For the list of configurations, see Telegraf Configuration Details for Supported Application Services.

   - For unsupported (non-curated plugins) application services, update the Telegraf configuration file or directory with the necessary inputs for Telegraf. For more information, see the Telegraf documentation Plugin Directory. For example, if you want to collect data for Ethernet device stats (device name is eth0) you can create the `ethools.conf` file in the Telegraf configuration directory and add the following content (for more information, see the documentation about Ethtool Input Plugin):

     ```
     [[inputs.ethtool]]
       # List of interfaces to pull metrics for
        interface_include = ["eth0"]
     ```

   - Whenever the `telegraf-utils` script is executed with Open Source Telegraf, the `cloudproxy-http.conf` file will be overwritten in the `telegraf.d` folder. You must maintain a different config file (for example, `postgres.conf`) for other custom input plugins to retain the configuration.

4 Start or restart Telegraf.

## Monitoring Applications using Open Source Telegraf on a Linux Platform

Use the helper script to monitor applications and operating systems on a Linux platform using open source Telegraf.

The helper script is tested only on the following operating systems:

- CentOS 7.x and CentOS 8.x
- RHEL 7.x and RHEL 8.x
- SUSE 12.x and SUSE 15.x

- OEL7.x and OEL 8.x
- Ubuntu 16.x, Ubuntu 18.x, Ubuntu 20.x, and Ubuntu 22.x
- VMware Photon Linux

### Prerequisites

- Install the jq package. For more information, see the official documentation for jq from https://stedolan.github.io/jq/download/.
- Ensure that VMTools version >=10.2 if a vCenter Server VM is monitored by VMware Aria Operations. `vmware-toolbox-cmd` is necessary only for vCenter Server VMs.
- Ensure that you have permissions to run the following commands and use the following packages:
  - Commands: /bin/bash, awk, sed, vmware-toolbox-cmd, jq, curl
  - Packages: coreutils (chmod, chown, cat), net-tools (ip, /dev/tcp, curl, wget)
- Only IPv4 is supported at present for cloud proxy.
- Ensure that Internet is activated.
- Ensure that cloud proxy is up and online in VMware Aria Operations. In the case of collector groups, ensure that at least one cloud proxy in the application monitoring high availability activated collector group is up and online in VMware Aria Operations.
- Verify that unzip is at 6.0-20.el7 or above.
- The uuidgen package must exist on the vCenter Server VMs, AWS EC2 instances, Azure VMs, or the physical server.
- vCenter Server VMs, AWS EC2 instances, Azure VMs, or the physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.

### Procedure

1. Download the helper script from cloud proxy located at `https://<CloudProxy-IP>/downloads/salt/telegraf-utils.sh`.

   **Note** Use the relevant cloud proxy **IP address** for <CloudProxy-IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. Navigate to the directory where the script is downloaded.
3. Activate execution permission of the script for Linux VM.

```
chmod +x telegraf-utils.sh
```

4 Run the helper script to update Telegraf configurations.

```
telegraf-utils.sh  opensource -c cloud_proxy_ip_or_collector_group_name -t token -d
telegraf_conf_dir -e telegraf_bin_path [-g gateway_url -a csp_auth_url]

Description of arguments:
      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector
Group Name
       -t : [REQUIRED] token - CSP Refresh Token of the user/account. For getting a new
token,
          follow - User/Organization Settings >> My Account >> API Tokens >> Generate a
New API Token.
          Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq

      -d : [REQUIRED] telegraf_conf_dir - Telegraf configuration directory and it is
required argument. ex: /etc/telegraf/telegraf.d
      -e : [REQUIRED] telegraf_bin_path - Path of telegraf executable ex: /usr/bin/telegraf

      -g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware
Aria Operations SaaS gateway URL
      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL
```

```
Example: /bin/bash telegraf-utils.sh opensource -t
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq -c 10.10.10.101 -d /etc/
telegraf/telegraf.d -e /usr/bin/telegraf
```

**Note** After you run the helper script, ensure that the respective configurations are set correctly in the given config directory (`-d` option) path with the name `cloudproxy-http.conf`. See Sample Configurations for more details. For managed vCenter Server VMs, AWS EC2 instances, or Azure VMs, you might see unmanaged configurations, because of one of the following reasons:

- vCenter Server VMs, AWS EC2 instances, or Azure VMs details are not available in VMware Aria Operations by the vCenter Server, AWS, or Azure adapters correspondingly. Wait for a minimum of one to two collection cycles after configuring the VMware Aria Operations vCenter Server cloud accounts.
- An incorrect `SAAS_REFRESH_TOKEN`.

**Note** By default, the InfluxDB output plugin is active in the `telegraf.conf` file and data is sent to the influxdb server so that you do not get multiple warning messages in the logs about the lack of configured influxdb server comment, the "[[outputs.influxdb]]" line should be commented. The following warning message is displayed:`W! [outputs.influxdb] When writing to [http://localhost:8086]: database "telegraf" creation failed: Post "http://localhost:8086/query": dial tcp [::1]:8086: connect: connection refused`

Example: `#[[outputs.influxdb]]`

**Note** Ensure that the input plugins in the `telegraf.conf` file are related to the corresponding operating system. See Telegraf Configuration Details for Operating Systems.

**Note** The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

5 Restart the Telegraf service.

```
systemctl restart telegraf
or
/usr/bin/telegraf -config /etc/telegraf/telegraf.conf -config-directory /etc/telegraf/
telegraf.d
```

### What to do next

**Managed VM object hierarchy**:

- If a vCenter Server VM is monitored by VMware Aria Operations, then the operating system and application objects fall under the respective **VM > OS object > 'application service' instance**.
- If an Azure VM is monitored by VMware Aria Operations, then the operating system and application objects fall under the respective **Azure VM > OS object > 'application service' instance**.
- If an AWS EC2 instance is monitored by VMware Aria Operations, then the operating system and application objects fall under the respective **AWS EC2 instance > OS object > 'application service' instance**.

**Machines not monitored by VMware Aria Operations**: If a vCenter Server VM, an AWS EC2 instance, an Azure VM, or a physical server is not monitored by VMware Aria Operations, then the operating system and application objects fall under **Environment > Endpoint > OS Object > 'application service' instance**.

**View Data Collected**: Data is collected and appears in the **Manage Telegraf Agents** page. To view the details, you can filter by **Agent Type > Open source agent** from the **Manage Telegraf Agents** page.

## Monitoring Applications using Open Source Telegraf on a Windows Platform

Use the helper script to monitor applications and operating systems on a Windows platform using open source Telegraf.

The helper script is tested only on Windows Server 2012, 2012 R2, 2016, 2019, and 2022.

### Prerequisites

- Verify that Windows PowerShell is at 4.0 or above.
- Ensure that VMTools version >=10.2 if a vCenter Server VM is monitored by VMware Aria Operations. `vmware-toolbox-cmd` is necessary only for vCenter Server VMs.
- Only IPv4 is supported at present for cloud proxy.
- Ensure that Internet is enabled.
- Ensure that cloud proxy is up and online in VMware Aria Operations. In the case of collector groups, ensure that at least one cloud proxy in the application monitoring high availability activated collector group is up and online in VMware Aria Operations.
- vCenter Server VMs, AWS EC2 instances, Azure VMs, or the physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.

### Procedure

1 Download the helper script from cloud proxy located at `https://<CloudProxy-IP>/downloads/salt/telegraf-utils.ps1`.

**Note** Use the relevant cloud proxy **IP address** for <CloudProxy-IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

If the script download fails with the following message: `The request was aborted: Could not create SSL/TLS secure channel`, follow the steps mentioned in Script Download Fails on a Windows Platform.

2 Navigate to the directory where the script is downloaded.

3 Run the helper script to update Telegraf configurations.

```
telegraf-utils.ps1  opensource -c cloud_proxy_ip_or_collector_group_name -t token -d
telegraf_conf_dir -e telegraf_bin_path [-g gateway_url -a csp_auth_url]

Description of arguments:
      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector
Group Name
      -t : [REQUIRED] token - [REQUIRED] - CSP Refresh Token of the user/account. For
getting a new token, follow -
           User/Organization Settings >> My Account >> API Tokens >> Generate a New API
Token.
           Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq

      -d : [REQUIRED] telegraf_conf_dir - Telegraf configuration directory and it is
required argument. ex: C:\Telegraf\telegraf-1.20.4\telegraf.d
      -e : [REQUIRED] telegraf_bin_path - Path of telegraf executable ex:
C:\Telegraf\telegraf-1.20.4\telegraf.exe

      -g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware
Aria Operations SaaS gateway URL
      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL
```

```
Example: .\telegraf-utils.ps1 opensource -c 10.10.10.101
-t gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq -d
C:\Telegraf\telegraf-1.20.4\telegraf.d -e C:\Telegraf\telegraf-1.20.4\telegraf.exe
```

**Note** After you run the helper script, ensure that the respective configurations are set correctly in the given config directory (`-d` option) path with the name `cloudproxy-http.conf`. See Sample Configurations for more details. For managed vCenter Server VMs, AWS EC2 instances, or Azure VMs, you might see unmanaged configurations, because of one of the following reasons:

- vCenter Server VMs, AWS EC2 instances, or Azure VMs details are not available in VMware Aria Operations by the vCenter Server, AWS, or Azure adapters correspondingly. Wait for a minimum of one to two collection cycles after configuring the VMware Aria Operations vCenter Server cloud accounts.
- An incorrect `SAAS_REFRESH_TOKEN`.

**Note**

- Do not use a space in the configuration path. Paths with spaces can be passed as a short name notation, such as `c:\PROGRA~1` for `c:\Program Files`.

**Note** Ensure that the input plugins in the `telegraf.conf` file are related to the corresponding operating system. See Telegraf Configuration Details for Operating Systems.

**Note** The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

4 Restart the Telegraf service.

```
telegraf.exe --config telegraf.conf --config-directory telegraf.d
```

Or you can make Telegraf a Windows service.

```
<Telegraf_executable_path> --config <Telegraf_config_file_path> --config-directory
<Telegraf_config_directory_path> --service install net start telegraf
```

For example:

```
& 'C:\Telegraf\telegraf-1.20.4\telegraf.exe'  --config
'C:\Telegraf\telegraf-1.20.4\telegraf.conf' --config-directory
'C:\Telegraf\telegraf-1.20.4\telegraf.d'  --service install
net start telegraf
```

### What to do next

**Managed VM object hierarchy**:

- If a vCenter Server VM is monitored by VMware Aria Operations, then the operating system and application objects will fall under the respective **VM > OS object > 'application service' instance**
- If an Azure VM is monitored by VMware Aria Operations, then the operating system and application objects will fall under the respective **Azure VM > OS object > 'application service' instance**.
- If an AWS EC2 instance of the VM is monitored by VMware Aria Operations, then the operating system and application objects will fall under the respective **AWS EC2 instance > OS object > 'application service' instance**

**Machines not monitored by VMware Aria Operations**: If a vCenter Server VM, an AWS EC2 instance, an Azure VM, or a physical server is not monitored by VMware Aria Operations, then the operating system and application objects will fall under **Environment > Endpoint > OS Object > 'application service' instance**.

**View Data Collected**: Data is collected and appears in the **Manage Telegraf Agents** page. To view the details, you can filter by **Agent Type > Open source agent** from the **Manage Telegraf Agents** page.

## Telegraf Configuration Details for Operating Systems

### Linux Operating Systems

To collect Linux OS related metrics and properties as it is in the managed Telegraf agent installation flow and to have localized object types, names, alerts, symptoms, metrics and properties, and so on, the `telegraf.conf` file (in case of installed telegraf it is located in `/etc/telegraf/telegraf.conf`) should be customized and have the following content:

```
# Read metrics about cpu usage
[[inputs.cpu]]
  ## Whether to report per-cpu stats or not
  percpu = true
  ## Whether to report total system cpu stats or not
  totalcpu = true
  ## If true, collect raw CPU time metrics
  collect_cpu_time = true
  ## If true, compute and report the sum of all non-idle CPU states
  report_active = true

# Read metrics about memory usage
[[inputs.mem]]
  # no configuration

# Read metrics about system load & uptime
[[inputs.system]]
  ## Uncomment to remove deprecated metrics.
```

```
  # fielddrop = ["uptime_format"]

# Read metrics about network interface usage
[[inputs.net]]
  ## By default, telegraf gathers stats from any up interface (excluding loopback)
  ## Setting interfaces will tell it to gather these explicit interfaces,
  ## regardless of status.
  ##
 # interfaces = ["eth0"]

# Read metrics about swap memory usage
[[inputs.swap]]
  # no configuration

# Read metrics about disk usage by mount point
[[inputs.disk]]
  ## By default stats will be gathered for all mount points.
  ## Set mount_points will restrict the stats to only the specified mount points.
  # mount_points = ["/"]

  ## Ignore mount points by filesystem type.
  # ignore_fs = ["tmpfs", "devtmpfs", "devfs", "iso9660", "overlay", "aufs", "squashfs"]

# Get the number of processes and group them by status
[[inputs.processes]]
  # no configuration

# Read metrics about disk IO by device
[[inputs.diskio]]
  ## By default, telegraf will gather stats for all devices including
  ## disk partitions.
  ## Setting devices will restrict the stats to the specified devices.
  # devices = ["sda", "sdb", "vd*"]
  ## Uncomment the following line if you need disk serial numbers.
  # skip_serial_number = false
```

### Windows Operating System

To collect Windows OS related metrics and properties as it is in a managed Telegraf agent installation flow and to have localized object types, names, alerts, symptoms, metrics and properties, and so on, the `telegraf.conf` file should be customized and have the following content.

For Windows OS, the `telegraf.conf` file default configurations were changed for Telegraf versions greater than or equal to 1.20.0 and includes Linux related input plugins, like `inputs.cpu, inputs.disk, inputs.diskio, inputs.kernel, inputs.mem, inputs.processes, inputs.swap, inputs.system`, and so on. They should be commented and those related to Windows should be uncommented.

```
[[inputs.win_perf_counters]]
PrintValid=true

[[inputs.win_perf_counters.object]]
```

```
  ObjectName = "Processor"
  Instances = ["*"]
  Counters = ["% Idle Time", "% Interrupt Time", "% Privileged Time", "% Processor Time", "%
User Time", "Interrupts/sec", "% DPC Time"]
  Measurement = "win.cpu"
  IncludeTotal = true

[[inputs.win_perf_counters.object]]
  ObjectName = "LogicalDisk"
  Instances = ["*"]
  Counters = ["% Disk Read Time", "% Disk Write Time", "% Free Space", "% Idle Time", "Avg.
Disk Bytes/Read", "Avg. Disk Bytes/Write", "Avg. Disk Queue Length", "Avg. Disk sec/Read",
"Avg. Disk sec/Write", "Avg. Disk Write Queue Length", "Avg. Disk Read Queue Length", "Free
Megabytes", "Split IO/Sec"]
  Measurement = "win.disk"

[[inputs.win_perf_counters.object]]
  ObjectName = "Memory"
  Counters = ["Available Bytes", "Cache Bytes", "Committed Bytes", "Cache Faults/sec",
"Demand Zero Faults/sec", "Page Faults/sec", "Pages/sec", "Transition Faults/sec", "Pool
Nonpaged Bytes", "Pool Paged Bytes"]
  Instances = ["------"]
  Measurement = "win.mem"

[[inputs.win_perf_counters.object]]
  ObjectName = "Network Interface"
  Counters = ["Bytes Received/sec", "Bytes Sent/sec", "Packets Outbound Discarded", "Packets
Outbound Errors", "Packets Received Discarded", "Packets Received Errors", "Packets Received/
sec", "Packets Sent/sec", "Connections Established"]
  Instances = ["*"]
  Measurement = "win.net"
  IncludeTotal = true

[[inputs.win_perf_counters.object]]
  ObjectName = "Paging File"
  Counters = ["% Usage"]
  Instances = ["*"]
  Measurement = "win.paging"
  IncludeTotal = true

  [[inputs.win_perf_counters.object]]
    ObjectName = "Process"
    Counters = ["% Privileged Time", "% Processor Time", "% User Time", "Elapsed Time",
"Handle Count", "IO Read Bytes/sec", "IO Read Operations/sec", "IO Write Bytes/sec", "IO
Write Operations/sec", "Private Bytes", "Thread Count", "Virtual Bytes", "Working Set",
"Working Set - Private"]
    Instances = ["_Total", "telegraf", "w3wp"]         # Replace this with a list of process
names that you want to monitor. "_Total" is all processes combined
    Measurement = "win.process"

[[inputs.win_perf_counters.object]]
  ObjectName = "System"
  Counters = ["Context Switches/sec", "Processes", "Processor Queue Length", "System Calls/
sec", "System Up Time", "Threads"]
  Instances = ["------"]
```

```
  Measurement = "win.system"

[[inputs.win_perf_counters.object]]
  ObjectName = "TCPv4"
  Counters = ["Connection Failures", "Connections Active", "Connections Established",
"Connections Passive", "Connection Reset", "Segments Received/sec", "Segments Retransmitted/
sec", "Segments Sent/sec"]
  Instances = ["------"]
  Measurement = "win.net.tcp"

[[inputs.win_perf_counters.object]]
  ObjectName = "TCPv6"
  Counters = ["Connection Failures", "Connections Active", "Connections Established",
"Connections Passive", "Connection Reset", "Segments Received/sec", "Segments Retransmitted/
sec", "Segments Sent/sec"]
  Instances = ["------"]
  Measurement = "win.net.tcp"

[[inputs.win_perf_counters.object]]
  ObjectName = "UDPv4"
  Counters = ["Datagrams No Port/sec", "Datagrams Received/Errors", "Datagrams Received/sec",
"Datagrams Sent/sec"]
  Instances = ["------"]
  Measurement = "win.net.udp"

[[inputs.win_perf_counters.object]]
  ObjectName = "UDPv6"
  Counters = ["Datagrams No Port/sec", "Datagrams Received/Errors", "Datagrams Received/sec",
"Datagrams Sent/sec"]
  Instances = ["------"]
  Measurement = "win.net.udp"
```

## Sample Configurations

If you do not use the helper script, you must provide specific configurations required on the end point to post metrics to cloud proxy. Based on whether the vCenter Server VM, the EC2 AWS instance, or the Azure VM is managed or unmanaged, the following configurations must be provided for open source Telegraf.You must provide correct values for the variables enclosed in <>.

If you do not use the helper script, you must download mandatory_tags.sh or mandatory_tags.bat and provide the path to the script that you have downloaded and the path to the Telegraf executable. For example, in the following code, you must provide the following: <Path to mandatory_tags.sh/bat> <path to telegraf executable>.

For Linux, download the mandatory_tags.sh from cloud proxy located at `https://<CP_IP>/downloads/salt/mandatory_tags.sh`.

For Windows, download the mandatory_tags.ps1 from cloud proxy located at `https://<CP_IP>/downloads/salt/mandatory_tags.ps1`.

**Note** For Windows and Linux, use the relevant cloud proxy IP address/FQDN for <CP_IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP

### Managed vCenter Server VM

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<VM_NAME/HOSTNAME>"
  omit_hostname = false

# Configuration for HTTP server to send metrics to
[[outputs.http]]
   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
   timeout = "5s"
   method = "POST"
   insecure_skip_verify = true
   data_format = "wavefront"
   ## Additional HTTP headers
   [outputs.http.headers]
     Content-Type = "text/plain; charset=utf-8"
     vmId = "<VM_MOR>"
     vcid = "<VC_ID>"
     hostname = "<VM_NAME/HOSTNAME>"
     uuid = ""
[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

### Managed AWS EC2 Instance

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
```

```
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<EC2_INSTANCE_NAME/HOSTNAME>"
  omit_hostname = true

# Configuration for HTTP server to send metrics to
[[outputs.http]]
  url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
  timeout = "5s"
  method = "POST"
  insecure_skip_verify = true
  data_format = "wavefront"
## Additional HTTP headers
[outputs.http.headers]
  Content-Type = "text/plain; charset=utf-8"
  vmId = "i-< EC2_INSTANCE_ID>"
  vcid = "<REAGON>"
  hostname = "<EC2_INSTANCE_NAME/HOSTNAME>"
  uuid = "aws"

[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

## Managed Azure VM

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<AZURE_VIRTUAL_MACHINE_NAME/HOSTNAME>"
  omit_hostname = true

# Configuration for HTTP server to send metrics to
[[outputs.http]]
   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
   timeout = "5s"
   method = "POST"
   insecure_skip_verify = true
   data_format = "wavefront"
   ## Additional HTTP headers
   [outputs.http.headers]
     Content-Type = "text/plain; charset=utf-8"
     vmId =
```

```
";subscriptions;<SUBSCRIPTION_ID>;resourceGroups;<AZURE_RESOURCE_GROUP>;providers;Microsoft.Co
mpute;virtualMachines;<AZURE_VIRTUAL_MACHINE_NAME>"
     vcid = "<SUBSCRIPTION_ID>"
     hostname = "<AZURE_VIRTUAL_MACHINE_NAME/HOSTNAME>"
     uuid = "azure"

[[inputs.exec]]
  commands = ["/bin/bash \"/etc/telegraf/telegraf.d/mandatory_tags.sh\" \"/usr/bb
in/telegraf\""]
  timeout = "5s"
  data_format = "influx"
```

### Unmanaged vCenter Server VMs, AWS EC2 Instances, and Azure VMs

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<VM_NAME/HOSTNAME>"
  omit_hostname = false

# Configuration for HTTP server to send metrics to
[[outputs.http]]
   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
   timeout = "5s"
   method = "POST"
   insecure_skip_verify = true
   data_format = "wavefront"
   ## Additional HTTP headers
   [outputs.http.headers]
     Content-Type = "text/plain; charset=utf-8"
     uuid = "<UUID>"
     ip = "<IP_ADDRESS>"
     hostname = "<VM_NAME/HOSTNAME>"
[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

## Telegraf Configuration Details for Supported Application Services

For application services and custom monitoring plugins supported by VMware Aria Operations, there are configuration details that you must follow.

## Table 4-58. Configuration Details for Supported Application Services

| Configuration Details for Supported Application Services |
| --- |
| **Active Directory**<br>Here are the configuration details:<br>`[[inputs.win_perf_counters]]`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "DirectoryServices"`<br>`  Counters = ["LDAP Active Threads","LDAP Client Sessions","LDAP Writes/sec","LDAP Searches/sec","LDAP Successful Binds/sec","LDAP New Connections/sec","LDAP Closed Connections/sec","LDAP UDP operations/sec","DS Threads in Use","DS Directory Writes/sec","DS Directory Reads/sec","DS Directory Searches/sec","DS Client Binds/sec","DS Server Binds/sec","DRA Pending Replication Synchronizations","DRA Sync Requests Made","DRA Sync Requests Successful","DRA Pending Replication Operations","DRA Inbound Objects/sec","DRA Inbound Bytes Total/sec","DRA Outbound Objects/sec","DRA Outbound Bytes Total/sec","Base Searches/sec","Database adds/sec","Database deletes/sec","Database modifys/sec","Database recycles/sec"]`<br>`  Instances = ["*"]`<br>`  Measurement = "ad.active.directory"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "Security System-Wide Statistics"`<br>`  Counters = ["NTLM Authentications","Kerberos Authentications","KDC AS Requests","KDC TGS Requests","Digest Authentications"]`<br>`  Instances = ["------"]`<br>`  Measurement = "ad.security.statistics"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "Database"`<br>`  Instances = ["*"]`<br>`  Counters = ["Database Cache % Hit","Database Cache Page Faults/sec","Database Page Faults/sec","Database Cache Page Fault Stalls/sec","Database Cache Size"]`<br>`  Measurement = "ad.database"`<br>`  #IncludeTotal=false #Set to true to include _Total instance when querying for all (*).`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "DFS Namespace Service Referrals"`<br>`  Instances = ["*"]`<br>`  Counters = ["Requests Processed","Requests Failed","Avg. Response Time"]`<br>`  Measurement = "ad.dfsn"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "DFS Replication Service Volumes"`<br>`  Instances = ["*"]`<br>`  Counters = ["Data Lookups","Database Lookups","Database Commits"]`<br>`  Measurement = "ad.dfs.replications"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "DNS"`<br>`  Counters = ["Dynamic Update Received","Dynamic Update Rejected","Recursive Queries","Recursive Queries Failure","Recursive Query Failure","Secure Update Failure","Total Query Received","Total Response Sent"]`<br>`  Instances = ["*"]`<br>`  Measurement = "ad.dns"` |
| **ActiveMQ** |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| Here are the configuration details: |

```
# Read Apache ActiveMQ information.
[[inputs.jolokia2_agent]]
  urls = ["http://localhost:8161/api/jolokia"]
  name_prefix = "activemq."
  plugin_name_override="activemq"
     username = "user"
     password = "user"

### JVM Generic

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_runtime"
    mbean = "java.lang:type=Runtime"
    paths = ["Uptime"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_memory"
    mbean = "java.lang:type=Memory"
    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]

  [[inputs.jolokia2_agent.metric]]
    name     = "jvm_garbage_collector"
    mbean    = "java.lang:name=*,type=GarbageCollector"
    paths    = ["CollectionTime", "CollectionCount"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name       = "jvm_memory_pool"
    mbean      = "java.lang:name=*,type=MemoryPool"
    paths      = ["Usage", "PeakUsage", "CollectionUsage"]
    tag_keys   = ["name"]
    tag_prefix = "pool_"

### ACTIVEMQ
  [[inputs.jolokia2_agent.metric]]
    name     = "queue"
    mbean    =
"org.apache.activemq:brokerName=*,destinationName=*,destinationType=Queue,type=Broker"
    paths    =
["QueueSize","EnqueueCount","ConsumerCount","DispatchCount","DequeueCount","ProducerCount","
InFlightCount"]
    tag_keys = ["brokerName","destinationName"]

  [[inputs.jolokia2_agent.metric]]
    name     = "topic"
    mbean    =
"org.apache.activemq:brokerName=*,destinationName=*,destinationType=Topic,type=Broker"
    paths    = ["ProducerCount","DequeueCount","ConsumerCount","QueueSize","EnqueueCount"]
    tag_keys = ["brokerName","destinationName"]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| ```
  [[inputs.jolokia2_agent.metric]]
    name     = "broker"
    mbean    = "org.apache.activemq:brokerName=*,type=Broker"
    paths    =
["TotalConsumerCount","TotalMessageCount","TotalEnqueueCount","TotalDequeueCount","MemoryLim
it","MemoryPercentUsage","StoreLimit","StorePercentUsage","TempPercentUsage","TempLimit"]
    tag_keys = ["brokerName"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
``` |
| **Apache HTTPD**<br>Here are the configuration details: |
| ```
# Read Apache status information (mod_status)
[[inputs.apache]]

  ## An array of URLs to gather from, must be directed at the machine
  ## readable version of the mod_status page including the auto query string.
  ## Default is "http://localhost/server-status?auto".
  ##
    ##urls = ["http://localhost/server-status?auto"]
  ##
  urls = ["http://127.0.0.1:80/server-status?auto"]
  ## Credentials for basic HTTP authentication.
    # username = "myuser"
    #password = "mypassword"

  ## Maximum time to receive response.
  # response_timeout = "5s"

  ## Optional SSL Config
      # ssl_ca = "/etc/telegraf/ca.pem"
    # ssl_cert = "/etc/telegraf/cert.pem"
    # ssl_key = "/etc/telegraf/key.pem"

  ## Use SSL but skip chain & host verification
    insecure_skip_verify = true
``` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Cassandra**<br>Here are the configuration details:<br>`# Read Cassandra metrics through Jolokia`<br>`[[inputs.cassandra]]`<br>`  context = "/jolokia/read"`<br>`  ## List of cassandra servers exposing jolokia read service`<br>`  servers = ["myuser:mypassword@10.10.10.1:8778","10.10.10.2:8778",":8778"]`<br>`  ## List of metrics collected on above servers`<br>`  ## Each metric consists of a jmx path.`<br>`  ## This will collect all heap memory usage metrics from the jvm and`<br>`  ## ReadLatency metrics for all keyspaces and tables.`<br>`  ## "type=Table" in the query works with Cassandra3.0. Older versions might`<br>`  ## need to use "type=ColumnFamily"`<br>`  metrics  = [`<br>`  "/java.lang:type=Memory/HeapMemoryUsage",`<br>`  "/org.apache.cassandra.metrics:type=Table,keyspace=*,scope=*,name=ReadLatency"`<br>`  ]` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Hyper-V**<br>Here are the configuration details:<br>[[inputs.win_perf_counters]]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Virtual Machine Health Summary"<br>Instances = ["------"]<br>Measurement = "hyperv.vm.health"<br>Counters = ["Health Ok", "Health Critical"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Dynamic Memory VM"<br>Instances = ["*"]<br>Measurement = "hyperv.vm.memory"<br>Counters = ["Physical Memory", "Added Memory", "Guest Visible Physical Memory"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Hypervisor Virtual Processor"<br>Instances = ["*"]<br>Measurement = "hyperv.hypervisor.virtual.processor"<br>Counters = ["% Guest Run Time", "% Hypervisor Run Time", "% Total Run Time"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Hypervisor Root Virtual Processor"<br>Instances = ["*"]<br>Measurement = "hyperv.hypervisor.root.virtual.processor"<br>Counters = ["% Guest Run Time", "% Hypervisor Run Time", "% Total Run Time"]<br>IncludeTotal = true<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Virtual IDE Controller (Emulated)"<br>Instances = ["*"]<br>Measurement = "hyperv.virtual.ide.controller"<br>Counters = ["Write Bytes/sec", "Read Bytes/sec", "Written Sectors/sec", "Read Sectors/sec"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Virtual Network Adapter"<br>Instances = ["*"]<br>Measurement = "hyperv.virtual.net.adapter"<br>Counters = ["Bytes/sec", "Bytes Received/sec", "Bytes Sent/Sec", "Packets Sent/sec", "Packets Received/sec", "Packets/sec"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Legacy Network Adapter"<br>Instances = ["*"]<br>Measurement = "hyperv.legacy.net.adapter"<br>Counters = ["Bytes Dropped", "Bytes Received/sec", "Bytes Sent/Sec"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Hypervisor Partition"<br>Instances = ["*"]<br>Measurement = "hyperv.hypervisor.partition"<br>Counters = ["Virtual Processors"]<br><br>[[inputs.win_perf_counters.object]]<br>ObjectName = "Hyper-V Virtual Storage Device"<br>Instances = ["*"]<br>Measurement = "hyperv.virtual.storage.device"<br>Counters = [ |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
                  "Maximum Bandwidth", "Read Bytes/sec", "Write Bytes/sec", "Queue Length",
                  "Lower Latency", "Minimum IO Rate", "Maximum IO Rate", "Latency",
"Throughput",
                  "Lower Queue Length", "Queue Length", "Normalized Throughput", "Write
Operations/Sec",
                  "Read Operations/Sec", "Write Bytes/sec", "Read Bytes/sec", "Error Count",
                  "Flush Count", "Write Count", "Read Count"
                ]

     [[inputs.win_perf_counters.object]]
       ObjectName = "Processor"
       Instances = ["*"]
       Counters = ["% Idle Time", "% Interrupt Time", "% Privileged Time", "% Processor
Time", "% User Time", "Interrupts/sec"]
       Measurement = "hyperv.host.cpu"
       IncludeTotal = true

     [[inputs.win_perf_counters.object]]
       ObjectName = "LogicalDisk"
       Instances = ["*"]
       Counters = ["% Disk Read Time", "% Disk Write Time", "% Free Space", "% Idle Time",
"Avg. Disk Bytes/Read", "Avg. Disk Bytes/Write", "Avg. Disk Queue Length", "Avg. Disk sec/
Read", "Avg. Disk sec/Write", "Avg. Disk Write Queue Length", "Free Megabytes", "Split IO/
Sec"]
       Measurement = "hyperv.host.disk"
       IncludeTotal = true

     [[inputs.win_perf_counters.object]]
       ObjectName = "Memory"
       Counters = ["Available Bytes", "Cache Bytes", "Committed Bytes", "Cache Faults/sec",
"Demand Zero Faults/sec", "Page Faults/sec", "Pages/sec", "Transition Faults/sec", "Pool
Nonpaged Bytes", "Pool Paged Bytes"]
       Instances = ["------"]
       Measurement = "hyperv.host.mem"

     [[inputs.win_perf_counters.object]]
       ObjectName = "Network Interface"
       Counters = ["Bytes Received/sec", "Bytes Sent/sec", "Packets Outbound Discarded",
"Packets Outbound Errors", "Packets Received Discarded", "Packets Received Errors",
"Packets Received/sec", "Packets Sent/sec", "Bytes Total/sec", "Current Bandwidth", "Output
Queue Length"]
       Instances = ["*"]
       Measurement = "hyperv.host.net"
       IncludeTotal = true

     [[inputs.win_perf_counters.object]]
       ObjectName = "System"
       Counters = ["Context Switches/sec", "Processes", "Processor Queue Length", "System
Calls/sec", "System Up Time", "Threads"]
       Instances = ["------"]
       Measurement = "hyperv.host.system"

     [[inputs.win_perf_counters.object]]
       ObjectName = "Process"
       Counters = ["% Privileged Time", "% Processor Time", "% User Time", "Elapsed Time",
"Handle Count", "IO Read Bytes/sec", "IO Read Operations/sec", "IO Write Bytes/sec", "IO
Write Operations/sec", "Private Bytes", "Thread Count", "Virtual Bytes", "Working Set",
"Working Set - Private"]
       Instances = ["_Total"]
       Measurement = "hyperv.host.process"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Java Plugin**<br>Here are the configuration details:<br>`[[inputs.jolokia2_agent]]`<br>`  # Prefix to attach to the measurement name`<br>`  name_prefix = "java."`<br>`  # Add agents URLs to query`<br>`  urls = ["http://localhost:8080/jolokia"]`<br>`  #username and password are mandatory for Jolokia 1.6 or later`<br>`  #username = <jolokia role username>`<br>`  #password = <jolokia role password>`<br>`  # response_timeout = "5s"`<br>`  ## Optional TLS config`<br>`  # tls_ca   = "/var/private/ca.pem"`<br>`  # tls_cert = "/var/private/client.pem"`<br>`  # tls_key  = "/var/private/client-key.pem"`<br>`  # insecure_skip_verify = false`<br>`  ### JVM Generic`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name  = "OperatingSystem"`<br>`    mbean = "java.lang:type=OperatingSystem"`<br>`    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name  = "jvm_runtime"`<br>`    mbean = "java.lang:type=Runtime"`<br>`    paths = ["Uptime"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name  = "jvm_memory"`<br>`    mbean = "java.lang:type=Memory"`<br>`    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name     = "jvm_garbage_collector"`<br>`    mbean    = "java.lang:name=*,type=GarbageCollector"`<br>`    paths    = ["CollectionTime", "CollectionCount"]`<br>`    tag_keys = ["name"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name       = "jvm_memory_pool"`<br>`    mbean      = "java.lang:name=*,type=MemoryPool"`<br>`    paths      = ["Usage", "PeakUsage", "CollectionUsage"]`<br>`    tag_keys   = ["name"]`<br>`    tag_prefix = "pool_"`<br>`  ### TOMCAT`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name     = "GlobalRequestProcessor"`<br>`    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"`<br>`    paths    = ["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]`<br>`    tag_keys = ["name"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name     = "JspMonitor"`<br>`    mbean    =`<br>`"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"`<br>`    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]`<br>`    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name     = "ThreadPool"`<br>`    mbean    = "Catalina:name=*,type=ThreadPool"`<br>`    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]`<br>`    tag_keys = ["name"]`<br>`  [[inputs.jolokia2_agent.metric]]`<br>`    name     = "Servlet"`<br>`    mbean    = "Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| `  paths    = ["processingTime","errorCount","requestCount"]`<br>`  tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]`<br>`[[inputs.jolokia2_agent.metric]]`<br>`  name     = "Cache"`<br>`  mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"`<br>`  paths    = ["hitCount","lookupCount"]`<br>`  tag_keys = ["context","host"]` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **JBoss Server**<br>Here are the configuration details:<br>`# Read JBoss status information (mod_status)`<br>`[[inputs.jolokia2_agent]]`<br>`  ##urls = ["http://localhost:8380/jolokia"]`<br>`  urls = ["http://127.0.0.1:8082/jolokia"]`<br>`  name_prefix = "jboss."`<br><br>`  ## Optional SSL Config`<br>`  # ssl_ca = "/etc/telegraf/ca.pem"`<br>`  # ssl_cert = "/etc/telegraf/cert.pem"`<br>`  # ssl_key = "/etc/telegraf/key.pem"`<br>`  ## Use SSL but skip chain & host verification`<br>`  insecure_skip_verify = true`<br><br>`  ### JVM Generic`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name = "OperatingSystem"`<br>`  mbean = "java.lang:type=OperatingSystem"`<br>`  paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name = "OperatingSystem"`<br>`  mbean = "java.lang:type=OperatingSystem"`<br>`  paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "jvm_runtime"`<br>`  mbean = "java.lang:type=Runtime"`<br>`  paths = ["Uptime"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "jvm_memory"`<br>`  mbean = "java.lang:type=Memory"`<br>`  paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "jvm_garbage_collector"`<br>`  mbean  = "java.lang:name=*,type=GarbageCollector"`<br>`  paths  = ["CollectionTime", "CollectionCount"]`<br>`  tag_keys = ["name"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "jvm_memory_pool"`<br>`  mbean  = "java.lang:name=*,type=MemoryPool"`<br>`  paths  = ["Usage", "PeakUsage", "CollectionUsage"]`<br>`  tag_keys  = ["name"]`<br>`  tag_prefix = "pool_"`<br><br>`  ### JBOSS`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "connectors.http"`<br>`  mbean  = "jboss.as:https-listener=*,server=*,subsystem=undertow"`<br>`  paths  = ["bytesReceived","bytesSent","errorCount","requestCount"]`<br>`  tag_keys = ["server","https-listener"]`<br><br>`  [[inputs.jolokia2_agent.metric]]`<br>`  name  = "connectors.http"` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
    mbean    = "jboss.as:http-listener=*,server=*,subsystem=undertow"
    paths    = ["bytesReceived","bytesSent","errorCount","requestCount"]
    tag_keys = ["server","http-listener"]

  [[inputs.jolokia2_agent.metric]]
    name     = "datasource.jdbc"
    mbean    = "jboss.as:data-source=*,statistics=jdbc,subsystem=datasources"
    paths    =
["PreparedStatementCacheAccessCount","PreparedStatementCacheHitCount","PreparedStatementCach
eMissCount"]
    tag_keys = ["data-source"]

  [[inputs.jolokia2_agent.metric]]
    name     = "datasource.pool"
    mbean    = "jboss.as:data-source=*,statistics=pool,subsystem=datasources"
    paths    = ["AvailableCount","ActiveCount","MaxUsedCount"]
    tag_keys = ["data-source"]

  [[inputs.jolokia2_agent.metric]]
    name = "thread.count"
    mbean  = "jboss.as:subsystem=*,thread-pool=*"
    paths =
["completedTaskCount","currentThreadCount","maxThreads","keepaliveTime","largestThreadCount"
,"activeCount","taskCount","rejectedCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount","TotalStartedThreadCount","DaemonThreadCount","PeakThreadCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Microsoft IIS**<br>Here are the configuration details:<br>`# Read MSIIS status information (mod_status)`<br>`[[inputs.win_perf_counters]]`<br><br>`[[inputs.win_perf_counters.object]]`<br>`# IIS, Web Service`<br>`ObjectName = "Web Service"`<br>`Counters = [`<br>`"Service Uptime", "Current Connections", "Bytes Sent/sec", "Total Bytes Sent", "Connection Attempts/sec",`<br>`"Bytes Received/sec", "Total Bytes Received", "Bytes Total/sec", "Total Bytes Transferred",`<br>`"Get Requests/sec","Total Get Requests", "Post Requests/sec","Total Post Requests",`<br>`"Put Requests/sec","Total Put Requests", "Delete Requests/sec","Total Delete Requests",`<br>`"Anonymous Users/sec", "NonAnonymous Users/sec", "Files Sent/sec", "Total Files Sent",`<br>`"Files Received/sec", "Total Files Received", "Files/sec", "Total Files Transferred",`<br>`"Not Found Errors/sec", "Locked Errors/sec", "Total Method Requests/sec"`<br>`]`<br>`Instances = ["*"]`<br>`Measurement = "iis.websvc"`<br>`IncludeTotal=true #Set to false to not include _Total instance.`<br><br>`[[inputs.win_perf_counters.object]]`<br>`# Web Service Cache / IIS`<br>`ObjectName = "Web Service Cache"`<br>`Counters = [`<br>`"Current Files Cached", "Active Flushed Entries", "Total Files Cached", "Total Flushed Files",`<br>`"File Cache Hits", "File Cache Misses", "File Cache Hits %", "File Cache Flushes",`<br>`"Current File Cache Memory Usage", "Maximum File Cache Memory Usage", "Current URIs Cached",`<br>`"Total URIs Cached", "Total Flushed URIs", "URI Cache Hits", "URI Cache Misses", "URI Cache Hits %",`<br>`"URI Cache Flushes", "Current Metadata Cached", "Total Metadata Cached", "Total Flushed Metadata",`<br>`"Metadata Cache Hits", "Metadata Cache Misses", "Metadata Cache Hits %", "Metadata Cache Flushes",`<br>`"Kernel: Current URIs Cached", "Kernel: Total URIs Cached", "Kernel: Total Flushed URIs",`<br>`"Kernel: URI Cache Hits", "Kernel: Uri Cache Hits/sec", "Kernel: URI Cache Misses", "Kernel: URI Cache Hits %",`<br>`"Kernel: URI Cache Flushes", "Cache Memory Usage"`<br>`]`<br>`Instances = ["*"]`<br>`Measurement = "iis.websvc.cache"`<br>`IncludeTotal=true #Set to false to not include _Total instance.`<br><br>`[[inputs.win_perf_counters.object]]`<br>`# IIS, ASP.NET`<br>`ObjectName = "ASP.NET"`<br>`Counters = ["Application Restarts","Request Wait Time","Requests Current","Requests Queued","Requests Rejected"]`<br>`Instances = ["*"]`<br>`Measurement = "iis.aspnet"` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
        #IncludeTotal=false #Set to true to include _Total instance when querying for all
(*).

      [[inputs.win_perf_counters.object]]
        # IIS, ASP.NET Applications
        ObjectName = "ASP.NET Applications"
        Counters = ["Cache Total Entries","Cache Total Hit Ratio","Output Cache
Entries","Errors Total/Sec","Requests Executing","Requests in Application Queue","Requests/
Sec"]
        Instances = ["*"]
        Measurement = "iis.aspnet.app"
        #IncludeTotal=false #Set to true to include _Total instance when querying for all
(*).

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR Exceptions"
        Counters = ["# of Exceps Thrown", "# of Exceps Thrown / Sec", "# of Filters / Sec",
"# of Finallys / Sec", "Throw to Catch Depth / Sec"]
        Instances = ["*"]
        Measurement = "iis.dotnet.exception"

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR Jit"
        Counters = ["% Time in Jit","IL Bytes Jitted / sec"]
        Instances = ["*"]
        Measurement = "iis.dotnet.jit"

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR Loading"
        Counters = ["% Time Loading"]
        Instances = ["*"]
        Measurement = "iis.dotnet.loading"

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR LocksAndThreads"
        Counters = ["# of current logical Threads","# of current physical Threads","# of
current recognized threads","# of total recognized threads","Queue Length / sec","Total #
of Contentions","Current Queue Length"]
        Instances = ["*"]
        Measurement = "iis.dotnet.lock"

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR Memory"
        Counters = ["% Time in GC","# Bytes in all Heaps","# Gen 0 Collections","# Gen
1 Collections","# Gen 2 Collections","# Induced GC","Allocated Bytes/sec","Finalization
Survivors","Gen 0 heap size","Gen 1 heap size","Gen 2 heap size","Large Object Heap
size","# of Pinned Objects"]
        Instances = ["*"]
        Measurement = "iis.dotnet.memory"

      [[inputs.win_perf_counters.object]]
        ObjectName = ".NET CLR Security"
        Counters = ["% Time in RT checks","Stack Walk Depth","Total Runtime Checks"]
        Instances = ["*"]
        Measurement = "iis.dotnet.security"

      [[inputs.win_perf_counters.object]]
        # HTTP Service request queues in the Kernel before being handed over to User Mode.
        ObjectName = "HTTP Service Request Queues"
        Instances = ["*"]
        Counters = ["CurrentQueueSize","RejectedRequests"]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| `Measurement = "iis.http.queues"`<br>`#IncludeTotal=false #Set to true to include _Total instance when querying for all (*).`<br><br>`[[inputs.win_perf_counters.object]]`<br>`# Process metrics, in this case for IIS only`<br>`ObjectName = "Process"`<br>`Counters = ["% Processor Time","Handle Count","Private Bytes","Thread Count","Virtual Bytes","Working Set"]`<br>`Instances = ["w3wp"]`<br>`Measurement = "iis.win.proc"`<br>`#IncludeTotal=false #Set to true to include _Total instance when querying for all (*).` |
| **Microsoft SQL Server**<br>Here are the configuration details:<br><br>`[[inputs.sqlserver]]`<br>`name_prefix = "MSSQL."`<br><br>`#servers = [`<br>`#"Server=<servername>;Port=1433;User Id=telegraf;Password=<mystrongpassword from step 2>;app name=telegraf;log=1;",`<br>`#"Server=<servername>;Port=1433;User Id=telegraf;Password=<mystrongpassword from step 2>;app name=telegraf;log=1;"`<br>`#]`<br>`servers = [`<br>`"Server=localhost;Port=1433;User Id=sa;Password=Password;app name=telegraf;log=1;"`<br>`]`<br><br>`namepass = ["Rows*writes*bytes*sec*", "Rows*reads*bytes*sec*", "Rows*writes*sec*", "Rows*reads*sec*", "Query*User*counter*", "Buffer*cache*hit*ratio*", "Page*life*expectancy*", "Page*lookups*sec*", "Page*reads*sec*", "Page*writes*sec*", "Lazy*writes*sec*", "Checkpoint*pages*sec*", "Log*Apply*Ready*Queue*", "Data*File*s*Size*KB*", "Log*File*s*", "XTP*Memory*Used*KB*", "Log*Flushes*sec*", "Write*Transactions*sec*", "Transactions*sec*", "Log*Flush*Wait*Time*", "Active*Transactions*", "Log*Bytes*Flushed*sec*", "Processes*blocked*", "User*Connections*", "Logins*sec*", "Logouts*sec*", "Active*Temp*Tables*", "Temp*Tables*Creation*Rate*", "Batch*Requests*sec*", "SQL*Compilations*sec*", "SQL*Re*Compilations*sec", "Stored*Procedures*Invoked*sec*", "Target*Server*Memory*KB*", "Total*Server*Memory*KB*", "SQL*Cache*Memory*KB*", "Log*Pool*Memory*KB*", "Connection*Memory*KB*", "Lock*Memory*KB*", "Memory*Grants*Pending*", "Active*memory*grant*amount*KB*", "Disk*Read*Bytes*sec*", "Disk*Read*IO*Throttled*sec*", "Disk*Read*IO*sec*", "Disk*Write*Bytes*sec*", "Disk*Write*IO*Throttled*sec*", "Used*memory*KB*", "CPU*usage*", "Free*Space*in*tempdb*KB*", "Version*Store*Size*KB*", "Transactions*", "Blocked*tasks*", "Active*requests*", "Queued*requests*", "Requests*completed*sec*", "Number*of*Deadlocks*sec*", "Lock*Wait*Time*ms*", "Lock*Waits*sec*", "Lock*Requests*sec*", "Average*Wait*Time*ms*", "Index*Searches*sec*", "Page*Splits*sec*", "Full*Scans*sec*", "CPU*", "Wait*time*ms*", "Wait*tasks*", "State*", "Recovery*Model*"]` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **MongoDB**<br>Here are the configuration details:<br>`# Read metrics from one or many MongoDB servers`<br>`[[inputs.mongodb]]`<br>`  ## An array of URI to gather stats about. Specify an ip or hostname`<br>`  ## with optional port add password. ie,`<br>`  ##   mongodb://user:auth_key@10.10.3.30:27017,`<br>`  ##   mongodb://10.10.3.33:18832,`<br>`  ##   10.0.0.1:10000, etc.`<br>`  ##servers = ["127.0.0.1:27017"]`<br>`  servers = ["localhost:27017"]`<br>`  gather_perdb_stats = true`<br>`  ## Optional SSL Config`<br>`    # ssl_ca = "/etc/telegraf/ca.pem"`<br>`    # ssl_cert = "/etc/telegraf/cert.pem"`<br>`    # ssl_key = "/etc/telegraf/key.pem"`<br>`  ## Use SSL but skip chain & host verification`<br>`    # insecure_skip_verify = false` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **MS Exchange Server**<br>Here are the configuration details: |

```
# Read MS Exchange status information (mod_status)
    [[inputs.win_services]]
      name_prefix = "msexchange."
      service_names = [
          "MSExchangeTransport",
          "MSExchangeTransportLogSearch",
          "MSExchangeUM",
          "MSExchangeUMCR",
          "MSExchangeThrottling",
          "MSExchangeServiceHost",
          "MSExchangeFastSearch",
          "MSExchangeRepl",
          "MSExchangeRPC",
          "MSExchangePop3",
          "MSExchangeSubmission",
          "MSExchangeDelivery",
          "MSExchangeMailboxReplication",
          "MSExchangeMailboxAssistants",
          "MSExchangeIS",
          "MSExchangeImap4",
          "MSExchangeHM",
          "MSExchangeFrontendTransport",
          "MSExchangeEdgeSync",
          "MSExchangeDiagnostics",
          "MSExchangeAntispamUpdate",
          "MSExchangeADTopology"
      ]


    [[inputs.win_perf_counters]]
      plugin_name_override="msexchange"
      [[inputs.win_perf_counters.object]]
        ObjectName = "Process"
        Counters = ["% Processor Time","Handle Count","Private Bytes","Thread
Count","Virtual Bytes","Working Set"]
        Instances = [
          "MSExchangeTransport",
          "MSExchangeTransportLogSearch",
          "umservice",
          "Microsoft.Exchange.UM.CallRouter",
          "MSExchangeThrottling",
          "Microsoft.Exchange.ServiceHost",
          "Microsoft.Exchange.Search.Service",
          "msexchangerepl",
          "Microsoft.Exchange.RpcClientAccess.Service",
          "Microsoft.Exchange.Pop3Service",
          "MSExchangeSubmission",
          "MSExchangeDelivery",
          "MSExchangeMailboxReplication",
          "MSExchangeMailboxAssistants",
          "Microsoft.Exchange.Store.Service",
          "Microsoft.Exchange.Imap4Service",
          "MSExchangeHMHost",
          "MSExchangeFrontendTransport",
          "Microsoft.Exchange.EdgeSyncSvc",
          "Microsoft.Exchange.Diagnostics.Service",
          "Microsoft.Exchange.AntispamUpdateSvc",
          "Microsoft.Exchange.Directory.TopologyService"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
        ]
        Measurement = "msexchange.process"

      [[inputs.win_perf_counters.object]]
        ObjectName = "Database"
        Counters = ["I/O Log Writes Average Latency","Log Threads Waiting","Log Record
Stalls/sec",
                "Database Cache % Hit","Database Page Fault Stalls/sec",
                    "I/O Database Writes Average Latency","I/O Database Reads Average
Latency"]
        Instances = ["*"]
        Measurement = "msexchange"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Database"
        Counters = [
          "I/O Log Writes Average Latency",
          "I/O Log Reads Average Latency",
          "Log Threads Waiting",
          "Log Record Stalls/sec",
          "Database Cache % Hit",
          "Database Page Fault Stalls/sec",
          "I/O Database Writes Average Latency",
          "I/O Database Reads Average Latency"
        ]
        Instances = ["*"]
        Measurement = "msexchange.database"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Database ==> Instances"
        Counters = [
          "Log Threads Waiting",
          "Log Record Stalls/sec",
          "I/O Database Writes Average Latency",
          "I/O Database Reads Average Latency"
        ]
        Instances = ["*"]
        Measurement = "msexchange.databaseinstance"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange ADAccess Domain Controllers"
        Counters = ["LDAP Read Time","LDAP Search Time","LDAP Searches Timed Out per
Minute","Long Running LDAP Operations/min"]
        Instances = ["*"]
        Measurement = "msexchange.adaccessdomaincontrollers"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeUMAvailability"
        Counters = ["Directory Access Failures","Total Inbound Calls Rejected by the UM
Service",
                "Total Inbound Calls Rejected by the UM Worker Process"]
        Instances = ["*"]
        Measurement = "msexchange.umavailability"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Availability Service"
        Counters = ["Average Time to Process a Free Busy Request"]
        Instances = ["*"]
        Measurement = "msexchange.availabilityservice"

      [[inputs.win_perf_counters.object]]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
        ObjectName = "MSExchange OWA"
        Counters = ["Current Unique Users", "Requests/sec", "Average Search Time"]
        Instances = ["*"]
        Measurement = "msexchange.owa"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeFrontEndTransport SmtpSend"
        Counters = ["Bytes Sent Total", "Messages Sent Total", "Message Bytes Sent Total"]
        Instances = ["*"]
        Measurement = "msexchange.frontendtransportsmtpsend"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeTransport SmtpSend"
        Counters = ["Message Bytes Sent Total", "Messages Sent Total"]
        Instances = ["_total"]
        Measurement = "msexchange.transportsmtpsend"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeTransport Queues"
        Counters = ["Active Mailbox Delivery Queue Length", "Retry Mailbox Delivery Queue
Length", "Submission Queue Length"]
        Instances = ["_total"]
        Measurement = "msexchange.transportqueues"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeIS Store"
        Counters = ["RPC Average Latency", "RPC Requests","RPC Operations/sec", ]
        Instances = ["*"]
        Measurement = "msexchange.isstore"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeIS Client Type"
        Counters = ["RPC Average Latency", "RPC Operations/sec" ]
        Instances = ["*"]
        Measurement = "msexchange.isclient"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Store Interface"
        Counters = ["RPC Latency average (msec)"]
        Instances = ["_total"]
        Measurement = "msexchange.isclientinterface"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange ADAccess Processes"
        Counters = ["LDAP Read Time", "LDAP Search Time"]
        Instances = ["*"]
        Measurement = "msexchange.adaccessprocesses"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange ActiveSync"
        Counters = ["Mailbox Search Total", "Requests/sec", "Ping Commands Pending",
"Current Requests",
                    "Average Request Time", "Sync Commands/sec"]
        Instances = ["*"]
        Measurement = "msexchange.activesync"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Active Manager Server"
        Counters = ["Server-Side Calls/sec", "Active Manager Database State writes to
Persistent storage/sec",
                    "GetServerForDatabase Server-Side Calls", "Total Number of Databases",
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
"Active Manager Role"]
        Instances = ["*"]
        Measurement = "msexchange.activemanagerserver"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeAutodiscover"
        Counters = ["Requests/sec"]
        Instances = ["*"]
        Measurement = "msexchange.autodiscover"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeWS"
        Counters = ["Requests/sec"]
        Instances = ["*"]
        Measurement = "msexchange.ws"

      [[inputs.win_perf_counters.object]]
        ObjectName = "Web Service"
        Counters = ["Current Connections","Connection Attempts/sec","Other Request Methods/
sec"]
        Instances = ["*"]
        Measurement = "msexchange.ws"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange WorkloadManagement Workloads"
        Counters = ["ActiveTasks","CompletedTasks","QueuedTasks"]
        Instances = ["*"]
        Measurement = "msexchange.workload"

      [[inputs.win_perf_counters.object]]
        ObjectName = "ASP.NET"
        Counters = ["Application Restarts","Worker Process Restarts","Request Wait Time"]
        Instances = ["*"]
        Measurement = "msexchange.aspnet"

      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange RpcClientAccess"
        Counters = ["RPC Averaged Latency", "RPC Requests", "Active User Count",
"Connection Count",
                    "RPC Operations/sec", "User Count"]
        Instances = ["*"]
        Measurement = "msexchange.rpcclient"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **MySQL**<br>Here are the configuration details: |

```
[[inputs.mysql]]

 ## specify servers via a url matching:
 ##  [username[:password]@][protocol[(address)]]/[?tls=[true|false|skip-verify]]
 ##  see https://github.com/go-sql-driver/mysql#dsn-data-source-name
 ##  e.g.
 ##    servers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"]
 ##    servers = ["user@tcp(127.0.0.1:3306)/?tls=false"]
 #
 ## If no servers are specified, then localhost is used as the host.
 ##servers = ["MS_USER:MS_PASSWORD@tcp(localhost:MS_PORT)/?tls=false"]
 servers = ["root:VMware1!@tcp(localhost:3306)/?tls=false"]
 ## the limits for metrics form perf_events_statements
 perf_events_statements_digest_text_limit  = 120
 perf_events_statements_limit              = 250
 perf_events_statements_time_limit         = 86400
 #
 ## if the list is empty, then metrics are gathered from all databasee tables

 #
 ## gather metrics from INFORMATION_SCHEMA.TABLES for databases provided above list
 gather_table_schema                       = true
 #
 ## gather thread state counts from INFORMATION_SCHEMA.PROCESSLIST
 gather_process_list                       = true
 #
 ## gather auto_increment columns and max values from information schema
 gather_info_schema_auto_inc               = true
 #
 ## gather metrics from SHOW SLAVE STATUS command output
 gather_slave_status                       = true
 #

 ## gather metrics from SHOW BINARY LOGS command output

   gather_binary_logs                        = false

 ## gather thread state counts from INFORMATION_SCHEMA.USER_STATISTICS
 gather_user_statistics                    = false
 #
 ## gather metrics from PERFORMANCE_SCHEMA.TABLE_IO_WAITS_SUMMARY_BY_TABLE
 gather_table_io_waits                     = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.TABLE_LOCK_WAITS
 gather_table_lock_waits                   = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.TABLE_IO_WAITS_SUMMARY_BY_INDEX_USAGE
 gather_index_io_waits                     = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.EVENT_WAITS
 gather_event_waits                        = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.FILE_SUMMARY_BY_EVENT_NAME
 gather_file_events_stats                  = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.EVENTS_STATEMENTS_SUMMARY_BY_DIGEST
 gather_perf_events_statements             = true
 #
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| ```<br>## Some queries we may want to run less often (such as SHOW GLOBAL VARIABLES)<br>interval_slow = "30m"<br><br>## Optional SSL Config<br><br># ssl_ca = "/etc/telegraf/ca.pem"<br><br># ssl_cert = "/etc/telegraf/cert.pem"<br><br># ssl_key = "/etc/telegraf/key.pem"<br><br>fieldpass = ["aborted_clients", "bytes_sent", "bytes_received",<br>"binlog_cache_size", "connection_errors_accept", "connection_errors_internal",<br>"connection_errors_max_connections", "queries", "threads_cached", "threads_connected",<br>"threads_running", "uptime", "delayed_insert_limit", "delayed_insert_timeout",<br>"delayed_queue_size", "max_connect_errors", "max_connections", "max_delayed_threads",<br>"max_error_count", "binary_size_bytes", "binary_files_count", "commands_show_processlist",<br>"access_denied", "bytes_received", "bytes_sent", "commit_transactions",<br>"concurrent_connections", "connected_time", "denied_connections", "denied_connections",<br>"rollback_transactions", "rows_fetched", "rows_updated", "select_commands",<br>"total_connections", "table_io_waits_total_fetch", "table_io_waits_total_insert",<br>"table_io_waits_total_update", "table_io_waits_total_delete", "connections",<br>"aborted_connects", "events_waits_seconds_total", "events_waits_total",<br>"innodb_print_all_deadlocks", "innodb_open_files", "innodb_buffer_pool_size",<br>"innodb_row_lock_time_avg", "innodb_row_lock_current_waits", "innodb_row_lock_time_max",<br>"innodb_table_locks", "innodb_row_lock_waits", "innodb_buffer_pool_dump_status",<br>"innodb_buffer_pool_load_status", "innodb_buffer_pool_pages_data",<br>"innodb_buffer_pool_bytes_data", "innodb_buffer_pool_pages_dirty",<br>"innodb_buffer_pool_bytes_dirty", "innodb_buffer_pool_pages_flushed", "innodb_checksums",<br>"events_statements_no_index_used_total", "*io_waits_seconds_total_fetch",<br>"read_high_priority", "write_concurrent_insert"]<br>```<br>**NGINX**<br>Here are the configuration details:<br>```<br># Read Nginx's basic status information (ngx_http_stub_status_module)<br>[[inputs.nginx]]<br><br>## An array of Nginx stub_status URI to gather stats.<br>urls = ["http://127.0.0.1:8081/nginx_status"]<br><br>## Optional SSL Config<br># ssl_ca = "/etc/telegraf/ca.pem"<br># ssl_cert = "/etc/telegraf/cert.pem"<br># ssl_key = "/etc/telegraf/key.pem"<br><br>## Use SSL but skip chain & host verification<br># insecure_skip_verify = false<br><br>## HTTP response timeout (default: 5s)<br>response_timeout = "5s"<br>```<br> |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **NTPD**<br>Here are the configuration details:<br>`[[inputs.ntpq]]`<br>`name_prefix = “ntpd.”`<br>`## If false, set the -n ntpq flag. Can reduce metric gather times.`<br>`dns_lookup = tr` |
| **OracleDB**<br>■ Complete the prerequisites mentioned in the Pre-Requirements for Application Services.<br>■ Download `wavefront_oracle_metrics.py` from cloud proxy to the Telegraf configuration directory. For example, `C:\Telegraf\telegraf-1.20.4\telegraf.d`.<br>  ■ Invoke-WebRequest `https://$CP_IP/downloads/salt/content-accessories/wavefront_oracle_metrics.py -OutFile C:\Telegraf\telegraf-1.20.4\telegraf.d\wavefront_oracle_metrics.py`<br>Here are the configuration details:<br>`# Read metrics exposed by chef`<br>`[[inputs.exec]]`<br>`# commands = ['python "wavefront_oracle_metrics.py script path" -u "OracleDB Username" -p "OracleDB Password" -s "OracleDB SID/DSN"']`<br>`commands = ['python "C:\Telegraf\telegraf-1.20.4\telegraf.d\wavefront_oracle_metrics.py" -u "orcl.vrops.es" -p "Qw123456" -s "https://VAP-test.vrops.es:1158/em"']`<br>`timeout = "180s"`<br>`data_format = "influx"`<br>`name_prefix = "oracledb."` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **PostgreSQL**<br>Here are the configuration details:<br>`# Read metrics from one or many postgresql servers`<br>`[[inputs.postgresql]]`<br>`  ## specify address via a url matching:`<br>`  ##   postgres://[pqgotest[:password]]@localhost[/dbname]\`<br>`  ##       ?sslmode=[disable\|verify-ca\|verify-full]`<br>`  ## or a simple string:`<br>`  ##   host=localhost user=pqotest password=... sslmode=... dbname=app_production`<br>`  ##`<br>`  ## All connection parameters are optional.`<br>`  ##`<br>`  ## Without the dbname parameter, the driver will default to a database`<br>`  ## with the same name as the user. This dbname is just for instantiating a`<br>`  ## connection with the server and doesn't restrict the databases we are trying`<br>`  ## to grab metrics for.`<br>`  ##`<br>`  address = "host=localhost user=postgres password=\t Password1! port=5432"`<br>`  ##address = "host=localhost user=postgres  password=\t Password1! port=5432"`<br>`  ## A  list of databases to explicitly ignore.  If not specified, metrics for all`<br>`  ## databases are gathered.  Do NOT use with the 'databases' option.`<br>`    # ignored_databases = ["postgres", "template0", "template1"]`<br>`  ## A list of databases to pull metrics about. If not specified, metrics for all`<br>`  ## databases are gathered.  Do NOT use with the 'ignored_databases' option.`<br>`    # databases = ["app_production", "testing"]`<br>`  ## Optional SSL Config`<br>`    # ssl_ca = "/etc/telegraf/ca.pem"`<br>`    # ssl_cert = "/etc/telegraf/cert.pem"`<br>`    # ssl_key = "/etc/telegraf/key.pem"`<br>`  ## Use SSL but skip chain & host verification`<br>`    # insecure_skip_verify = false` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Pivotal**<br>Here are the configuration details: |

```
# Read Tomcat status information (mod_status)
[[inputs.jolokia2_agent]]
  ## urls = ["http://localhost:8080/jolokia"]
  urls = ["http://localhost:8080/jolokia"]
  name_prefix = "pivotalserver."

  ## Optional SSL Config
    # ssl_ca = "/etc/telegraf/ca.pem"
    # ssl_cert = "/etc/telegraf/cert.pem"
    # ssl_key = "/etc/telegraf/key.pem"

  ## Use SSL but skip chain & host verification
    # insecure_skip_verify = false

  ### JVM Generic

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_runtime"
    mbean = "java.lang:type=Runtime"
    paths = ["Uptime"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_memory"
    mbean = "java.lang:type=Memory"
    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]

  [[inputs.jolokia2_agent.metric]]
    name     = "jvm_garbage_collector"
    mbean    = "java.lang:name=*,type=GarbageCollector"
    paths    = ["CollectionTime", "CollectionCount"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name       = "jvm_memory_pool"
    mbean      = "java.lang:name=*,type=MemoryPool"
    paths      = ["Usage", "PeakUsage", "CollectionUsage"]
    tag_keys   = ["name"]
    tag_prefix = "pool_"

  ### TOMCAT
  [[inputs.jolokia2_agent.metric]]
    name     = "GlobalRequestProcessor"
    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"
    paths    = ["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name     = "JspMonitor"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"
    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]
    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]
    name     = "ThreadPool"
    mbean    = "Catalina:name=*,type=ThreadPool"
    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name     = "Servlet"
    mbean    = "Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"
    paths    = ["processingTime","errorCount","requestCount"]
    tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]
    name     = "Cache"
    mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"
    paths    = ["hitCount","lookupCount"]
    tag_keys = ["context","host"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **RabbitMQ**<br>Here are the configuration details:<br>`# Read metrics from one or many RabbitMQ servers via the management API`<br>`[[inputs.rabbitmq]]`<br>`arc_service_id = "145d46c3-3f5d-449b-af93-f32a150c4965"`<br>`  ##url = "http://localhost:15672"`<br>`  url = "http://10.218.135.219:15672"`<br>`  metric_exclude = ["federation"]`<br>`  # name = "rmq-server-1" # optional tag`<br>`  username = "admin"`<br>`  password = "admin"`<br>`  ## Optional SSL Config`<br>`  # ssl_ca = "/etc/telegraf/ca.pem"`<br>`  # ssl_cert = "/etc/telegraf/cert.pem"`<br>`  # ssl_key = "/etc/telegraf/key.pem"`<br>`  ## Use SSL but skip chain & host verification`<br>`  # insecure_skip_verify = false`<br>`  ## Optional request timeouts`<br>`  ##`<br>`  ## ResponseHeaderTimeout, if non-zero, specifies the amount of time to wait`<br>`  ## for a server's response headers after fully writing the request.`<br>`  # header_timeout = "3s"`<br>`  ##`<br>`  ## client_timeout specifies a time limit for requests made by this client.`<br>`  ## Includes connection time, any redirects, and reading the response body.`<br>`  # client_timeout = "4s"`<br>`  ## A list of nodes to pull metrics about. If not specified, metrics for`<br>`  ## all nodes are gathered.`<br>`  # nodes = ["rabbit@node1", "rabbit@node2"]` |
| **Riak**<br>Here are the configuration details:<br>`# Read Riak status information (mod_status)`<br>`[[inputs.riak]]`<br>`  # Specify a list of one or more riak http servers`<br>`  ##servers = ["http://localhost:8098"]`<br>`  servers = ["http://127.0.0.1:8098"]` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **SharePoint Server**<br>Here are the configuration details:<br>`[[inputs.win_services]]`<br>`  name_prefix = "sharepoint."`<br><br>`  service_names = [`<br>`  "OSearch15",`<br>`  "SPAdminV4",`<br>`  "SPSearchHostController",`<br>`  "SPTimerV4",`<br>`  "SPTraceV4",`<br>`  "SPUserCodeV4",`<br>`  "SPWriterV4"`<br>`]`<br><br>`[[inputs.win_perf_counters]]`<br>`  plugin_name_override="sharepoint"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "Process"`<br>`  Instances = ["mssearch", "WSSADMIN", "hostcontrollerservice", "OWSTIMER", "wsstracing", "SPUCHostService", "SPWRITER"]`<br>`  Counters = ["% Processor Time", "Thread Count"]`<br>`  Measurement = "sharepoint.process"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "Web Service"`<br>`  Instances = ["*"]`<br>`  Counters = ["Bytes Sent/sec", "Bytes Received/sec"]`<br>`  Measurement = "sharepoint.webservice"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "SharePoint Foundation"`<br>`  Counters = ["Object Cache Hit %", "Executing Sql Queries", "Executing Time/Page Request", "Reject Page Requests Rate", "Incoming Page Requests Rate", "Active Threads", "Object Cache Hit Count", "Current Page Requests", "Sql Query Executing  time", "Responded Page Requests Rate"]`<br>`  Instances = ["_total"]`<br>`  Measurement = "sharepoint.Foundation"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "SharePoint Records Management Counters"`<br>`  Counters = ["Search results processed / sec"]`<br>`  Instances = ["------"]`<br>`  Measurement = "sharepoint.RecordsManagementCounters"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "SharePoint Foundation Search Gatherer Projects"`<br>`  Counters = ["Crawls in progress", "Filtered Text Rate", "Filtered Office Rate", "Filtered HTML Rate", "Accessed File Rate", "Accessed HTTP Rate", "File Errors Rate", "HTTP Errors Rate"]`<br>`  Instances = ["_total"]`<br>`  Measurement = "sharepoint.FoundationSearchGathererProjects"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`  ObjectName = "SharePoint Foundation Search Schema Plugin"`<br>`  Counters = ["Total Documents"]`<br>`  Instances = ["_total"]`<br>`  Measurement = "sharepoint.FoundationSearchSchemaPlugin"` |

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation BDC Online"
    Counters = ["Total calls failed", "Total calls per second"]
    Instances = ["------"]
    Measurement = "sharepoint.FoundationBDCOnline"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search Gatherer"
    Counters = ["Changes Processed", "Threads Committing Transactions", "Time Outs",
"Active Queue Length", "Idle Threads", "Heartbeats Rate"]
    Instances = ["------"]
    Measurement = "sharepoint.FoundationSearchGatherer"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search Indexer Plugin"
    Counters = ["Propagation Rate", "Average Query Latency", "Queries Succeeded", "Queries
Failed", "Queries"]
    Instances = ["_total"]
    Measurement = "sharepoint.FoundationSearchIndexerPlugin"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search Query Processor"
    Counters = ["Security Descriptor Cache Misses"]
    Instances = ["------"]
    Measurement = "sharepoint.FoundationSearchQueryProcessor"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search FAST Content Plugin"
    Counters = ["Batches Failed Timeout", "Submission Timeouts", "Items Failed Total",
"Items Failed Timeout"]
    Instances = ["------"]
    Measurement = "sharepoint.FoundationSearchFASTContentPlugin"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search Archival Plugin"
    Counters = ["Queues Committing", "Queues Waiting", "Queues Filtering", "Queues
Available"]
    Instances = ["_total"]
    Measurement = "sharepoint.FoundationSearchArchivalPlugin"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation BDC Metadata"
    Counters = ["Cache misses per second", "Cache hits per second"]
    Instances = ["_total"]
    Measurement = "sharepoint.FoundationBDCMetadata"

  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation Search Gatherer Databases"
    Counters = ["Documents in the crawl history", "Documents in the crawl queue"]
    Instances = ["_total"]
    Measurement = "sharepoint.FoundationSearchGathererDatabases"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| **Tomcat Server**<br>Here are the configuration details: |

```
# Read Tomcat status information (mod_status)
[[inputs.jolokia2_agent]]
  ##urls = ["http://localhost:8080/jolokia"]
  urls = ["http://localhost:8083/jolokia"]
  name_prefix = "tomcat."

  ## Optional SSL Config
    # ssl_ca = "/etc/telegraf/ca.pem"
    # ssl_cert = "/etc/telegraf/cert.pem"
    # ssl_key = "/etc/telegraf/key.pem"

  ## Use SSL but skip chain & host verification
    # insecure_skip_verify = false

  ### JVM Generic

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]

  [[inputs.jolokia2_agent.metric]]
    name = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_runtime"
    mbean = "java.lang:type=Runtime"
    paths = ["Uptime"]

  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_memory"
    mbean = "java.lang:type=Memory"
    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]

  [[inputs.jolokia2_agent.metric]]
    name     = "jvm_garbage_collector"
    mbean    = "java.lang:name=*,type=GarbageCollector"
    paths    = ["CollectionTime", "CollectionCount"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name       = "jvm_memory_pool"
    mbean      = "java.lang:name=*,type=MemoryPool"
    paths      = ["Usage", "PeakUsage", "CollectionUsage"]
    tag_keys   = ["name"]
    tag_prefix = "pool_"

  ### TOMCAT
  [[inputs.jolokia2_agent.metric]]
    name     = "GlobalRequestProcessor"
    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"
    paths    = ["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name     = "JspMonitor"
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |

```
    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"
    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]
    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]
    name     = "ThreadPool"
    mbean    = "Catalina:name=*,type=ThreadPool"
    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]
    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]
    name     = "Servlet"
    mbean    = "Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"
    paths    = ["processingTime","errorCount","requestCount"]
    tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]
    name     = "Cache"
    mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"
    paths    = ["hitCount","lookupCount"]
    tag_keys = ["context","host"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

## Table 4-58. Configuration Details for Supported Application Services (continued)

| Configuration Details for Supported Application Services |
| --- |
| Oracle WebLogic |
| **WebSphere**<br>Here are the configuration details:<br>`[[inputs.httpjson]]`<br>`  servers = ["http://localhost:9080/wasmonitor/stats"]`<br>`  method = "POST"`<br>`  name_override="websphere"`<br>`  response_timeout = "60s"`<br>`  tag_keys = ["name", "process", "node", "cell", "mbeanIdentifier"]`<br>`  [inputs.httpjson.headers]`<br>`  Authorization = "Basic d2F2ZWZyb250OndhdmVmcm9udA=="`<br>`[[inputs.exec]]`<br>`  commands = ['C:\\VMware\\UCP\\ucp-minion\\bin\\run-py-script.bat "C:\VMware\UCP\ucp-minion\lib\site-packages\ucp\application_availability.py" --config="C:\VMware\UCP\salt\conf\grains" "java" "http://localhost:9080"']`<br>`  timeout = "180s"`<br>`  data_format = "influx"`<br>`  name_prefix = "websphere."`<br>`  [inputs.exec.tags]`<br>`  server = "http://localhost:9080/wasmonitor/stats"` |

## Table 4-59. Configuration Details for Custom Monitoring Plugins

| |
|---|
| **Ping check**<br>Here are the configuration details:<br>`[[inputs.ping]]`<br>`## Hosts to send ping packets to.`<br>`urls = [www.vmware.com]`<br><br>`## Number of ping packets to send per interval.  Corresponds to the "-c"`<br>`## option of the ping command.`<br>`# count = 1`<br><br>`## Time to wait between sending ping packets in seconds.  Operates like the`<br>`## "-i" option of the ping command.`<br>`# ping_interval = 1.0`<br><br>`## If set, the time to wait for a ping response in seconds.  Operates like`<br>`## the "-W" option of the ping command.`<br>`# timeout = 1.0`<br><br>`## If set, the total ping deadline, in seconds.  Operates like the -w option`<br>`## of the ping command.`<br>`# deadline = 10`<br><br>`## Interface or source address to send ping from.  Operates like the -I or -S`<br>`## option of the ping command.`<br>`# interface = ""` |
| **TCP check**<br>Here are the configuration details:<br>`[[inputs.net_response]]`<br>`name_override="netresponse"`<br>`## Server address (default localhost)`<br>`address = "10.10.10.10:443"`<br>`protocol = "tcp"`<br><br>`## Set timeout`<br>`# timeout = "1s"`<br><br>`## Set read timeout (only used if expecting a response)`<br>`# read_timeout = "1s"` |

## Table 4-59. Configuration Details for Custom Monitoring Plugins (continued)

| |
|---|
| **HTTP check**<br>Here are the configuration details:<br><br>```<br># HTTP/HTTPS request given an address a method and a timeout<br>[[inputs.http_response]]<br>name_override="httpresponse"<br>## List of url to query.<br>address = https://10.10.10.10/suite-api/api/adapters<br>method = "GET"<br><br>## Set http_proxy.<br>## Telegraf uses the system wide proxy settings if it's is not set.<br># http_proxy = http://localhost:8888<br><br>## Set response_timeout (default 5 seconds)<br># response_timeout = "5s"<br><br>## Whether to follow redirects from the server (defaults to false)<br># follow_redirects = false<br><br>## Optional name of the field that will contain the body of the response.<br>## By default it is set to an empty String indicating that the body's content won't be added<br># response_body_field = ''<br><br>## Optional substring or regex match in body of the response (case sensitive)<br># response_string_match = "\"service_status\": \"up\""<br># response_string_match = "ok"<br># response_string_match = "\".*_status\".?:.?\"up\""<br><br><br>## Optional SSL Config<br># ssl_ca = "/etc/telegraf/ca.pem"<br># ssl_cert = "/etc/telegraf/cert.pem"<br># ssl_key = "/etc/telegraf/key.pem"<br><br><br>## Use SSL but skip chain & host verification<br>insecure_skip_verify = true<br><br>## HTTP Request Headers (all values must be strings)<br>[inputs.http_response.headers]<br>accept = "application/json" Authorization = "OpsToken a16f7a2b-<br>b033-48bc-9827-2daf8e205537::ec11ee5f-8623-4558-a904-8b4ea3f6f47d"<br>```<br><br> |

## Table 4-59. Configuration Details for Custom Monitoring Plugins (continued)

| |
|---|
| **Custom scripts**<br>Here are the configuration details:<br>**For Linux Platforms:**<br>`# Read metrics from command that can output to stdout`<br>`[[inputs.exec]]`<br>`name_prefix = "executescript."`<br>`name_override = "output"`<br>`# commands = ["<prefix> <file path> <argument>"]`<br>`commands = ["python /opt/scripts/argument.py 10"]`<br><br>`## Data format to consume.`<br>`data_format = "value"`<br>`data_type = "integer"`<br><br>`## Timeout for each command to complete.`<br>`timeout = "300s"`<br><br>`[inputs.exec.tags]`<br><br>`file_path = "/opt/scripts/argument.py"`<br>`script_name = "my-py-script on Centos7-VM"`<br><br>**For Windows Platforms:**<br>`# Read metrics from command that can output to stdout`<br>`[[inputs.exec]]`<br><br>`name_prefix = "executescript."`<br>`name_override = "output"`<br><br>`# commands = ["<prefix> <file path> <argument>"]`<br>`commands = ["Powershell -File C:\\\\VMware\\\\Scripts\\\\Arguments.ps1 vmtoolsd"]`<br><br>`## Data format to consume.`<br>`data_format = "value"`<br>`data_type = "integer"`<br><br>`## Timeout for each command to complete.`<br>`timeout = "300s"`<br><br>`[inputs.exec.tags]`<br><br>`file_path = "C:\\VMware\\Scripts\\Arguments.ps1"`<br>`script_name = "my-script-for-vmtool on Windows2022-PhysicalServer"` |

## Table 4-59. Configuration Details for Custom Monitoring Plugins (continued)

| |
|---|
| **Linux Process**<br>Here are the configuration details (via regular expression):<br>`# Monitor process cpu and memory usage`<br>`[[inputs.procstat]]`<br>`  ## pattern as argument for pgrep (ie, pgrep -f <pattern>)`<br>`  pattern = ".*ucp-.*"`<br>`  fieldpass =  ["running", "cpu_usage", "memory_usage"]`<br>`  [inputs.procstat.tags]`<br>`  search_pattern = "regex_.#!AsTeRiSk!#ucp-.#!AsTeRiSk!#"`<br>Here are the configuration details (via executable name):<br>`# Monitor process cpu and memory usage`<br>`[[inputs.procstat]]`<br>`  ## executable name (ie, pgrep <exe>)`<br>`  exe = "top"`<br>`  fieldpass =  ["running", "cpu_usage", "memory_usage"]`<br>`  [inputs.procstat.tags]`<br>`  search_pattern = "exec_top"`<br>Here are the configuration details (via process ID):<br>`# Monitor process cpu and memory usage`<br>`[[inputs.procstat]]`<br>`  ## PID file to monitor process`<br>`  pid_file = "/var/run/vmtoolsd.pid"`<br>`  fieldpass =  ["running", "cpu_usage", "memory_usage"]`<br>`  [inputs.procstat.tags]`<br>`  search_pattern = "pidfile_/var/run/vmtoolsd.pid"` |
| **Windows Service**<br>Here are the configuration details (via Windows service name):<br>`# Monitor process cpu and memory usage`<br>`[[inputs.procstat]]`<br>`## Windows service name`<br>`win_service = "Dhcp"`<br>`pid_finder = "native"`<br>`fieldpass =  ["running", "cpu_usage", "memory_usage"]`<br>`[inputs.procstat.tags]`<br>`search_pattern = "exec_Dhcp"`<br>`[[inputs.win_services]]`<br>`service_names = ["Dhcp"]`<br>`name_override="procstat"`<br>`[inputs.win_services.tags]`<br>`search_pattern = "exec_Dhcp"` |

## Troubleshooting (Open Source Telegraf)

### Windows Services

If you configure multiple services and the service names include underscores, then the associated objects have the same display name.

#### Problem

While monitoring a Windows Services plugin using open source Telegraf, if there are multiple services with names such as, `svc1_name1` and `svc1_name2` with an underscore in the service names, then both the objects have the same display name.

If you want the display names to have different names, follow these steps.

#### Solution

1. Manually edit the service name in VMware Aria Operations after the Windows Services object is created so that the correct service name is displayed as the display name of the Service object.
2. Navigate to **Environment > Inventory > Objects**. Select the Service object with the incorrect name and edit it.
3. Add the value as entered in the **Service Name** field to the **Display Name** field.

   For example, while monitoring a Windows service with the service name **`SMS_SITE_COMPONENT_MANAGER`**, if the display name of the current object is **`Services exec SMS on Windows_OS_on_<windows hostname>`**, where **`SMS`** is the display name, change the display name to **`Services exec SMS_SITE_COMPONENT_MANAGER on Windows_OS_on_<windows hostname>`**.
4. Click **OK**.

## Monitoring Physical Servers

If you want to start monitoring physical servers, use product-managed Telegraf agents. Use the helper script to install the product-managed agents. For more information, see Install/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform . Existing configurations for user-managed Telegraf agents that are used to monitor physical servers is supported, data collection will continue, and additional steps are not required.

## Service and Application Discovery

You can discover services and applications using the Service Discovery adapter.

## Service Discovery

Service discovery helps you discover services running in each VM and then builds a relationship or dependency between the services from different VMs. You can view basic metrics based on the services you want to monitor. You can also use the service discovery dashboards to monitor the services.

Service discovery helps you determine the kind of services running on each VM in your environment. You can find out which VM is a part of a service, the impact of shutting down or moving a VM, the impact of an incident, and the right escalation path for a problem. You can also determine which VMs are used to migrate a service and which services are impacted by a planned outage on a VM or an infrastructure component.

## Application Discovery

Applications can also be discovered on a VM with services connected to each other and those that talk to each other. You can discover predefined and custom applications. VMware Aria Operations for Logs and are the predefined applications. You can view the VMs connected to their services, how the services talk to each other, and those that are connected to different VMs.

To discover and monitor services and add and view applications, follow these steps in VMware Aria Operations:

- Configure Service and Application Discovery. For more information, see Configure Service and Application Discovery.
- Manage Services and View Applications. For more information, see Manage Services and View Applications.
- Monitor services using dashboards. For more information, see Service Discovery Dashboards.
- View the services discovered. For more information, see Discovered Services.

## Supported Platforms and Products for Service Discovery

Service discovery supports specific platforms and product versions.

You can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

### Supported Product Versions for Credential-Based Service Discovery

- For ESXi, vCenter Server, and VMware Cloud on AWS versions, see the VMware Product Interoperability Matrix.
- VMware Tools: For details, see KB 75122.

### Supported Product Versions and Other Pre-Requisites for Credential-Less Service Discovery

For information, see KB 78216.

### Operating System Versions

| Operating Systems | Version |
|---|---|
| Windows | Windows 7, Windows Server 2008/R2, and above. |
| Linux | Photon, RHEL, CentOS, SUSE Linux Enterprise Server, OEL, and Ubuntu (all Linux operating systems must be based on kernel version 2.6.25 or above). |

## Supported Services

Service and application discovery supports several services that are supported in VMware Aria Operations. The supported services are listed here.

### Supported Services:

- Active Directory
- Apache HTTP
- Apache Tomcat
- Cassandra
- DB2
- Exchange Client Access Server
- Exchange Edge Transport Server
- Exchange Hub Transport Server
- Exchange Mailbox Server
- Exchange Server
- Exchange Unified Messaging Server
- GemFire
- IIS
- JBoss
- MS-SQL DB
- MySQL DB
- Nginx
- Oracle DB
- RabbitMQ
- SharePoint
- SharePoint Application Server

- SharePoint Server
- SharePoint Web Server
- SRM vCenter Replication Management Server
- SRM vCenter Replication Server
- Sybase DB
- tc Server
- vCenter Site Recovery Manager Server
- vCloud Director
- VMware vCenter
- VMware vCenter (Appliance)
- VMware View Server
- vROps Analytics
- vROps Collector
- vROps GemFire
- vROps Postgres Data
- vROps Postgres Repl
- vROps UI
- vRLI Daemon
- vRLI vInternalization
- vRLI UI
- WebLogic
- WebSphere

## Configure Service and Application Discovery

To discover applications and services and their relationships and to access basic monitoring, you can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

### Prerequisites

- You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to discover services.

For credential-based service discovery, the configured vCenter Server user must have the following privileges:

- key: VirtualMachine.GuestOperations.ModifyAliases, Localization: Guest operations -> Guest operation alias modification
- key: VirtualMachine.GuestOperations.QueryAliases, Localization: Guest operations -> Guest operation alias query
- key: VirtualMachine.GuestOperations.Modify, Localization: Guest operations -> Guest operation modifications
- key: VirtualMachine.GuestOperations.Execute, Localization: Guest operations -> Guest operation program execution
- key: VirtualMachine.GuestOperations.Query, Localization: Guest operations -> Guest operation queries

For credential-less service discovery, the configured vCenter Server user must have the following privileges:

- key: VirtualMachine.Namespace.Management, Localization: Service Configuration -> Manage service configurations
- key: VirtualMachine.Namespace.ModifyContent, Localization: Service Configuration -> Modify service configuration
- key: VirtualMachine.Namespace.Query, Localization: Service Configuration -> Query service configurations
- key: VirtualMachine.Namespace.ReadContent, Localization: Service Configuration -> Read service configuration

- The ESXi instance that hosts the VMs where services should be discovered, must have HTTPS access to port 443 from the cloud proxy on which the service discovery adapter instance is configured.
- Verify that the following types of commands and utilities are used:

| Type | Commands and Utilities |
| --- | --- |
| UNIX Operating Systems | |
| Service Discovery | `ps`, `ss`, and `top` |
| Performance Metrics Collection | : `awk`, `csh`, `ps`, `pgrep`, and `procfs` (file system) |
| Windows Operating Systems | |

| Type | Commands and Utilities |
| --- | --- |
| Service Discovery | `wmic`, `netstat`, `findstr`, `net`, `reg`, and `sort`<br>**Note** PowerShell Scripts are now used instead of batch scripts in case of credential-less discovery starting with VMware Tools version 12.3.0 or above. See Deprecated Features for Windows Client for more details.<br>**Note** If you are using credential-less discovery and VMware Tools version 12.3.0 and above on a Windows VM, ensure that the following modules are installed on the OS::<br>■ Microsoft.PowerShell.Core (default)<br>■ Microsoft.PowerShell.Utility<br>■ Microsoft.PowerShell.Management<br>■ SmbShare (for net share)<br>■ netstat utility<br>Also, ensure that the version of the OS is Windows Server 2008R2 SP2 and above. |
| Performance Metrics Collection | `wmic`, `typeperf`, and `tasklist` |

- User Access Restrictions
  - For Linux operating systems, ensure that the user is a root or member of the *sudo* users group.

    **Note** For non-root users, the `NOPASSWD` option must be activated in `/etc/sudoers` file to avoid the metrics collector scripts from waiting for the interactive password input.

    Steps to activate the `NOPASSWD` option for a particular sudo user:

    1 Login to the specific VM as a root user.

    2 Run the `sudo visudo` command that opens an editor.

    3 In the command section, add username ALL=(ALL) NOPASSWD:<ss path>, <awk path>, <netstat path>. The username must be replaced with an existing user name for which this option is activated. Example: vmware ALL=(ALL) NOPASSWD: `/usr/sbin/ss`, `/usr/bin/netstat`, `/user/bin/awk`.

    When you perform the Execute Script action and you need to use `command/utilities`, for those commands that need a sudo user password provision, the full path of `command/utility` must be added to the NOPASSWD commands list.

    4 Save the file and close it. It is automatically reloaded.

- To discover services on Windows, the local administrator account must be configured.

  **Note** Services will not be discovered for administrator group members that are different from the administrator account itself if the policy setting `User Account Control: Run all administrators in Admin Approval Mode` is turned on. As a workaround, you can turn off this policy setting to discover services. However, if you turn the policy setting off, the security of the operating system is reduced.

- To discover services on Windows Active Directory, the domain administrator account must be configured.

- The system clock must be synchronized between the VMware Aria Operations nodes, the vCenter Server, and the VM if service discovery is working in credential-based mode and guest alias mapping is used for authentication.
- The configured user must have read and write privileges to the temp directory (execute privilege is also required on this directory in Linux systems). For Windows systems, the path can be taken from the environment variable *TEMP*. For Linux systems, it is `/tmp` and/or `/var/tmp`.
- The SSO Server URL must be reachable from the VMware Aria Operations node on which the service discovery adapter is located.
- For more information about supported platforms and versions, see Supported Platforms and Products for Service Discovery.

**Note** If more than one VMware Aria Operations instance is monitoring the same vCenter Server and service discovery is activated for those VMware Aria Operations instances, then service discovery might be unstable, which is a known VMware Tools problem. As a result, guest operations might fail to execute.

Procedure

1. From the left menu, click **Configure > Applications and Services**, and then from the right panel, click the **Service Discovery** tile.
2. From the **Service Discovery** page, click the **Configure Service Discovery** option.
3. From the **Integrations** page, click the vCenter Server instance from the list and then select the **Service Discovery** tab.
4. To activate service discovery in this vCenter Server, activate the **Service Discovery** option.
5. To activate application discovery in this vCenter Server, select the **Enable Application Discovery** check box.
6. You can choose to add credentials by selecting the **Use alternate credentials** check box.
   a. Click the plus sign and enter the details in the **Manage Credentials** dialog box, which include a credential name and a vCenter user name and password. In addition, enter the user name and password for Windows, Linux, and SRM and click **OK**.

7 Alternatively, if you are using the default user name and password, enter a default user name and password for Windows, Linux, and SRM.

8 Enter a password for the guest user mapping.

9 Click **Save**.

**Note** If you specify a non-root user for Linux, services are not discovered unless you activate the option Use Sudo (Linux Non-root user) while editing the associated Service Discovery adapter instance after you create the vCenter Cloud Account. This option is deactivated by default, which means the root user is expected by default when you configure the vCenter Cloud Account.

10 Edit the cloud account created for service discovery.

11 In the **Advanced Settings** section, activate the **Application Discovery** field to discover predefined and custom applications.

12 In the **Advanced Settings** section, to configure credential-less service discovery, select **Activated** from the **Credential-less service discovery status** field.

13 Click **Save**.

#### What to do next

You can manage services supported by VMware Aria Operations on specific VMs.

## Manage Services

You can manage services supported by VMware Aria Operations on the specific VMs.

### Where You Manage Services

From the left menu, select **Environment > Applications**. From the **Applications** panel, select **Manage SDMP Services**. You can also navigate to the `Manage SDMP Services` tab by selecting **Configure > Applications and Services**, and then from the right panel, select the **Service Discovery** tile. Select the `Manage Services` option from the **Service Discovery** page.

You can view specific details from the options in the data grid.

Table 4-60. Datagrid Options

| Options | Description |
| --- | --- |
| VM Name | Name of the VM. |
| Operating System | Operating system installed on the VM. |
| Services Discovered | Displays the names of discovered services or `None`, if services are not discovered on the VM. |
| Service Monitoring | Displays the current value of the VM's service monitoring setting. If set, services are discovered and service performance metrics are calculated every 5 minutes. Otherwise, service discovery is performed every 24 hours. |

### Table 4-60. Datagrid Options (continued)

| Options | Description |
| --- | --- |
| Authentications Status | VM authentication status for service discovery. The possible values are:<br>■ Unknown<br>■ Failed<br>■ Guest Alias<br>■ Common Credentials<br>■ Credential-less |
| Power State | Power status of the VMs. The possible values are:<br>■ Powered On<br>■ Powered Off<br>■ Suspended<br>■ Unknown |
| Collection State | Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the collection state icon. To manage an adapter instance to start and stop collection of data, from the left menu, click **Environment > Inventory**. |
| Collection Status | Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the collection status icon. To manage an adapter instance to start and stop collection of data, from the left menu, click **Environment > Inventory**.<br>You can view a message for VMs with a failed authentication status in a tool tip when you point to the collection status icon. |
| vCenter Name | Name of the vCenter Adapter instance to which that VM resource belongs. |

### Table 4-61. Toolbar and Row Options

| Options | Description |
| --- | --- |
| VM Actions | Displays a list of actions. For more information, see Actions in VMware Aria Operations. |
| **Horizontal Ellipsis > Provide Password** | Select VMs from the list, click the horizontal ellipsis and then click **Provide Password** to provide a user name and password for the selected VMs to discover the services. |
| **Horizontal Ellipsis > Activate Service Monitoring** | Select VMs from the list, click the horizontal ellipsis and then click **Activate Service Monitoring** to activate frequent service discovery and service performance metrics calculation (every 5 minutes).<br>**Note** Selecting too many VMs will potentially result in vCenter Server degradation which is a known issue. |

Table 4-61. Toolbar and Row Options (continued)

| Options | Description |
| --- | --- |
| **Horizontal Ellipsis > Deactivate Service Monitoring** | Select VMs from the list, click the horizontal ellipsis and then click **Deactivate Service Monitoring** to deactivate frequent service discovery and service performance metrics calculation. Service discovery defaults to the 24-hour cycle. |
| Check box to Select/Deselect all | Selects/clears all VM object selections. |
| **Vertical Ellipsis > Go To Details** | Navigates to the **Summary** tab for the selected VM. |
| Filter | You can search through the list of VMs according to the following criteria: VM Name, Operating System, Power State, Authentication Status, and Service. |

## Create Application Definition

Create application definitions, when you want to discover applications based on the user defined rules such as object names, tags, or properties.

### Procedure

1 From the left menu, click **Configure > Application and Services**.

2 In the **Applications and Services** page, click **Application Discovery: Rule Based** tile.

3 Click **Add Application Definition**.

4 In the **Create Application Definition** page, enter the required values.

Table 4-62. Create Application Definition

| Option | Description |
| --- | --- |
| Application Name | Enter a name for the application. |
| Application Prefix | Enter a prefix for the application. Prefix can help you identify the discovered applications by name. It will be prepended to the application name. |
| Select criteria for grouping discovered applications - This section allows you to define the criteria based on which the objects are grouped into an application. | |
| Group For Object Type | Select an object type to group objects by |
| Select | Select the criteria to group the objects by. The options available are:<br>■ Properties<br>■ Object Name<br>■ Tag Category |
| Add | Click add to add another criteria |
| Reset | Click reset to clear the criteria |

Table 4-62. Create Application Definition (continued)

| Option | Description |
| --- | --- |
| Advanced Settings - This section allows you to specify information for additional filters. | |
| Select the Object Type that matches all of the following criteria | Use the drop-down list to select the object type which matches the defined criteria |
| Select | Select the object type. The options available are:<br>■ Metrics<br>■ Relationships<br>■ Properties<br>■ Object Name<br>■ Tag<br>Based on your selection, you must specify the additional attributes. |
| Add | Click Add to add another criteria |
| Reset | Click Reset to clear the defined criteria |
| Add New Criteria | Click Add New Criteria to add another criteria |
| Remove Criteria | Click Remove Criteria to delete the criteria |

5 Click **Create**.

The created application definition is displayed in the **Application Discovery: Rule Based** page.

# Manage Applications

You can discover applications based on user defined criteria such as virtual machine names, tags, or properties (rule-based applications). For example, if the application definition is assigned to a cluster compute resource called 'xyz', then the applications would be discovered only within that cluster. The object scope assignments can be made to both service based application definitions and rule based application definitions.

## Where You Manage Applications

From the left menu, select **Environment > Applications**. From the **Applications** panel, select **Manage Applications**.

You can view direct or inherited assignments by selecting either **All** or **Self** respectively, in the sliding option at the top of the UI. You must select any of the environment objects in order to see application definitions.

You can view specific details from the options in the data grid.

### Table 4-63. Datagrid Options

| Options | Description |
| --- | --- |
| Object Name | Name of the Object. |
| Direct Assignments | Displays the number of application definitions assigned directly to this object. |
| Inherited Assignments | Displays the number of application definitions that are inherited from the object's ancestors. |
| Object Type | Displays the object type. |

### Table 4-64. Assign Application Definition To Scope of Objects

| Options | Description |
| --- | --- |
| **Vertical Ellipsis** | Displays Assign Applications Definitions page. |
| Select rules from the list below to activate application discovery for selected objects | Displays the list of Application Definition Name and Application Definition Type. Select the required definition, the options available are Service-based application definition and Rule-based application definition.<br>**Note** You can use the search option to look for application definitions name. |
| > | Expand the arrow next to the object to view the definitions assigned to the selected object and the discovered applications count. You can also view the inherited assignments for the selected object. |
| Save | Navigates to the Summary tab for the selected Object. |
| **Vertical Ellipsis > Go To Details** | Navigates to the **Summary** tab for the selected objects. |

## View Applications

You can view applications created by VMware Aria Operations on specific VMs.

You can view all resources of the Application type, including:

- VMs connected to their services, how the services talk to each other and are connected to different VMs, using the Service Discovery adapter.
- Discovered applications using the Application Management Pack.

### Where You View the Applications

From the left menu, select **Environment > Applications** to view the **Applications Home** page.

You can view the list of applications from **Applications Home** page. Click an application to view the application in the right pane. Select an application row and then click the **Vertical Ellipsis > Go To Details** to view the object details. You can use the Filter option to search through the list of applications according to the following criteria :

- Name
- Object Type

### Options to View the Applications

After you click on an application from the **Applications Home** page, you can view details of the application or service in the right pane.

| Options | Description |
|---|---|
| View Sphere | Displays the VMs connected to their services in a spherical view. |
| View Graph | Displays the VMs connected to their services as a graph. |
| View List | Displays the VMs connected to their services in a list view. |
| View Links | You can view the links between various services. |

# Discovered Services

You can discover services using the Service Discovery adapter.

## Discovered Services

You can view discovered services, the number of VMs on which each discovered service is running, and you can configure service discovery.

### Where You View the Discovered Services

From the left menu, click **Configure > Applications and Services**. From the right panel, click the **Service Discovery** tile.

### Discovered Services

You see a list of services that are discovered and the number of VMs that have the services running. You see this section after you have configured Service Discovery and the services are discovered.

### Known Services

You see a list of all the services supported and those that can be discovered.

### Custom Services

You can add a service by clicking **Add Service Definition**. You can add either a process name or Regex from the **Add Custom Service** dialog box.

Table 4-65. Add Custom Service

| Options | Description |
| --- | --- |
| Type | Specify the type as either process or regex.<br>**Process**:<br>The process name must exactly match the name that you see in the guest OS when running commands `ps` in Linux and `wmic` in Windows. Specify a single port for each service.<br>The following characters are not supported: **`,`**, **`\`**, and **`#`**.<br>**Regex**<br>Enter a regular expression that corresponds to the command line (or at least name) of the service, that you see in the guest OS when you run the following commands: `ps` in Linux and `wmic` in Windows.<br>For example, to discover Cassandra services, enter **`cass.*dra`** as the regex.<br>The following characters are not supported: **`,`** and **`\n`**.<br>**Note** If you specify a service using regex and the service reboots at a later time and all the ports are modified, the current service displays as offline and the new service is discovered again. If the new service has any port overlap with the previous one, a new resource will not be discovered and the previous one will continue collecting data as before. |
| Process Name | Enter a process name. |
| Port | Enter the port information. |
| Display Name | Enter the display name. |

**Note** Custom service can be discovered via Service Discovery if there is a permanent listening TCP port or if there is an established UDP connection.

# Discovered Applications

You can discover applications based on grouping the services discovered, using the Service Discovery adapter.

## Discovered Applications

You can discover predefined applications and custom applications. VMware Aria Operations for Logs and are the predefined applications.

### Where You View the Discovered Applications

From the left menu, click **Configure > Applications and Services**, and then from the right panel click the **Application Discovery: Service Based** tile.

### Custom Applications

You can define custom applications. Click **Add Application Definition** to add a custom application.

Table 4-66. Add Custom Application

| Option | Description |
| --- | --- |
| Name | Enter a name for the application. |
| Prefix | Enter a prefix for the application. |
| Services | Select a service from the list.<br>If the Service Discovery adapter discovers the service and if the services are connected to each other, a new application is discovered. The new application appears in the **Applications Home** page. Navigate to **Environment > Applications**. |

### Known Applications

You see a list of the predefined applications supported. Select the predefined application, click the vertical ellipsis, and then click **Preview**. From the relevant Application pane, you can view the application services that can be discovered, and if the connections are identified, they will form an application. You can view the applications that are discovered by clicking **Environment > Applications > Applications Home**.

**Note** If the same service instance is a member of more than one application (both known and custom), then the service is a part of the application that has the greatest number of services configured in the application definition.

The discovered set of services that communicate with each other, should match at least 70% of the defined application. Only the matching ones are filtered out based on whether one service is defined on more than one application.

## Discovered Rule-Based Applications

VMware Aria Operations discovers applications based on user defined criteria such as object names, tags, or properties.

### Where You View the Application Discovery Definitions

From the left menu, click **Configure > Application and Services**. From the right pane, click the **Application Discovery: Rule Based** tile.

### Application Discovery: Rule Based

The rule based application discovery provides you the following options.

- Specify grouping criteria for each rule, it can be tags, object names, or properties
- Apply additional filters to the rule to limit the scope of objects considered for grouping

You can discover applications based on the user defined criteria such as object names, tags, or properties. After the rules are created you can assign these rules to specific objects. VMware Aria Operations discovers applications based on these rules and displays them in Applications Home page.

### Edit or Delete Application Rule

You can modify or delete the application definition. To edit or delete the application definition, click the vertical ellipsis next to the application and select **Edit** or **Delete**. If you select the edit option you are redirected to the **Edit Application Definition** page, here you can modify the application definition and save the changes. If you select the delete option, you are prompted with a confirmation message, click **Yes** to delete the application definition. Please note that if you delete the application definition, the applications that were discovered by that definition are also deleted.

## Alert for Service Unavailability

When a service is unavailable, an alert is triggered for the specific VM.

### Alert for Service Unavailability

On a VM that is monitored, if one of the services is down, in the next collection cycle an alert is triggered.

| Alert Name | Symptom |
|---|---|
| One or more monitored service(s) are unavailable on the virtual machine. | Service is not available.<br>When the service is available again, the symptom disappears. |

The alert is canceled in the following scenarios:

- When all the discovered services are available again in the monitored VM.
- If the service is not available within 7 days.
- If you deactivate service monitoring for the monitored VM.

**Where You Find the Alert**

From the **Manage Services** page, ensure that the VM is monitored and one or more service(s) are unavailable on the VM. Select the VM, click **Show Details** to go to the summary page. Click **Alerts** from the toolbar, and then click the **Alerts** tab.

### Property for Service Unavailability

You can view the property called Status for a service that has been discovered on the VM. For more information, see the topic called Services Properties.

## Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

## Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

### Table 4-67. Virtual Machine Metrics

| Metric Name | Description |
| --- | --- |
| Guest OS Services\|Total Number of Services | Number of out-of-the-box and user-defined services discovered in the VM. |
| Guest OS Services\|Number of User Defined Services | Number of user-defined services discovered in the VM. |
| Guest OS Services\|Number of OOTB Services | Number of out-of-the-box services discovered in the VM. |
| Guest OS Services\|Number of Outgoing Connections | Number of outgoing connection counts from the discovered services. |
| Guest OS Services\|Number of Incoming Connections | Number of incoming connection counts to the discovered services. |

## Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

### Table 4-68. Service Summary Metrics

| Metric Name | Description |
| --- | --- |
| Summary\|Incoming Connections Count | Number of incoming connections. |
| Summary\|Outgoing Connections Count | Number of outgoing connections. |
| Summary\|Connections Count | Number of incoming and outgoing connections. |
| Summary\|Pid | Process ID. |

## Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

### Table 4-69. Service Performance Metrics

| Metric Name | Description |
| --- | --- |
| Performance metrics group\|CPU | CPU usage in percentage. |
| Performance metrics group\|Memory | Memory usage in KB. |
| Performance metrics group\|IO Read Throughput | IO read throughput in KBps. |
| Performance metrics group\|IO Write Throughput | IO write throughput in KBps. |

### Service Type Metrics

Service discovery discovers metrics for service type objects.

**Table 4-70. Service Type Metrics**

| Metric Name | Description |
| --- | --- |
| Number of instances | Number of instances of this service type. |

# Configuring Business Applications

A Business Application is a set of interconnected applications, services and hosts, which are configured to offer a service to the organisation. Business Applications can be internal, like organization email system or customer-facing, like an organization website. The Business Applications page in VMware Aria Operations is where you see Business Applications and their health.

A Business Application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. You create a Business Application to determine how your environment is affected when one or more components in the Business Application experiences problems, and to monitor the overall health and performance of the Business Application.

VMware Aria Operations can also display applications discovered through Service Discovery, and by importing applications from VMware Aria Operations for Networks SaaS and Ensemble.

## The Business Application Page

The Business Applications page is where you view all your Business Applications and their health in a sortable table. The health of the Business Application is determined by the aggregated health of the underlying objects. The **Health** column in the table displays the status.

### Where You View the Business Applications Page

From the left menu, select **Environment > Business Applications** to view the Business Applications page. All the available Business Applications and applications are displayed in the table.

The table displays the following types of objects:

Table 4-71. Business Applications Page

| Business Applications Source | Business Applications Type |
| --- | --- |
| Auto discovered and monitored applications which have more than one node | ■ Imported through VMware Aria Operations for Networks SaaS. |
| Imported from VMware Aria Hub. | |
| Manually created Business Applications. | ■ Can contain custom groups.<br>■ Can contain any object to tier relationships.<br>■ Can contain applications. |

**Note** Applications discovered using Service Discovery which have only one node are not displayed.

## Columns in the Business Application Table

The Business Applications table displays the following columns:

| Column Name | Description |
| --- | --- |
| Vertical Ellipses | Offers options for the user to perform actions on the Business Application. You can also Enable Synthetic Monitoring, Disable Synthetic Monitoring, and Configure Synthetic Monitoring for the Business Application. |
| Health | Displays the health of the Business Application by showing different icons. On hovering the cursor over the icon, you can see the health percentage. |
| Name | Displays the name of the Business Application and an icon indicates the type of Business Application. |
| Description | Displays the description of the Business Application you entered when creating it. |
| Synthetic Monitoring State | Displays one of the following states for Synthetic Monitoring:<br>■ Enabled<br>■ Disabled<br>■ Configured<br>■ Not configured |
| Application Tags | Displays the application tags you entered when you created the Business Application. |
| Business Criticality | Displays one of the following criticality states for the Business Application:<br>■ Medium<br>■ Critical<br>■ Low |
| Environment | Displays the environment where the Business Application resides in. |

## Actionable Options on the Business Applications Page

### Table 4-72. Actionable Options on the Business Applications Page

| Option | Description |
| --- | --- |
| Vertical Ellipses | Click the vertical ellipses beside a Business Application to perform one of the following actions:<br>■ Edit (Not available for VMware Aria Operationscreated Business Applications)<br>■ Delete<br>■ Go to Details. This opens the Business Applications Summary tab in the Object Browser.<br>■ Troubleshoot with Workbench<br>■ Configure Synthetic Monitoring<br>■ Enable Synthetic Monitoring<br>■ Disable Synthetic Monitoring<br>■ Edit Synthetic Monitoring |
| Filter | Filter the objects in the table by name. |
| Add Business Application | Click to create a new Business Application |
| Show Columns | Click to hide/show columns in the table. |
| Preview | Click on a Business Application name in the table to see a visual preview of the objects and their relationships. |

# Add Business Application

A Business Application is a first class citizen with tiers. Business Applications are groups of related objects in your environment that mimic an application in your business. You can create application to application relationships. Once you add a Business Application to VMware Aria Operations, you can use the Business Applications page to track the health of objects in the application. You can optimize capacity, cost, and run What-if Analysis on Business Applications.

## How Business Applications Work

The tier is a convenient way to organize objects that perform a specific task in an application in your organization. For example, you can group all of your database servers together in a tier. Business Applications need not contain tier objects. Application or custom group objects can be connected to a Business Application without a tier object, and a tier object can contain no objects.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the Business Application. However, if the Business Application is created by VMware Aria Operations, the changes in the actual application will be reflected in the Business Application object without manual edit.

Construct a Business Application to view a particular segment of your business. The Business Application shows how the performance of one object affects other objects in the same Application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data for your business, you see a yellow, orange, or red status if the application health is degrading. From the Business Applications page, you can investigate which server is causing or exhibiting the problem.

## The Add Business Application Screen

When you click **Add**, VMware Aria Operations displays a blank canvas on the left pane, and options to select objects on the right pane. The **Select Members** section in the right pane has two tabs. The **Inventory** tab displays a list of all objects in you environment in a tree structure. You can search for objects here. The **Custom Groups** tab displays all the custom groups in your environment. You can use the search bar to search for custom groups.

Drag and drop objects from the right pane to the tier objects on the left pane. You can also drag and drop objects directly to the Business Application. Click **Save** to create the Business Application.

The new Business Application canvas is what you use to build your Business Application.

### Table 4-73. New Business Application Canvas

| Option | Description |
| --- | --- |
| Vertical icon | Change the layout of the canvas from horizontal to vertical. |
| Undo icon | Undo the last action. |
| Redo icon | Redo the last action. |
| Fit View icon | Fit the Business Application structure within the boundary of the canvas. Use this option after you zoom into the canvas. |
| Add Tier icon | Add a new tier to the Business Application |
| Edit Tier icon | Available when you click a tier. Click the edit icon in a tier to edit the name and description of the tier. |

### Table 4-73. New Business Application Canvas (continued)

| Option | Description |
| --- | --- |
| Edit Business Application Details icon | Available when you click a Business Application. Click the edit icon in the Business Application to open the Business Applications Detail dialog box. Alternatively, click the DETAILS link next to the name of the Business Application. |
| Preview | Available for custom group and application objects when you have added objects to a tier. Click the preview icon after you have added an Application to a tier to see how services and objects within the application are related. This interactive element is displayed on the right pane. You can view links, switch to spherical view or list view. If you have too many objects to can use the filter option to find objects. |

The **Select Members** pane is where you select objects to add to the Business Application. Drag an object to add to a tier in the canvas. Applications and custom groups can be nested under a Business Application. Custom groups can be added to applications and vice versa.

To find an object, search by name. Each object listed includes identifier information to help distinguish between objects of similar names.

### Table 4-74. Select Members Pane

| Option | Description |
| --- | --- |
| Search | Available in the **Inventory** and **Custom Groups** tabs. Search for objects to add to the Business Application. Filter your search by clicking the filter icon. In the **Inventory** tab, by using the filter option, you can choose to see only objects which are Applications. |
| Inventory | Browse the inventory to select objects to add to the Business Application using drag and drop. |
| Custom Groups | Browse custom groups to add to the Business Application using drag and drop. You can nest custom groups or add them to the Business Application directly. You can add custom groups to Applications. |

The Business Applications Details dialog box is where you add information to build out a model of your Business Application resources through applications and other infrastructure objects in your environment.

### Table 4-75. Business Application Details

| Option | Description |
| --- | --- |
| Description | Provide a description for your Business Application. |
| Application Tag | Provide a tag for your Business Application. |

Table 4-75. Business Application Details (continued)

| Option | Description |
| --- | --- |
| Business Criticality | This is your Business Application business criticality. Select between Medium, Critical and Low. Default is Medium. |
| Environment | This is where your Business Application is deployed. Select from one of the following options:<br>■ DR<br>■ Development<br>■ Production<br>■ Staging<br>■ Test |

## Configure, Activate, or Deactivate Synthetic Monitoring

Synthetic monitoring can report API monitoring and check results in Prometheus metrics format which is pushed to VMware Aria Operations for further analysis, visualization, and alerting purposes.

### Where you find Synthetic Monitoring

From the left menu, click **Environment > Business Application**. On the Business Application page, click the vertical ellipses next to a Business Application name to configure, activate, or deactivate Synthetic Monitoring. Deatcivate Synthetic Monitoring appears only for a Business Application which has Synthetic Monitoring already activated. Activate Synthetic Monitoring appears only for a Business Application which has Synthetic Monitoring already configured.

### Synthetic Monitoring Configuration Page

You can add up to 10 APIs for monitoring. Click the **ADD NEW API** button to add new APIs. The API configuration options appears in the bottom pane.

### ADD NEW API Options

When you click the **ADD NEW API** button, you get the following configuration options:

| Option | Description |
| --- | --- |
| API Name | Provide a name for the API call. |
| API Request | Select the **HTTP Method** and provide the **Public Endpoint URL**. Options for the API type are:<br>■ GET<br>■ POST |
| Advanced Settings | Set number of API call retry attempts and timeout value. |

| Option | Description |
| --- | --- |
| Headers | In the **Headers** tab, select the content type of the body from the drop-down menu. The options are:<br>■ application/json<br>■ application/xml<br>Provide inputs for the **Key** and **Value** and click **TEST**. |
| Body | In the **Body** tab, provide the body for **API Request**. |
| Parameters | Enter **KEY** and **VALUE** properties |
| Authentication | In the **Authentication** tab, select a type of authentication for the API from the drop-down menu. The options are:<br>■ None<br>■ Basic Authentication<br>■ Bearer Token<br>■ OAuth Authentication<br>When you select a type of authentication apart from None, you get the options to provide credentials. The credentials depend on the authentication type.<br>Once you provide the authentication type, click **TEST** to check if VMware Aria Operations can reach the API with the values you provided.<br>**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The VMware Aria Operations administrator must assign the credentials from the **Orphan and Unassigned** page. For more information, see the topic, Managing Orphaned and Unassigned Content. |

Click **TEST** to test the API request. Click **Save** after you provide these values..

VMware Aria Operations monitors the APIs at a frequency displayed in the **Frequency** property in the top pane after you activate Synthetic Monitoring for the Business Application. Each Business Application can have up to 10 API requests registered. VMware Aria Operations can have up to 20 Business Applications with Synthetic Monitoring activated.

You can view the results of synthetic monitoring in the Synthetic Monitoring tab. See, the topic, Synthetic Monitoring Tab.

# Application Integration

You can integrate Application Performance Monitoring tools to discover applications in VMware Aria Operations.

The integrations for Application Performance Monitoring tools are enabled using the following applications.

- AppDynamics. For product documentation, see VMware Aria Operations Management Pack for AppDynamics.

- New Relic. For product documentation, see VMware Aria Operations Management Pack for New Relic.
- DataDog. For product documentation, see VMware Aria Operations Management Pack for Datadog.
- Dynatrace. For product documentation, see VMware Aria Operations Management Pack for Dynatrace.

## Application Discovery

Integrating the Application Performance Monitoring tools to discover applications, enhances the troubleshooting abilities by retrieving the application topology and key metrics.

### How Can I View Applications Available for Integration

From the left menu, click **Environment** > **Applications** > **Application Integration**. Click **Discover Applications** to view the applications available for integration and the applications that are successfully integrated. For details, see Integrating Applications.

### Where Can I Find the Discovered Applications

Once the application is integrated, the discovered applications are displayed in the **Data Sources** > **Integrations** page after a few collection cycles.

Table 4-76. Application Integration Page Options

| Options | Descriptions |
| --- | --- |
| All Filters | You can filter the discovered applications by Name, Object Type, and Adapter Type. |
| Application Name | Displays the name of the application. Click this link to view the details in the **Summary** page. |
| Object Type | Displays the type of object. |
| Source | Displays the source of the application. |
| Adapter Type | Displays the name of the adapter. |
| Active Alerts | Displays the active alerts associated with the application. |

## Integrating Applications

You can integrate Application Performance Monitoring tools to discover applications in VMware Aria Operations. These integrations enhance the troubleshooting abilities by retrieving the application topology and key metrics.

The **Application Integration** page displays the applications available for integration and the applications that are successfully integrated. The integrations for Application Performance Monitoring tools are enabled using the following applications.

- AppDynamics

- New Relic
- DataDog
- Dynatrace

### How Can I Download Applications for Integration

From the left menu, click **Environment** > **Applications** > **Application Integration**.

Click **Configure Now** and in the **Integrations** > **Account Types** page, click the required management pack to automatically download it. For details, see The Integrations Page in VMware Aria Operations.

You can also download the management pack from the **Repository** tab under **Data Sources** > **Integrations**. For details, see The Repository Tab.

### How Can I Add Accounts for Solutions

From the left menu, click **Data Sources** > **Integrations** > **Accounts**. Click **Add Account**, and then select the solution you want to manage.

You can also add accounts from the **Repository** tab under **Data Sources** > **Integrations**. For details, see The Repository Tab.

## Configuring Ping Adapter Instances

In VMware Aria Operations, you can configure the Ping functionality to verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN.

- If you have multiple adapter instances running on different collectors and both are pinging the same address, you can still get statistics from both the adapter instances for the same IP.
- The FQDN names are checked for validity, the FQDN validation relies on RFC1034 and RFC1123, and only top level domains of the internet are validated. The `.local` domain is not supported as it does not fall into the list of top-level domains in the Domain Name System (DNS) of the Internet.

#### Procedure

1. In the left menu, click **Data Sources** > **Integrations**.
2. In the Accounts tab, click **Add Account**
3. Click **Other** to filter the list of accounts. The Ping account tile is displayed after filtering out other tiles.
4. Click the Ping adapter instance.
5. Click **Yes** in the diaglog box in the dialog box which opens. This will install the management pack.

6 Configure the Ping adapter instance.

| Option | Description |
| --- | --- |
| **Name** | Enter a name for the adapter instance. |
| **Description** | Enter the description of the adapter instance. |
| **Unique Name** | Specify the name for the adapter instance. You can use the name to view the metrics published for the adapter instance. |
| **Address List** | Specify the IP address, IP address range, and the FQDN which must be pinged. |
| **Configuration Filename** | Specify the name of the configuration file. The configuration file contains the IP addresses, CIDR information, and FQDN details as a comma-separated file. |
| **Collectors/Groups** | Select the collector from which this adapter instance must run. |
| **Validate Connection** | Click to check whether the connection is successful or not. |
| **Advanced Settings** | To configure the advanced settings, click the drop-down menu. |
| **Wait Interval Time (second)** | Specify the time interval in seconds to wait before running the next batch. Range: 0-300 seconds. |
| **Batch Size** | Specify the number of request packets to send to each target. Range: 20-100. |
| **Interval (millisecond)** | Specify the time the fping waits between successive packets to an individual targets. Greater or equal to 2000 milli seconds. |
| **DNS Name Resolve Interval** | Specify the time at which you must resolve the DNS name for the next cycle. Minimum value is 15 minutes. |
| **Packet Size** | Specify the byte size of the packet when you ping. Range: 56-65536 bytes. |
| **Don't Fragment** | Select **False** to fragment the packet and **True** to not fragment the packet. |
| **Generate FQDN Child IPs** | Select **True** to create IP objects by resolved names and add as child of FQDN. |

7 Click **ADD**.

Results

After you configure the Ping adapter instance, you can view the adapter details from **Data Sources > Integrations > Repository**.

# VMware Aria Automation

VMware Aria Automation extends operational management capabilities of the VMware Aria Operations platform to provide the cloud aware operational visibility of the cloud infrastructure. VMware Aria Automation helps you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use VMware Aria Automation to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with VMware Aria Operations.
- Import and synchronize existing cloud accounts from VMware Aria Automation to VMware Aria Operations.
- Manage the workload placement of VMs that are part of the clusters managed by VMware Aria Automation.
- Integrate and troubleshoot vSphere endpoint issues associated with VMware Aria Automation using the VMware Aria Operations dashboard.

**Note** In this release we support only vSphere endpoints.

## Advanced Workload Placement for Allocation Model

VMware Aria Operations supports allocation aware advanced workload placement between VMware Aria Automation and VMware Aria Operations for initial placement of virtual machines. The placement recommendation is based on allocation settings that are defined in the VMware Aria Operations policy. The advanced workload placement awareness for allocation works with the existing demand model.

### Need for Allocation Model Awareness

The advanced workload placement based on demand model, depends on the actual demand for resources in a cluster and datastore. If the allocation is based only on the utilization of resources, then it might result in over allocation or over provisioning of resources in clusters. To avoid this, VMware Aria Operations provides allocation model awareness. Allocation model awareness addresses issues related to over allocation or over provisioning by providing you the option of setting the appropriate overcommit ratios in the VMware Aria Operations policy.

The advanced workload placement feature uses the demand model by default and cannot be turned off. To activate the allocation model in addition to the demand model, in VMware Aria Operations configure the appropriate overcommit ratios in policy settings for the preferred clusters, datastores, and datastore clusters.

### Prerequisites

- VMware Aria Operations is configured as an endpoint in VMware Aria Automation,this happens automatically in case of VMware Aria Operations
- Advanced placement policy in the Cloud zone must be activated
- vCenter Cloud account instance on which the initial provision is done should be same across VMware Aria Operations and VMware Aria Automation
- Overcommit ratios in VMware Aria Operations policy must be configured for the following:
  - For CPU and Memory overcommit ratio - Cluster Compute Resource

- For Disk overcommit ratio - Datastore and Datastore Cluster

### How to Activate Advanced Placement Policy

To activate the advanced placement policy, in VMware Aria Automation Cloud Assembly Service, navigate to **Infrastructure** > **Cloud Zone** > **Summary** tab and set the placement policy to **ADVANCED**.

**Note** If VMware Aria Operations returns no recommendations, then under the placement policy, you can specify if you want VMware Aria Automation to fall back to its default placement using the toggle option.

### How to Activate Allocation Awareness in Advanced Workload Placement Feature

To activate Allocation awareness in VMware Aria Operations perform the following actions.

1. From the left menu, click **Configure** >**Policies**.
2. Select **Policy Definition**.
3. Select the Active policy which is assigned to the Cluster Compute resource under Cloud Zones and make the required changes.
4. Click **Edit Policy** and navigate to the **Capacity** tile.
5. Select the **Object Type** as Cluster Compute Resource and activate the Allocation Model.
6. Set the overcommit ratio as per your requirement and click **Save**.
7. Repeat Step 3 to Step 5 for datastore cluster, vSAN datastores and local datastores.

After the configuration is completed, VMware Aria Automation Cloud makes a provisioning request to VMware Aria Operations, and the advanced workload placement engine calculates the recommendation and shares it with VMware Aria Automation. To know the changes after the configuration, view the Allocation metrics and Demand metrics for the cluster.

**Note** The allocation model awareness enhancement is limited to advanced workload placement feature. This capability is not extended to workload optimization feature.

## Integrate VMware Aria Automation SaaS Service with VMware Aria Operations SaaS Service

The integration of the VMware Aria Operations service and VMware Aria Automation Cloud service happens automatically if your organization has access to both the services. If you add VMware Aria Automation Cloud services first and then add VMware Aria Operations cloud services, VMware Aria Operations autoconfigures the VMware Aria Automation Cloud accounts.

### How Cloud Integration Works

Once the integration of VMware Aria Automation cloud service with VMware Aria Operations is complete. VMware Aria Operations service provides the following information about the integrated VMware Aria Automation cloud account:

- Import cloud accounts defined in VMware Aria Automation Cloud to VMware Aria Operations.
- View cloud zones defined in VMware Aria Automation cloud in VMware Aria Operations.
- Modifications done to Cloud Zones in VMware Aria Automation cloud are reflected in VMware Aria Operations.
- View the integrated cloud adapter instance in VMware Aria Automation inventory objects.
- List the objects related to the integrated cloud adapter in VMware Aria Operations inventory.
- Access to the objects in your VMware Aria Automation cloud account is based on your user role. An organization administrator has access to all the objects in your cloud environment, but an organization member has limited access to the objects in your cloud environment.
- Updates done to the objects or data of VMware Aria Automation cloud are reflected in VMware Aria Operations as well.
- After the VMware Aria Automation cloud and VMware Aria Operations integration is complete, all the dashboards of VMware Aria Automation are created in VMware Aria Operations.

You can perform reset and upgrade operation on VMware Aria Automation cloud accounts. If you reset the VMware Aria Automation account, you create an environment and all the data related to the old setup is removed from the system.

If you upgrade the VMware Aria Automation cloud account, the historical data is retained in the VMware Aria Operations account.

## Importing Accounts from VMware Aria Automation

You can import and synchronize existing cloud accounts from VMware Aria Automation to VMware Aria Operations. Click **Import Accounts from VRA** > **Import Accounts** to list all the cloud accounts associated with vCenter Server, Amazon AWS, and Microsoft Azure that are not managed by VMware Aria Operations. You can select and import these accounts into VMware Aria Operations directly with existing credentials as defined in VMware Aria Automation or add or edit the credentials before the import process. The **Import Accounts from VRA** option is hidden from the user until the integration with VMware Aria Automation is activated from the integration page under **Data Sources > Integrations > Accounts** or **Repository** tabs.

#### Prerequisites

- Verify that VMware Aria Automation is activated from **Data Sources > Integrations > Accounts** in VMware Aria Operations.
- Verify that you know the vCenter Server credentials that have sufficient privileges to connect and collect data.

- Verify that the user has privileges of Organizational Owner and Cloud Assembly administrator set in VMware Aria Automation.

Procedure

1 From the left menu, go to **Data Sources > Integrations > Accounts** tab, click on the horizontal ellipses, and then select **Import Accounts from VRA**.

2 From the **Import Accounts** page, select the cloud account you want to import.

3 To override an existing credential from VMware Aria Automation.

   - Select the existing credential from the **Credential** drop-down menu and click **Save**.
   - To add a new credential, click the plus icon next to the **Credential** drop-down menu and enter the credential details and click **Save**.

4 Select the collector/group from the drop-down menu.

5 Click **Validate** to verify that the connection is successful.

6 Click **Import**.

Results

The imported cloud account is listed in the **Data Sources > Integrations > Accounts** page. After the data collection for the cloud account is complete the configuration status changes from **Warning** to **OK**.

## Supported VMware Aria Automation Versions

VMware Aria Automation is supported on vRealize Operations 8.6 version. Workload placement for day 1 operations is supported from vRealize Automation 7.3 onwards with vRealize Operations 6.6 and above. Workload placement for day 2 operations is supported from vRealize Automation 7.5 onwards with vRealize Operations 7.0 and above.

## Object Types

VMware Aria Automation brings in cloud accounts and their relationships from VMware Aria Automation into VMware Aria Operations for operational analysis. You can use the following items in the virtual infrastructure as object types in VMware Aria Operations.

- Cloud Zone
- Blueprint
- Project
- Deployment
- Cloud Account
- User
- Organization

- Cloud Automation Services World

## Workload Placement

In VMware Aria Operations, you can configure VMware Aria Automation instances to work with VMware Aria Operations instances. Using VMware Aria Operations you can monitor the placement of existing workloads and optimize the resource usage.

### Prerequisites

- Verify that the user has privileges of Organizational Owner and Cloud Assembly Administrator set in VMware Aria Automation.
- You must know the vCenter Server credentials and have the necessary permissions to connect and collect data.
- VMware Aria Operations must have the same vCenter Cloud Account configured to match with VMware Aria Automation.
- Ensure that integration is activated for VMware Aria Operations and VMware Aria Automation.

### Procedure

1. From the left menu, click **Home** >**Optimize** > **Workload Placement**.
2. Click the **View** filter drop-down menu and select the **VRA Managed** objects.

   All the Cloud Zones related to the vCenter Server are displayed in VMware Aria Operations.
3. Click the **Cloud Zone** you want to optimize.
4. Based on the operational intent, click **Optimize Now**.

   The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.
5. If you are satisfied with the projected results of the optimization action, click **NEXT**.
6. Review the optimization moves, then click **BEGIN ACTION**.

   In the scope of VMware Aria Automation integration, VMware Aria Operations sends a move migration request directly to VMware Aria Automation. In the earlier versions, the migration request was sent to the vCenter Server.

### What to do next

To verify that the optimization action is complete, select **Administration** from the left menu, and click **Recent Tasks** . In the **Recent Tasks** page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

# Pricing for VMware Aria Automation Components in VMware Aria Operations

After you integrate VMware Aria Automation private cloud adapter instances with VMware Aria Operations, you can calculate the cost of deployments, projects, and virtual machines of the selected cloud adapter. Pricing provides an overview of the costs related to the cloud environment, cloud resources, and the costs associated with the project.

## How the Pricing Works in VMware Aria Automation

- VMware Aria Operations understands the constructs defined in VMware Aria Automation and calculates the CPU, RAM, Storage and Additional prices for Projects, Deployments, and virtual machines.
- A single project can have multiple deployments and a single deployment can have multiple virtual machines associated with the deployment.
- Pricing for multiple virtual machines associated with the deployment is the sum of all the resources associated with individual virtual machines.
- If a single project has multiple deployments, then the project pricing is equal to the sum of individual deployments. The deployment can have multiple virtual machines and resources associated with it.
- On day one, the pricing is equal to the cost of resources defined in VMware Aria Operations.
- On day two, the price is calculated using the following formula.
  - Cost of resources for the present day – Cost of resources for the previous day
- If in case the pricing does not happen as per the definition, then the partial price is set to true, and the pricing is calculated based on the previous days price.
- In VMware Aria Operations, the following new dashboards are included to view the pricing details for the VMware Aria Automation instances.
  - Cloud Automation Environment Overview
  - Cloud Automation Project Cost Overview
  - Cloud Automation Resource Consumption Overview
  - Cloud Automation Top-N Dashboard

## Data Collection Enhancements in VMware Aria Automation for Pricing in VMware Aria Operations

The following enhancements have been made for the data collection process from VMware Aria Automation for pricing purposes.

- Collect cloud zones with relation to clusters and resources pools from VMware Aria Automation to VMware Aria Operations.
- Collect Projects from VMware Aria Automation with relation to deployments.

- Include project, cloud zone, and blueprint as properties in virtual machines that are deployed in VMware Aria Automation.

### Upfront Price Support for VMware Cloud on AWS Resources

VMware Aria Operations supports upfront pricing for VMware Cloud on AWS resources in the following ways:

- VMware Aria Operations supports upfront pricing for VMware Cloud on AWS only if rate-based pricing is configured in VMware Aria Automation for VMware Cloud on AWS resources.
- VMware Aria Operations does not support cost-based computation for VMware Cloud on AWS resources.

### Upfront Pricing Support for Metering Policy

VMware Aria Operations supports upfront pricing for metering policy in the following ways:

- VMware Aria Operations supports tag-based metering policy for cost calculation using virtual machines with specific key and value. Virtual machines can be charged on a per day basis.
- VMware Aria Operations supports metering policy with one-time charges for virtual machines on a daily basis.
- VMware Aria Operations supports metering policy for specific operating systems.
- VMware Aria Operations supports custom properties in metering policy for calculating cost of resources in virtual machines.

## Managing Public Cloud Endpoints with VMware Aria Automation Integration

With the management pack, you can monitor deployments made to public cloud endpoints such as Amazon Web Services (AWS), Google Cloud Platform, and Microsoft Azure. You can monitor the performance, health, utilization, and availability attributes of deployments made to public cloud endpoints.

This integration supports the management packs for public clouds – MP for AWS, GCP, and Azure. Hence, the public cloud MPs are a prerequisite for this enhancement.

will display the VMware Aria Automation deployed by public cloud resources as long as the same resources are also monitored by for the respective public cloud MP.

This enhancement also shows all the cloud accounts and cloud zones that are part of public cloud end points.

**Note** To monitor the VMware Aria Automation resources deployed to AWS, GCP, and Azure, you must ensure that the cloud accounts are configured in both VMware Aria Automation and .

### Import Account Functionality

Users may import AWS, Azure, and vCenter accounts automatically through the **Import Account** option. However, for the GCP account you should manually add the cloud account through **Add Account** option. The credential used to configure the adapter must match the one used in VMware Aria Automation. The traversal specification of VMware Aria Automation has been enhanced to include details about AWS, Azure, and GCP accounts.

## Cloud Zones in VMware Aria Operations

Cloud zones help you to group a set of compute resources and assign capability tags to the zone. The cloud zone is based on accounts/regions, so you must have at least one cloud account configured before you can create a cloud zone. Cloud zones define where and how blueprints configure deployments. You can have one or many cloud zones assigned to each project based on priority and limits.

### How Cloud Zones Work

After you integrate VMware Aria Automation with VMware Aria Operations, you can retrieve cloud zones into VMware Aria Operations. The **Cloud Zones** option is hidden from the user until the integration with VMware Aria Automation is activated from the integration page under **Administration** > **Management**.

The Cloud Zones option is activated in VMware Aria Operations, only if the following conditions are met.

- VMware Aria Automation instance is integrated successfully in VMware Aria Operations **Administration** > **Management**>**Integrations**.
- VMware Aria Automation objects are discovered in VMware Aria Operations.
- VMware Aria Automation accounts and Aria Operations vCenter Cloud Accounts are synchronized.

All the Cloud Zone objects which are existing in VMware Aria Automation environment, are discovered in VMware Aria Operations. Cloud zones, whose dependent clusters are not discovered in VMware Aria Operations, are not represented in Capacity Overview, Reclaim, and Workload Optimization pages.

### Cloud Zones List

You can view the list of cloud zones that exist in your environment. In this view, you can click a cloud zone to display all the resources and objects that are associated with the cloud account. When you click the Cloud Zone, you are directed to the standard object summary page of the cloud account.

#### Where You Find Cloud Zones

Select **Environment** in the menu and click **Cloud Zones** tab.

#### Cloud Zone Tab Options

| Option | Description |
|---|---|
| Name | Displays the name of the selected cloud zone. |
| Cloud Account | Displays the cloud accounts associated with the cloud zone. |
| Resources | Displays the cloud account resources associated with the cloud zone.<br>**Note** If the resource field is empty, it means VMware Aria Operations does not have a corresponding vCenter Cloud Account for that associated Cloud Zone. Add a new vCenter Cloud Account manually or use the Import Cloud Account option from the Cloud Account page. |
| Capability Tags | Displays the capability tags associated wit the cloud zone. |

# vSAN

You can make vSAN operational in a production environment by using dashboards to evaluate, manage, and optimize the performance of vSAN objects and vSAN-activated objects in your vCenter Server system.

vSAN extends the following features:

- Discovers vSAN disk groups in a vSAN datastore.
- Identifies the vSAN-activated cluster compute resource, host system, and datastore objects in a vCenter Server system.
- Automatically adds related vCenter Server components that are in the monitoring state.
- Support for vSAN datastores in workload optimization with cross-cluster rebalance actions.
  - You can move VMs from one vSAN datastore to another vSAN datastore.
  - You can optimize the container if all the vSAN clusters are not in resync state.
  - VMs with different storage policies for each disk or VMs with different types of storage for each disk will not be moved.
  - You can generate a rebalance plan only if sufficient disk space is available at the destination vSAN datastore (The vSAN datastore slack space will also be considered).
  - The storage policy assigned to the VM will be considered during the workload optimization (Compatibility check is performed against the storage policy).
  - VM migration from vSAN datastore to vSAN stretched clusters is not supported.

## Configure a vSAN Adapter Instance

When configuring an adapter instance for vSAN, you add credentials for a vCenter Server. In the earlier versions of VMware Aria Operations, the vSAN solution was installed as part of the VMware Aria Operations installation. Now, in case of a new installation the vSAN solution is pre-bundled as part of VMware Aria Operations OVF, you must install the vSAN solution separately.

### Prerequisites

Only vCenter Server systems that are configured for both the vCenter adapter and the vSAN adapter appear in the inventory tree under the vSAN and Storage Devices. Verify that the vCenter Server that you use to configure the vSAN adapter instance is also configured as a vCenter adapter instance for the VMware vSphere® solution. If not, add a vCenter adapter instance for that vCenter Server.

You must open port 5989 between the host and any VMware Aria Operations node on which the vSAN adapter resides. This is applicable when the vSAN version in vSphere is 6.6 or lower.

You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to monitor the vSAN and Storage Devices.

To know how to install the Native Management Packs, see The Repository Tab.

### Procedure

1. From the left menu, select **Data Sources** and then click **Integrations > ADD**.
2. From the **Account Types** page, select the vCenter Server instance from the list and then click the **vSAN** tab.
3. To use the vCenter Server for activating vSAN, move the **vSAN configuration** option to the right.

   **Note** Once the vSAN adapter instance is activated and saved, the activated vSAN configuration option is not visible.

4. The credentials provided for the vCenter Server instance are also used for vSAN adapter instance. If you do not want to use these credentials, you can click **Use alternate credentials** option.
   a. Click the plus sign next to the Credential field and enter the details in the **Manage Credentials** dialog box.
   b. Enter the credential name, vCenter user name, and password and click **OK**.
5. Choose **Enable SMART data collection**, to activate SMART data collection for physical disk devices.
6. Click **Add**.

   The vSAN configuration is activated for the cloud account.
7. Click **Test Connection** to validate the connection with your vCenter Server instance.
8. Accept the vCenter Server security certificate.
9. Click **Save Settings**.

### Results

The adapter is added to the Adapter Instance list and is active.

### What to do next

To verify that the adapter is configured and collecting data from vSAN objects, wait a few collection cycles, then view application-related data.

- Inventory. Verify that all the objects related to the vSAN instance are listed. Objects should be in the collecting state and receiving data.
- Dashboards. Verify that vSAN Capacity Overview, Migrate to vSAN, vSAN Operations Overview, and Troubleshoot vSAN, are added to the default dashboards.
- Under **Environment > vSAN and Storage Devices**, verify that the vSAN hierarchy includes the following related vCenter Server system objects:
  - vSAN World
  - Cache Disk
  - Capacity Disk
  - vSAN-activated vCenter Server clusters
  - vSAN Fault Domains (optional)
  - vSAN-activated Hosts
  - vSAN Datastores
  - vSAN Disk Groups
  - vSAN Datastore related VMs
  - vSAN Witness Hosts (optional)

## Verify that the Adapter Instance is Connected and Collecting Data

You configured an adapter instance of vSAN with credentials for a vCenter Server. Now you want to verify that your adapter instance can retrieve information from vSAN objects in your environment.

To view the object types, from the left menu, click **Environment > Inventory > Adapter Instances > vSAN Adapter Instance > <*User_ Created_ Instance*>**.

### Table 4-77. Object Types that vSAN Discovers

| Object Type | Description |
|---|---|
| vSAN Adapter Instance | The VMware Aria Operations Management Pack for vSAN instance. |
| vSAN Cluster | vSAN clusters in your data center. |
| vSAN Datastore | vSAN datastores in your data center. |
| vSAN Disk Group | A collection of SSDs and magnetic disks used by vSAN. |
| vSAN Fault Domain | A tag for a fault domain in your data center. |
| vSAN Host | vSAN hosts in your data center. |

Table 4-77. Object Types that vSAN Discovers (continued)

| Object Type | Description |
| --- | --- |
| vSAN Witness Host | A tag for a witness host of a stretched cluster, if the stretched cluster feature is activated on the vSAN cluster. |
| vSAN World | A vSAN World is a group parent resource for all vSAN adapter instances. vSAN World displays aggregated data of all adapter instances and a single root object of the entire vSAN hierarchy. |
| Cache Disk | A local physical device on a host used for storing VM files in vSAN. |
| Capacity Disk | A local physical device on a host used for read or write caching in vSAN |

The vSAN adapter also monitors the following objects discovered by the VMware vSphere adapter.

- Cluster Compute Resources
- Host System
- Datastore

### Procedure

1. In the menu, click **Administration** and then in the left pane, click **Configuration > Inventory** .
2. In the list of tags, expand **Adapter Instances** and expand **vSAN Adapter Instance**.
3. Select the adapter instance name to display the list of objects discovered by your adapter instance.
4. Slide the display bar to the right to view the object status.

   | Object Status | Description |
   | --- | --- |
   | Collection State | If green, the object is connected. |
   | Collection Status | If green, the adapter is retrieving data from the object. |

5. Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

### What to do next

If objects are missing or not transmitting data, check to confirm that the object is connected. Then check for related alerts.

To ensure that the vSAN adapter can collect all performance data, the Virtual SAN performance service must be activated in vSphere. For instructions on how to activate the service, see Turn on Virtual SAN Performance Service in the VMware Virtual SAN documentation.

If the Virtual SAN performance service is deactivated or experiencing issues, an alert is triggered for the vSAN adapter instance and the following errors appear in the adapter logs.

```
ERROR com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
   - Failed to collect performance metrics for Disk Group
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
   - vSAN Performance Service might be turned OFF.
com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics
   - (vim.fault.NotFound)
   {
   faultCause = null,
   faultMessage = (vmodl.LocalizableMessage)
      [
         com.vmware.vim.binding.impl.vmodl.LocalizableMessageImpl@98e1294
      ]
   }
```

## vSAN Log Analytics Enhancements

When VMware Aria Operations is integrated with VMware Aria Operations for Logs, you can view and troubleshoot VMware Aria Operations for Logs object issues within VMware Aria Operations. Earlier you could troubleshoot issues related only to vCenter objects, but now you can troubleshoot issues related to vSAN also.

The enhancements to vSAN log analytics include use of specific queries to retrieve log information for the following vSAN objects:

- vSAN Cluster
- Witness Host
- Disk Group
- Cache Disk
- Capacity Disk

### Where You Find vSAN Object Logs

Navigate to the vSAN Object Details page, and click the **Logs** tab.

**Note** If you are not logged in to VMware Aria Operations for Logs, then VMware Aria Operations prompts you to log in to VMware Aria Operations for Logs with your login credentials.

VMware Aria Operations uses special queries for each object type. Using the special queries for vSAN objects, you can perform the following actions:

- View interactive analytics for the selected vSAN object.
- Retrieve log details for the vSAN object.
- Analyze and troubleshoot issues related to the vSAN object.

# VMware Aria Operations for Networks

The VMware Aria Operations for Networks adapter activates integration of VMware Aria Operations with VMware Aria Operations for Networks. VMware Aria Operations for Networks provides network visibility and analytics to minimize risk during application migration, optimize network performance, manage and scale VMware NSX-T, VMware NSX for vSphere, vCenter on VMware Cloud on AWS, VMware SD-WAN by VeloCloud, and Kubernetes deployments.

- Cloud Integration can be done, when VMware Aria Operations for Networks and VMware Aria Operations services are present in the same organization.
- For the cloud integration to be possible, you must have VMware Aria Operations for Networks and VMware Aria Operationsservices in the same geographical location.

This adapter gets problem events from VMware Aria Operations for Networks and publishes the alerts in VMware Aria Operations. Alerts are mapped correctly to the common objects between VMware Aria Operations for Networks and VMware Aria Operations. Common objects supported in this adapter are vCenter Server, VMware NSX-T, and VMware NSX for vSphere. For the common objects, supports launch-in-context to VMware Aria Operations for Networks. This allows the user to perform deep network troubleshooting with the VMware Aria Operations for Networks as the context.

The VMware Aria Operations for Networks adapter only supports VMware Aria Operations for Networks versions 5.2 and above. The VMware Aria Operations for Networks adapter can be installed and configured with the on-prem and SaaS versions of . The VMware Aria Operations for Networks adapter does not support cross platform configuration, it should be on-prem to on-prem VMware Aria Operations for Networks and SaaS to VMware Aria Operations for Networks SaaS.

# NSX

The NSX adapter allows you to retrieve alerts and findings from NSX to VMware Aria Operations.

The NSX adapter supports adapter configuration using vIDM for NSX versions 3.0 and above. The roles and permissions associated with the vIDM users collecting the NSX adapter data is:

| Roles | Permissions |
| --- | --- |
| Enterprise Admin | Collect all data. |
| Network engineer | ■ Collect all the NSX resources except the Load Balancer and collect limited routers data.<br>Router data collected:<br>■ Tier 0 router connected to logical switch.<br>■ Tier 1 router created from vCloud Director. |

| Roles | Permissions |
| --- | --- |
| ■ Security Engineer<br>■ Security Operator<br>■ Auditor | Collect all data except the load balancer. |
| ■ LB Admin<br>■ LB Auditor<br>■ Netxpartner Admin | Cannot collect any data. |

## Configuring the NSX Adapter

### Prerequisites

To view the roles and permissions associated with the VIDM users collecting the NSX adapter, see NSX

- To use Principal Identities authentication in VMware Aria Operations you must have created a Principal Identity user in NSX.
- Ensure that you have the client certificate and the key to authenticate the Principal Identity users in VMware Aria Operations.

### Procedure

1. From the left menu, click **Data Sources** > **Integrations**.
2. On the **Accounts** tab, click **Add**, and on the **Accounts Types** page, select **NSX**.
3. Configure the adapter instance.

| Option | Action |
| --- | --- |
| Name | Enter the name for the NSX instance as you want it to appear in VMware Aria Operations. |
| Description | Enter any additional information that helps you manage your instances. |

4 Configure the connection.

| Option | Action |
| --- | --- |
| Virtual IP/NSX Manager | Enter the FQDN, the IP address, or the Virtual IP of the NSX manager. Both IPv4 address and IPv6 address formats are supported. |
| Credential | Select the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the NSX environment, click the plus sign.<br>1 From the **Manage Credential** dialog, you can click the **Credential Kind** drop-down and select **NSX Client Certificate Crentials** or **NSX Credentials**.<br>The NSX Client Certificate allows you to use Principal Identities user or certificate-based client for authentication and the NSX Credential allows you to use local administrator or VMware Identity Manager for authentication.<br>2 Enter the credential details based on your selection.<br>For NSX Credentials, enter the following details.<br>■ Credential Name - The name by which you are identifying the configured credentials.<br>■ User Name - The user name of the NSX instance.<br>■ Password - The password of the NSX instance.<br>For NSX Client Certificate Credentials, enter the following details.<br>■ Credential Name - The name by which you are identifying the configured credentials.<br>■ Client Certificate Data - Enter the value of the client certificate data associated with the principal user.<br>■ Client Key Data - Enter the value of the client key data associated with the principal user. |

| Option | Action |
| --- | --- |
| Collector /Group | Determine which VMware Aria Operations collector or collector group is used to manage the account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.<br>**Review and Accept Untrusted Certificates**<br>When you upgrade VMware Aria Operations from an earlier version to VMware Aria Operations 8.6, the NSX adapter moves to a warning state and the data collection stops. This happens only when you have an adapter that presents a self-signed certificate, or a certificate signed by an untrusted Certification Authority.<br>To continue the adapter configuration, you must Validate the Connection, where you are prompted to review and accept the certificate.<br>**Note** If you have a multi-node cluster configuration, then you are prompted to review and accept the certificate for each node. |
| Advanced Settings | |
| Exclude Resources from Monitoring | Select specific NSX resources that you want to exclude from monitoring in VMware Aria Operations. Once the exclusion is applied, the monitoring for the specified resources will be disabled.<br>**Note** Earlier the Management Pack for NSX would collect all the data and monitor all the resource kinds associated with the Management Pack, but now it allows you to exclude specific resources from monitoring. |
| Exclude Services from Monitoring | Select specific NSX management services that you want to exclude from monitoring in VMware Aria Operations. Once the exclusion is applied, the monitoring for the specified services will be disabled.<br>**Note** Earlier the Management Pack for NSX would collect all the data and monitor all the services associated with the Management Pack, but now it allows you to exclude specific services from monitoring. |

| Option | Action |
| --- | --- |
| Auto Discovery | Set the **Auto Discovery** to True or False.<br>■ **Auto Discovery - True** - Activates auto-discovery of the new objects added to the monitored system. By default, this field is set to True.<br>■ **Auto Discovery - False** - Deactivates auto discovery, you must manually discover the objects from the system which you want to monitor . |
| Import Alerts from NSX | You can directly import alerts from VMware NSX (formerly known as VMware NSX-T) into VMware Aria Operations by setting the **Import Alerts from NSX** field to **True**.<br>Depending on the number of alerts and the system's performance, importing alerts into VMware Aria Operations may take some time.<br>By default, this field is set to **False**.<br>**Note** Direct Alert integration is not supported on VMware Cloud on AWS and VMware Cloud on Dell. |

5 Click **Add**.

The adapter instance is added to the list.

#### What to do next

Verify that the adapter is configured and is collecting data.

## Support for Principal Identities Authentication for the NSX Management Pack

VMware Aria Operations supports authentication of Principal Identities (PI) using the NSX Management Pack. The Principal Identities (PI) are unique users in NSX who can create an object and ensure that the object can only be modified or deleted by the same identity. The authentication of principal identities is only supported through client certificate. The principal identities authentication is local to NSX Manager, so it does not require VMware Identity Manager, and it is possible to assign a predefined Role-based access control (RBAC) role to the principal identity.

Principal Identities are generally used by third-party applications or cloud management platforms such as Open stack, and Pivotal Container Services (PKS) to ensure that an administrator does not modify the NSX configuration which can generate a mismatch between their view of the NSX environment and the actual configuration.

# Configuring Cloud Federation Adapter

You can configure the adapter instances and add the VMware Aria Operations and VMware Aria Operations SaaS instances that you want to monitor.

### Procedure

1 From the left menu, click **Data Sources > Integrations**.

2 Under the Accounts tab, click **Add Account**.

3 On the Account Types page, click the **Cloud Federation Adapter**.

4 Enter the display name and the description of the adapter.

5 In the **Organization ID** field, enter the VMware Cloud Services ID, if you are creating an instance that monitors VMware Aria Operations SaaS.

**Note** The Organization ID must be left empty if you are creating an instance that monitors VMware Aria Operations.

6 Enter the Host name or the IP address of the VMware Aria Operations or VMware Aria Operations SaaS source.

**Note** While entering the URL for the VMware Aria Operations SaaS source, ensure you enter the URL in the format `http://www.host-name.com/ariaops`.

7 Select and add the credential you want to use from the drop-down menu. To add new credentials, click the plus sign.

From the Credential Kind drop-down:

- If you want to monitor the VMware Aria Operations instances, Select the API Token Credentials and enter the credentials.
- If you want to monitor VMware Aria Operations instances, Select the Principal Credentials and enter the credentials.

8 From the **Collector / Group** drop-down, select Default Collector Group to monitor the VMware Aria Operations instances.

To monitor VMware Aria Operations instances, select the appropriate cloud proxies from the **collector / group** drop-down.

9 Click **Validate Connection** to verify if the configuration is successful.

10 Click **Add**.

### Results

The Cloud Federation Adapter is added to the list.

## Google Cloud Platform

The VMware Aria Operations Management Pack for Google Cloud Platform allows you to dive into the key performance indicators for your Google Cloud Platform environment. Each adapter instance of this Management Pack has diagnostic dashboards and collects metrics and properties from Google Cloud.

This Management Pack supports the following Google Cloud Platform products:

- Compute Engine
- Container Engine
- Cloud Storage
- Cloud VPN
- Big Query

## Supported GCP Services

The Management Pack for Google Cloud Platform supports the following services.

| Service | Object | Description |
|---|---|---|
| Dataproc | Dataproc Cluster | A managed Spark and Hadoop service that allows you to take advantage of open source data tools for batch processing, querying, streaming, and machine learning |
| Cloud Load Balancing (HTTP(S) Load Balancers only) | Load Balancer | Distributes traffic across multiple instances of your applications and regions. This helps in reducing the risk of performance issues that your applications experience. |
| CE Instance (K8s Nodes included) | Persistent Disk | Provides secure and customizable compute service that lets you create and run virtual machines on Google's infrastructure |
| Google Kubernetes Engine | ■ K8s Clusters<br>■ K8s Container<br>■ K8s Pods | Provides managed environment for running containerized apps |
| Big Query | ■ Big Query Dataset<br>■ Big Query Table | Provides data warehouse for business agility and insights |
| Cloud VPN | ■ VPN Gateways<br>■ VPN Tunnels | Allows you to connect your infrastructure to Google Cloud Platform (GCP) on your terms, from anywhere. (Part of Hybrid Connectivity) |
| VPC Network | VPC Network | Provides virtual network for Google Cloud resources and cloud-based services |
| Cloud Storge | Storage Buckets | Provides object storage that's secure, durable, and scalable |
| Cloud SQL | Cloud SQL | Provides fully managed database for MySQL, PostgreSQL, and SQL Server |

| Service | Object | Description |
| --- | --- | --- |
| Memorystore | ■ Memorystore Redis<br>■ Memorystore Memcached<br>  ■ Memcached Node | Provides in-memory database for managed Redis and Memcached |
| Cloud Spanner | Cloud Spanner | Provides cloud-native relational database with unlimited scale and 99.999% availability |
| Sole-tenant Node Group | Sole-tenant Node Group | Provides dedicated hardware for compliance, licensing, and management |
| Filestore | Filestore | Provides file storage that is highly scalable and secure |
| Node Pool | Node Pool | Provides collection of CE Instances (Node) that are created by the K8 Cluster. |
| Cloud Bigtable | ■ Cloud BigTable Cluster<br>■ Cloud BigTable<br>■ Cloud BigTable Table | Provides cloud-native wide-column database for large-scale, low-latency workloads |
| Firebase Realtime Database | Firebase Realtime Database | Provides NoSQL database for storing and syncing data in real time |
| Firestore Database | Firestore Database | Provides cloud-native document database for building rich mobile, web, and IoT applications |
| Cloud Run | ■ Cloud Run<br>■ Cloud Run Revision<br>■ Cloud Run Service | Provides a fully managed environment for running containerized applications |
| App Engine | ■ App Engine<br>■ App Engine Service | Provides a serverless application platform for application and back end |

## Charges for GCP Metrics

GCP charges you for the metrics you collect. You can reduce costs by selecting only the metrics that are most helpful and filtering out those that are of less interest.

By default, the GCP requests data every five minutes. Every collection cycle makes one Cloud Monitoring call per metric, for an object.

For information about metric costs, see Google Cloud's Monitoring Charges.

Based on the costs associated with running the adapter, you can take advantage of some of the features that limit the amount of data you collect from GCP.

- Turn off auto discovery and use manual discovery.
- Go to **Advanced Settings** and select only those services that are critical to your system.
- Subscribe only to specific critical regions.

- Use acceptlist and blocklist filtering to select object import by name.
- Use acceptlist or blocklist to selectively import specific service instances using instance names. Using regex, you can filter services by providing the partial name of the instance.

## Configuring Google Cloud Platform

After the installation, configure an instance for this management pack.

### Prerequisites

- Install the Management Pack for Google Cloud Platform.
- Obtain the service account JSON file. See, Creating a Service Account.
- Determine the services for which you collect metrics. See, Supported GCP Services. The default value * includes all services in your subscription. If you do not want to subscribe to all services, you can specify service identifiers in the Services field.
- Determine the regions to which you subscribe. The default value * includes all regions in your subscription. If you do not want to subscribe to all regions, you can specify region identifiers in the Regions field.

Table 4-78.

| Region-Friendly Name | Region Identifier |
|---|---|
| Asia Pacific (Taiwan) | asia-east1 |
| Asia Pacific (Hong Kong) | asia-east2 |
| Asia Pacific (Tokyo) | asia-northeast1 |
| Asia Pacific (Osaka) | asia-northeast2 |
| Asia Pacific (Seoul) | asia-northeast3 |
| Asia Pacific (Mumbai) | asia-south1 |
| Asia Pacific (Singapore) | asia-southeast1 |
| Asia Pacific (Sydney) | australia-southeast1 |
| EU (Finland) | europe-north1 |
| EU (Belgium) | europe-west1 |
| EU (London) | europe-west2 |
| EU (Frankfurt) | europe-west3 |
| EU (Netherlands) | europe-west4 |
| EU (Zürich) | europe-west6 |
| South America(Canada) | northamerica-northeast1 |

Table 4-78. (continued)

| Region-Friendly Name | Region Identifier |
| --- | --- |
| South America (Osasco) | southamerica-east1 |
| Iowa (Central) | us-central1 |
| US East (South Carolina) | us-east1 |
| US East (Northern Virginia) | us-east4 |
| US West (Oregon) | us-west1 |
| US West (Los Angeles) | us-west2 |
| US West (Salt Lake City) | us-west3 |
| ASIA (Multi-regions) | asia |
| EU (Multi-regions) | eu |
| US (Multi-regions) | us |

- Determine any denylist or allowlist filters. These filters use regular expressions to filter in or out specific objects by name. For example, a allow list filter of `.*indows.*` allows only objects with a name including "indows". A denylist filter of `.*indows.*` filters out all objects with that string in their name.

Procedure

1 From the left menu, click **Data Sources > Integrations**.

2 On the Accounts tab, click **Add Account**, and on the Accounts Types page, select **GCP**.

3 Configure the adapter instance.

| Option | Action |
| --- | --- |
| Name | The name for the adapter instance. |
| Description | (Optional) The description of the adapter instance. |

4 Configure the connection.

| Option | Action |
| --- | --- |
| Project Id | Enter the Google Cloud Platform service account project Id. |
| Services | Select the services from which you want to capture metrics. If you want to collect metrics for specific services, then click the drop-down icon and select one or more services. If you do not select any of the services, the metrics for all the services are collected. For more information on supported GCP services, see Supported GCP Services. |

| Option | Action |
| --- | --- |
| Regions | Select the regions you want to subscribe to. If you want to subscribe to specific regions, then click the drop-down icon and select one or more regions. If you want to subscribe to all the regions, do not select any of the regions in this field. |
| Credential | Select the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access this management pack environment, click the plus sign.<br>■ Enter a name for the credential values you are creating. This is not the name of the adapter instance, but a friendly name for the Service account JSON file credential.<br>■ Enter the service account private key downloaded as a JSON file.<br>■ (Optional) Enter any required local proxy information for your network.<br>**Note**<br>■ The proxy information activates VMware Aria Operations that is deployed on a local network to have public network access for GCP to collect statistics for its accounts.<br>■ VMware Aria Operations supports only basic auth. |
| Collector / Group | Select the collector on which you want to run the adapter instance. A collector gathers objects into its inventory for monitoring. The collector specified by default is selected for optimal data collection. |

5 Click the arrow to the left of the **Advanced Settings** to configure advanced settings.

| Option | Action |
| --- | --- |
| Support Auto Discovery | Set this option to true for automatic discovery of the Google Cloud Platform services. If you set this value to false, when you create an adapter instance you must perform a manual discovery of services. |
| Allowed List Regex | Add regular expressions to allow only objects with names that fit the criteria you specify. |
| Blocked List Regex | Add regular expressions to filter objects by name. |
| Actions Enabled | Activate this option to run actions specific to the GCP CE instance using VMware Aria Operations. By default, this option is set to **true**.<br>For details on running actions, see Power On, Power Off, and Reboot Actions. |

| Option | Action |
| --- | --- |
| Enable CE GPU Metric Collection | This option allows collection of GPU metrics for GCP Compute Engine (CE) instances. By default, this value is set to **False**.<br>If this option is set to **True**, then the VMware Aria Operations collects GPU metrics for all CE instances as GCP does not distinguish between normal CE instances and GPU instances.<br>**Note** Refer to the instructions in the GCP documentation to activate Cloud GPU metrics before enabling VMware Aria Operations for data collection. |
| Enable CE SQL Server Metric Collection | This option allows collection of SQL server metrics for GCP Compute Engine (CE) instances. By default, this value is set to **False**.<br>If this option is set to **True**, then the VMware Aria Operations collects SQL server metrics for all CE instances as GCP does not distinguish between normal CE instances and SQL server instances. |

6 To initiate the authentication request, click **Validate Connection**.

7 Click **Save**.

The adapter instance is added to the list.

### What to do next

Verify that the adapter is configured and collecting data by viewing application-related data.

| Where to View the Information | Information to View |
| --- | --- |
| Collection status and state in the **Integrations** page | The strings *Collecting* or *Data receiving* appear approximately 10 minutes after you have configured the adapter. |
| Environment Overview | The objects related to the Google Cloud Platform are added to the inventory. |
| Dashboards | The Google Cloud Platform dashboards are added to VMware Aria Operations. |

## Creating a Service Account

To configure the Management Pack for Google Cloud Platform, you must create a service account in Google Cloud Platform and download the private key as a JSON file. To create the service account, you must have the Service Account Admin role (roles/iam.serviceAccountAdmin) or the Editor primitive role (roles/editor). For read-only access, the service account requires the project level viewer role (Viewer - primitive role on GCP).

To monitor the Google Cloud Platform account using this Management Pack, activate the following APIs:

- BigQuery API
- Compute Engine API
- Cloud Storage and Google Cloud Storage JSON API
- Kubernetes Engine API
- Stackdriver Monitoring API (The Stackdriver monitoring API (monitoring.googleapis.com) is required to monitor time-series metric data).

To activate these APIs:

1 In the Cloud Console, navigate to **APIs & Services** for your project.

2 In the **Library** page, search for the above APIs.

3 Select the Service API you want to activate.

4 Click **Enable**.

When the APIs are activated and the service account has the correct set of roles and associated permissions, this Management Pack can retrieve Google Cloud Platform data. When creating a service account, you must select a Google Cloud Platform project as Google Cloud Platform does not allow the service account to belong directly under the Google Cloud Platform Organization.

Procedure

1 In the Cloud Console, navigate to **IAM & Admin > Service Accounts Management > Create Service Account Key**

2 From the Service account list, select **New service ccount**.

3 In the **Service account name** text box, enter a name.

4 From the Role list, The Role field authorizes the service account to access resources. Select **Project > Owner** or select the required services in read-only (as a viewer).

5 Click **Save**.

6 Download the service account private key as a JSON file.

## View Objects for Google Cloud Platform

You can use the inventory tree to browse and select objects. The inventory tree shows a hierarchical arrangement of the GCP objects by region.

Procedure

1 In the left pane of VMware Aria Operations, click the **Environment** icon.

- Click **Accounts** to view the GCP adapter instances.
- Click **Regions** to to view the GCP regions.

2 To view the child objects, expand the regions and then expand the regions per account.

**Note** All the account-specific objects related to a region are grouped under the region per account section.

3 To display information about the object, select an object in the inventory tree.

# Integration of VMware Aria Operations and VMware Aria Hub

VMware Aria Hub is a platform that unifies data across the VMware Cloud Management services. By unifying this data, Ensemble provides visibility into your hybrid and multi-cloud infrastructure, as well as the applications that run on that infrastructure. VMware Aria Hub is a common data layer for VMware products. Once the integration is complete, VMware Aria Operations can create applications which are automatically discovered by VMware Aria Hub, as Business Applications.

## VMware Aria Hub Object Relationship in VMware Aria Operations

The VMware Aria Hub adapter in VMware Aria Operations collects application information from VMware Aria Hub, along with the associated tiers and tier members. The tier members are vCenter Virtual Machines and AWS EC2 instances. VMware Aria Operations does not create the tier member objects, but creates a relationship with tiers after you enable the vCenter and AWS adapters. Virtual Machine and EC2 members will be in Application/Tier relationship only when the respective Cloud Accounts is added in VMware Aria Operations.

To see the relationship between the VMware Aria Hub application, tiers and tier members, from the left menu, click **Environment > Business Applications**. All the imported VMware Aria Hub Business Applications are listed. Click on an VMware Aria Hub Business Application to view more information, including the relationship. You cannot edit the Business Applications which VMware Aria Operations imports from VMware Aria Hub. However, any relationship change made in VMware Aria Hub is automatically reflected in the corresponding VMware Aria Operations Business Application.

## Integrating the VMware Aria Hub Service

The integration of the VMware Aria Operations and VMware Aria Hub services happens automatically if you can access both the services.

To view the applications which are imported by the adapter, from the left menu, click **Environment > Inventory > Adapter Instances > VMware Aria Hub > Adapter Instance > <Adapter_Instance>**. By default, the collection cycle is 30 minutes.

## Imported VMware Aria Hub Business Application Properties:

VMware Aria Operations collects the following properties from VMware Aria Hub applications:

- Creation Time

- Last Modification Time
- Source

Go to the object browser and click **All Objects > VMware Aria Hub Adapter**. VMware Aria Operations lists all the Business Applications which are imported from VMware Aria Hub. Click on a Business Application name to see the summary page of the Business Application. Under the Metrics tab, you can see a list of properties which VMware Aria Operations collects from VMware Aria Hub for the application. The discovered VMware Aria Hub applications can leverage Synthetic Monitoring. For more details, see Configure, Enable, or Disable Synthetic Monitoring.

# Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria Operations

5

VMware supports its partners to host and sell cloud services built on VMware technology using VMware Cloud Director (VCD), also known as Cloud Director Service (CDS) in the SaaS environment. Cloud Director Service provides constructs to segment the virtual infrastructure and offers it as a service to tenants of these partners.

There are several variants of infrastructure that are sold by these partners such as, 'Pay as you go', 'Raw capacity' also known as 'Allocation based', 'Raw capacity with minimum guarantee' also known as 'Reservation based'. The combination of these can be offered to same tenants and it becomes challenging to track usage over a period and charge appropriately. It becomes critical for the cloud service providers as the tenants demand transparency in billing, and the cloud service providers must offer it. Chargeback addresses this by accurately metering the infrastructure. It provides options to configure different models for pricing this metered infrastructure.

**Note** The Chargeback capabilities are available only when you have Cloud Director Service configured in VMware Aria Operations.

This chapter provides information about configuring Chargeback for Cloud Director Service.

**Note** VMware Cloud Director (VCD) is known as Cloud Director Service (CDS) in the SaaS environment.

Read the following topics next:


- Activating Chargeback Capabilities in VMware Aria Operations
- Register VMware Aria Operations in a VMware Cloud Director Instance
- Migration of Chargeback Data to VMware Aria Operations in Case of an Upgrade
- Managing Chargeback Administration Settings
- Configuring VCD Pricing Details
- Creating Notification Rules for Tenant in VMware Aria Operations
- Configuring Tenant Email in VMware Aria Operations
- Managing Chargeback Reports

- Generating Bills in VMware Aria Operations
- Scheduling Bill Generation Using Automation Central
- Viewing Chargeback Summary

# Activating Chargeback Capabilities in VMware Aria Operations

To activate Chargeback capabilities in VMware Aria Operations, the cloud service providers must install VMware Cloud Director Management Pack.

Perform the following steps to activate the capabilities related to Chargeback in VMware Aria Operations.

**Note** If your VMware Aria Operations is integrated with VMware Chargeback in this version of VMware Aria Operations, to activate VMware Aria Operations Chargeback capabilities, you must migrate your data from VMware Chargeback. For more details, see Migration of Chargeback Data to VMware Aria Operations in Case of an Upgrade . After migration, VMware Aria Operations Chargeback capabilities will be activated.

### Procedure

1. From the left menu, click **Data Sources > Integrations**. Click the **Repository** tab. The page displays installed integrations as clickable cards.
2. Get the VMware Cloud Director adapter from the Repository Tab.
3. Additionally, you can configure a VMware Cloud Director instance. For more information see, Add an adapter instance. Optionally, you can add NSX and VMware Aria Operations Management Pack for Cloud Director Availability adapters. For more information, see Configuring the NSX Adapter.
4. To register the VMware Aria Operations plugin, see Register VMware Aria Operations in a VMware Cloud Director Instance.

### Results

All chargeback capabilities are now visible in VMware Aria Operations.

### What to do next

For details on setting up VMware Aria Operations, see Getting Started with VMware Aria Operations.

For more information on VMware Aria Operations, you can refer to the following guides:

- Configuring VMware Aria Operations
- Using VMware Aria Operations
- VMware Aria Operations API Programming Guide

# Register VMware Aria Operations in a VMware Cloud Director Instance

You can register VMware Aria Operations in a VMware Cloud Director instance, for tenant view capabilities.

Procedure

1 Click **Manage Registration** and enter the credentials to register your VMware Aria Operations in the VMware Cloud Director instance.

2 You must complete one of the following:

- Enter a username and password for a user with VMware Aria Operations register/unregister privileges, or
- Skip the username and password and select the **Use Collection Credentials** checkbox.

# Migration of Chargeback Data to VMware Aria Operations in Case of an Upgrade

When you upgrade to this version and later versions of VMware Aria Operations, and you have VMware Chargeback, you have the option to switch to VMware Aria Operations Chargeback and use all the capabilities from one place. This can be achieved through the migration process. The existing settings and configurations will migrate to VMware Aria Operations.

## Steps to Migrate Chargeback Data to VMware Aria Operations

1 From the VMware Chargeback UI, in the left menu, click **Migration**.

2 From the **Migration** page, you will see the list of tasks to be migrated with the Start and End time. You can also view the status of each task that is being migrated. Click **Initiate Migration**.

**Note** Click on each task row to view information about the number of objects that have been migrated, the number of objects that have been ignored, the migrated resources, error messages if any, and so on.

List of tasks in Chargeback and their corresponding location in VMware Aria Operations

Table 5-1. Migration Details from Chargeback to VMware Aria Operations

| Chargeback Migration Steps and Location in Chargeback | Location in VMware Aria Operations |
|---|---|
| **Initialization** | NA |
| **Email Configuration**<br>**Admin Setting > Configure > Configure Email** | **Configure > Alerts > VCD Tenant Email Configuration** |

Table 5-1. Migration Details from Chargeback to VMware Aria Operations (continued)

| Chargeback Migration Steps and Location in Chargeback | Location in VMware Aria Operations |
| --- | --- |
| **Pricing Policies**<br>**Pricing > Configuration** | **Configure > Policies > Policy Definition**.<br>**Pricing Policies Assignment**<br>All pricing policies from Chargeback will be converted to separate policies under the default Policy in VMware Aria Operations. They will have the same name as they have in the Chargeback UI with the pricing card UUID added as a suffix. After migration Active Polices are the ones that have at least one object assignment in Chargeback.<br>**Pricing Policies Assignment Migration**<br>■ If there is just one active policy in VMware Aria Operations before migration, then all the policy assignments from Chargeback will be migrated to VMware Aria Operations accordingly.<br>■ If there is more than one active policy in VMware Aria Operations, then after migration, the policy assignment of Organization and Organization VDC objects should be done manually. In this scenario it is recommended to proceed with policy assignments immediately to ensure continuous pricing calculations. |
| **Bill Schedules** | Bill schedules migration should be done manually. For more details about scheduling bills, see Scheduling Bill Generation Using Automation Central. |
| **Notification Rules**<br>**Alerts > Tenant Alerts > Notification Rules** | **Configure > Alerts > VCD Tenant Notifications** |
| **Report Schedules**<br>**Reports > Tenant Reports** | **Visualize > Reports > Manage**<br>Filter by **VCD Tenants** and then view the **Schedules** column to view the Scheduled reports. |
| **Tenant Reports**<br>**Reports > Tenant Reports** | **Visualize > Reports > Generated Reports**<br>Filter by **VCD Tenants** and view the generated reports. |
| **Generated Bills**<br>**Pricing > Bills** | **Chargeback > Bills** |

Table 5-1. Migration Details from Chargeback to VMware Aria Operations (continued)

| Chargeback Migration Steps and Location in Chargeback | Location in VMware Aria Operations |
|---|---|
| **Bill Retention Settings**<br>**Admin Settings > Settings > Data Retention** | **Administration > Chargeback > Data Retention > Generated Bills** |
| **Finalization**<br>**Admin Settings > Support > Self Health** | **Administration > Audit > Chargeback Migration Audit**<br>■ Users created with the role of VCD Tenant Admin will be removed.<br>■ All VCD instances which had VMware Chargeback registered will now be switched to VMware Aria Operations.<br>  ■ Manual registration may be required in certain cases. These instances will be reported in the Finalization step summary.<br>■ VMware Aria Operations Chargeback capabilities will be activated.<br>■ A detailed report of migration can be found at **Administration > Audit > Chargeback Migration Audit**. For more details, see Chargeback Migration Audit. |

# Managing Chargeback Administration Settings

Cloud service providers can manage the access of metrics and pages by configuring specific metrics or pages that the tenants are permitted to access. Cloud service providers can also define storage pricing policies and manage historical data related to data retention in VMware Aria Operations.

Ensure that you have activated VMware Cloud Director, for details, see Activating Chargeback Capabilities in VMware Aria Operations.

## Managing Access to Metrics

Cloud service providers can manage the access of metrics to the tenants in VMware Aria Operations. You can configure specific metrics that your tenants are permitted to access.

### Procedure

1 To manage access of metrics, from the left menu, click **Administration > Chargeback**.

2 In the **Manage Metrics** tab, select an object type for which you want to activate or deactivate metric access from the **Select Object Type** list.

**Note** By default, all metrics are deactivated.

3 To activate or deactivate a particular metric, select the metric and select **Activate** or **Deactivate** from the **Actions** drop-down list. You can also activate or deactivate metric by clicking the drop-down arrow against a specific metric and selecting **Activate** or **Deactivate**.

**Note** To activate or deactivate all metrics, click the Select All icon and select **Activate** or **Deactivate** from the **Actions** drop-down list.

4 Click **Save**.

## Managing Access to Pages

Cloud service providers can manage the access of pages to the tenants in VMware Aria Operations. You can configure specific pages that your tenants are permitted to access.

### Procedure

1 To manage access of pages, from the left menu, click **Administration > Chargeback**.

2 Click the **Manage Pages** tab, and select the pages for which you want to grant access to the tenant.

3 Click **Save**.

## Defining Storage Pricing

Cloud service providers can charge for storage either based on storage policies or independent of it in VMware Aria Operations.

Storage prices are calculated based on usage by storage policies that are independent of underlying VMs and templates, or by aggregating the usage from underlying VMs and templates. The difference is that the former considers indirect disks such as log disks and swap disks, whereas the latter considers only the storage used directly by the VMs.

### Procedure

1 To define price setting, from the left menu, click **Administration > Chargeback**.

2 Click the **Price Settings** tab, and select an option to aggregate storage charges for PAY AS YOU GO (PAYG) org-VDCs.

   a To charge independent of storage policies, select **Aggregate storage charges from VMs, Media, vAPP templates and independent disks for PAYG org-VDCs**.

   b To charge based on storage policies, select **Aggregate storage charges from Storage profiles for PAYG org-VDCs**.

3 Click **Save**.

**Note** The **Default Base Rate** option under **Policy Definition > VCD Pricing > Storage Rate** is considered only when you have selected Aggregate storage charges from Storage profiles for payg Org-VDCs in the **Pricing Settings** tab.

## Managing Historical Data

You can view and manage historical data retention related to system settings. You can tune these settings to manage storage space consumed by VMware Aria Operations.

### Procedure

1. To view and manage historical data, from the left menu, click **Administration > Chargeback**.
2. Click the **Data Retention** tab, and enter the number of days for which you want to retain the generated bills.

   **Note** You can retain the generated bills to a maximum of 1440 days.

3. Click **Save**.

# Configuring VCD Pricing Details

You can edit the parent policy settings and configure your OVDC settings using the **VCD Pricing** tab in VMware Aria Operations.

If you want to copy the VCD pricing settings from the policy currently being edited to another policy, click **Copy local changes to other policy** and select the policy to which you want to copy the settings. The copied pricing configuration will override any existing local pricing configuration in the target policy.

### Procedure

1. From the left menu, click **Configure > Policies > Policy Definition**, and then select the required policy or click **Add** to add a new policy.
2. Select the required policy, and in the right pane, click **Edit Policy**.
3. In the <policy name> [Edit] workspace, click the **VCD Pricing** card.
4. Click the Lock icon to override parent policy settings.
5. Select if you want to activate or deactivate the pricing engine.
6. Select the pricing policy type from the drop-down menu. The pricing policy type determines your billing model based on the Organization VDC type.
7. Configure the OVDC settings.
   - **Compute Rate**: Click the Lock icon to edit the parent policy settings. Select if you want to charge based on the Sizing Policy or based on CPU/Memory Rate.

     The following options appear if you select **CPU/Memory Rate**.

     - CPU Count Based/CPU GHz Based: Select if you want to charge the CPU rate based on GHz or CPU count.
     - Charge Period: The Charge Period indicates the frequency of charging.

- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
- Default Base Rate (per CPU count): Enter a valid number for Default Base Rate.
- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.

  **Note** This field appears when you charge the CPU based on GHz.

- Default Base Rate (per GHz): Enter a valid number for Default Base Rate.

  **Note** This field appears when you charge the CPU based on GHz.

- Other Base Rate Slabs: Using slabs, you can optionally charge different rates depending on the number of CPUs used. Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 5 in Greater than or equal and 5 as Base Rate, it means if the usage is 5 CPU and above, then the Base Rate of 5 will be applied for the whole usage.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- Fixed Cost (per charge period per VM): Enter a valid number. Fixed costs do not depend on the units of charging.

Enter details for **Memory Rate**:

- Charge Period: The Charge Period indicates the frequency of charging.
- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
- Default Base Rate (per GB): Enter a valid number for Default Base Rate.
- Base Rate Slab: Using slabs, you can optionally charge different rates depending on the memory allocated. Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.

- Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 GB and above, then the Base Rate of 10 will be applied for the whole usage.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- Charge Overage: This rule applies only to the allocation pool model, where guaranteed resources are some percent of total available resources. In this case, for the guaranteed percent, a normal base rate will be applied. If usage goes beyond the guaranteed percent, the rate mentioned in overage will be applied (for the delta usage which is higher than guaranteed).
- Overage Memory Rate: Enter an overage memory rate. This field appears only when you activate the **Charge Overage** option. If usage goes beyond the guaranteed percent, the rate entered in overage will be applied (for the delta usage which is higher than guaranteed).
- Fixed Cost (per charge period per VM): Enter a valid number. Fixed costs do not depend on the units of charging.

The following options appear if you select **Sizing Policy**. Sizing policies are a way of defining template VM sizes such as Small, Medium, and Large, in terms of CPU and Memory.

- Add: Click **Add** and enter the following details:
  - Sizing Policy Name: Select the sizing policy name from the drop-down.
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
  - Base Rate: Enter a base rate.

  The slabs that you add appear in the table below. To edit, delete, or add new slabs, click the vertical ellipses and select the desired option.

- **Storage Rate**: Click **Storage Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Click **Create Storage Policy** and enter the following details:
    - Storage Policy Name: Select a storage policy name from the drop-down list.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.

- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
- Default Base Rate (per GB): Enter a default base rate.

  **Note** The Default Base Rate option is considered only when you have selected **Aggregate storage charges from Storage profiles for payg Org-VDCs** option under **Administration > Chargeback > Pricing Settings** tab.

The storage policy that you add appears in the table below. To edit, delete, or add new slabs, click the vertical ellipses and select the desired option. Using slabs, you can optionally charge different rates depending on the storage allocated.

- **Network Rate**: Click **Network Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - **Network Data**
    - External Network Transmit (per MB): Enter the rates for external network transmit.
    - External Network Receive (per MB): Enter the rates for external network receive.
  - **Network Transmit Rate (Bandwidth)**
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Usage: Select the usage based on which you want to charge.
    - Default Base Rate (per MBps): Enter a default base rate.
    - Add Slab: Using slabs, you can optionally charge different rates depending on the network data consumed. Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 Mbps and above, then the Base Rate of 10 will be applied for the whole usage.

      The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Network Receive Rate (Bandwidth)**
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Usage: Select the Usage based on which you want to charge.
    - Default Base Rate (per MBps): Enter a default base rate.

- Add Slab: Using slabs, you can optionally charge different rates depending on the network data consumed. Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 Mbps and above, then the Base Rate of 10 will be applied for the whole usage.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Advanced Network Rate**: Click **Advanced Network Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Edge Gateway Size: Enter Base Rate, Period, and Fixed Cost for sizes Compact, Large, X-Large, and Quad Large.

    Chargeback allows you to define the size of the edge gateway and assign differential prices based on the edge size.
  - Distributed Firewall: Enter Charge Period and Base Rate (per count).
  - Edge Services: Enter Charge Period and Base Rate (per count) for IP Count, and enter Charge Period and Base Rate (if enabled) for HA, DHCP, IPV6, IP Sec, LB, NAT, SSL VPN, L2 VPN, Firewall, Static Routing, BGP Routing, and OSPF Routing.

    IP Count is the unique IP count available on the external network of the Org-VDC. Pricing can be performed based on the count of these IPs.

    Apart from the basic data transfer, there are additional value-added services offered in Cloud Director Service in combination with NSX. All the network services associated with specific edges such as HA, DHCP, IPV6, IP Sec, Load Balancer, NAT, SSL VPN, L2 VPN, Firewall, Static Routing, BGP Routing, and OSPF Routing are considered for charging based on whether these services are activated or not. If services are activated for a specific day and a base rate is applied for that service, then that particular service gets charged for that specific day. If the service is deactivated on any day then the base rate will not be applied.
  - Network Service Pricing (NSXsT Only)
    - Firewall Charges (per firewall rule count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.

    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- L2VPN charges (per L2VPN count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- Load Balancer Charges (per load balancer count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- NSX Advanced Load Balancer (Throughput): Enter Charge Period, Usage, Default Base Rate (per MBps), and click **Add Slab** to enter the following base rate slab details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    Cloud service providers can charge for consumption of NSX Advanced Load Balancer based on throughput. You can create slabs with different rates for throughput values.

    **Note** The unit of charging for 'throughput' is 'mbitsps'.

The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

**Note** The default base rate will be applied for any range of usage that is not covered in the above set of slabs.

- **Guest OS Rate**: Click **Guest OS Rate** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for any additional charges that are to be applied on Virtual machines based on their Discovered Guest Operating System by vCenter.
  - Click **Create Guest OS Rate** and enter the following details:
    - Guest OS Name: Enter a guest OS name.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

    The guest OS rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Cloud Director Availability**: Click **Cloud Director Availability** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Per Replica Charge: This section is used to set pricing for replications created from Cloud Director Availability. You can charge for each replication object, based on the SLA profile they belong to. For charging replications without any SLA Profile assigned, please enter None as the SLA Profile name.

    Click **Create Per Replication Charge** and enter the following details:
    - Replication SLA Profile Name: Enter a profile name.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Base Rate: Enter a base rate.

    The replication charges that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - Storage Usage Charge: This section is used to set additional pricing for storage used by Cloud Director Availability replications in Cloud Director. The storage usage defined in this tab will be added additionally to the Storage Policy Base Rate.

    Click **Create Storage Usage Charge** and enter the following details:
    - Storage Policy Name: Select a storage policy name from the drop-down list.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Base Rate: Enter a base rate.

The storage charges that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **vCenter Tag Rate**: Click **vCenter Tag Rate** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for any additional charges that have to be applied to the VMs based on their discovered Tags from vCenter.
  - **Fixed Rate**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Alternate Pricing Policy**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Pricing Policy: Select the alternate pricing policy name.
    - Priority: Select the priority for the alternate pricing policy. When the metadata or tag-based charges overlap, setting a priority allows you to define which policy should be processed first.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **VCD Metadata Rate**: Click **VCD Metadata Rate** in the left pane and then click the Lock icon to edit the parent policy settings. The metadata based prices are available in bills only if the Enable Metadata option is enabled in the VMware Aria Operations Management Pack for VMware Cloud Director. This section is used to charge for any additional charges that have to be applied to the VMs based on their discovered metadata from the Cloud Director.
  - **Fixed Rate**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Alternate Pricing Policy**: Click **Create Base Rate** and enter the following details:
  - Metadata Tag Key: Enter a tag key.
  - Metadata Tag Value: Enter a tag value.
  - Pricing Policy: Select the alternate pricing policy name.
  - Priority: Select the priority for the alternate pricing policy. When the metadata or tag-based charges overlap, setting a priority allows you to define which policy should be processed first.

  The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **One Time Fixed Cost**: Click **One Time Fixed Cost** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for one time incidental charges on Virtual machines, such as creation/setup charges, or charges for one-off incidents like installation of a patch. These costs do not repeat on a recurring basis.
  - VM Creation: Enter the one time fixed cost.
  - VCD Metadata: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter the one time fixed cost.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - vCenter Tag: Click **Create vCenter Tag** and enter the following details:
    - vCenter Tag Key: Enter a tag key.
    - vCenter Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter the one time fixed cost.

    The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Rate Factors**: Click **Rate Factors** in the left pane and then click the Lock icon to edit the parent policy settings. Use Rate Factors to either increase or discount the prices against individual resources consumed by the Virtual Machines or by whole charges against the Virtual Machine.
  - VCD Metadata: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.

        - Metadata Tag Value: Enter a tag value.
        - Change the price of: Select the required option from the drop-down list.
        - By applying a factor of: Enter a valid number. For example, if you want to increase the price of the CPU which has a tag 'Tag1-Value1' by 20% then select CPU from the **Change the price of** drop-down menu and enter 1.2 in **By applying a factor of**.

      The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

    - vCenter Tag: Click **Create vCenter Tag** and enter the following details:
        - vCenter Tag Key: Enter a tag key.
        - vCenter Tag Value: Enter a tag value.
        - Change the price of: Select the required option from the drop-down list.
        - By applying a factor of: Enter a valid number. For example, if you want to increase the price of the CPU which has a tag 'Tag1-Value1' by 20% then select CPU from the **Change the price of** drop-down menu and enter 1.2 in **By applying a factor of**.

      The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Tanzu Kubernetes Clusters**: Click **Tanzu Kubernetes Clusters** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for Tanzu K8s clusters and objects below an Org VDC based on certain attributes of K8s like CPU, Storage, Memory, and so on.
    - **Cluster Fixed Cost**: Enter the following details:
        - Charge Period: The Charge Period indicates the frequency of charging.
        - Fixed Cost (per count): Enter a valid number.
    - **Cluster CPU Rate**: Enter the following details:
        - Charge Period: The Charge Period indicates the frequency of charging.
        - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
        - Default Base Rate (per GHz): Enter a default base rate.
        - Base Rate Slab: Click **Add Slab** and enter the following details:
            - Greater than or equal: Enter a valid number.
            - Base Rate: Enter a base rate.
            - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

          The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Cluster Memory Rate**: Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
  - Default Base Rate (per GB): Enter a default base rate.
  - Base Rate Slab: Click **Add Slab** and enter the following details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Cluster Storage Rate**: Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Default Base Rate (per GB): Enter a default base rate.
  - Base Rate Slab: Click **Add Slab** and enter the following details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

  **Note** The default base rate will be applied to any range of usage that is not covered in the above set of slabs.

- **CSE Kubernetes Cluster**: Click **CSE Kubernetes Cluster** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for K8s clusters and objects below an Org VDC based on certain attributes of K8s like CPU, Storage, Memory, and so on.
  - **Cluster Fixed Cost**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Fixed Cost (per count): Enter a valid number.
  - **Cluster CPU Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.

- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Default Base Rate (per GHz): Enter a default base rate.
- Base Rate Slab: Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Cluster Memory Rate**: Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
  - Default Base Rate (per GB): Enter a default base rate.
  - Base Rate Slab: Click **Add Slab** and enter the following details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Cluster Storage Rate**: Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
  - Default Base Rate (per GB): Enter a default base rate.
  - Base Rate Slab: Click **Add Slab** and enter the following details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

**Note** The default base rate will be applied to any range of usage that is not covered in the above set of slabs.

- **Additional Fixed Cost**: Click **Additional Fixed Cost** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to add any other costs that have to be applied at the Org-VDC level. This can be used for charges such as overall tax, overall discounts, and so on. These can be applied to specific Org-VDCs based on Org-VDC metadata.
  - **Fixed Cost (is applied at Org-VDC level)**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Fixed Cost: Enter a valid number.
  - **VCD Metadata (Addition cost is applied at Org-VDC level)**: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Price: Enter a valid price.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **VCD Metadata One Time (Addition cost is applied at Org-VDC level)**: Click **Create VCD Metadata One Time** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter a one time fixed cost.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Organization**: Click **Cloudian Storage** in the left pane and then click the Lock icon to edit the parent policy settings. A third-party storage service that allows the cloud service providers to provide additional storage to the tenants and charge them for the same. You can charge for the consumption of cloudian storage by creating slabs with different rates for storage values.

  Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.

- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Default Base Rate (per GB): Enter a default base rate.
- Base Rate: Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

**Note** The bills for cloudian storage will be generated only at the organization level.

8 Click **Preview** to review your changes and then, click **Save**.

### What to do next

You can assign policies to the required vCenter/Cluster Compute Resource under **Configure > Policies > Policy Assignment**. For details, see Assigning Policies.

**Note** For details on configuring vCenter pricing, see Configuring vCenter Pricing Details.

## Creating Notification Rules for Tenant in VMware Aria Operations

Cloud service providers can add, manage, and edit your notification rules for tenants in VMware Aria Operations.

### Prerequisites

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances. For details on configuring outbound plug-ins, see Adding Outbound Notification Plug-Ins.

**Note** Chargeback allows you to only add a Standard Email Plug-In for outbound alerts.

### Procedure

1 To manage your notifications, from the left menu, click **Configure > Alerts**, and then in the right pane, click **VCD Tenant Notifications**. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.

2 Enter the following notification details.

| Option | Description |
| --- | --- |
| Name | Provide a name for the notification rule. |
| Description | Describe the notification rule. |
| Notification Status | Either activate or deactivate a notification setting. Deactivating a notification will stop the alert notification for that setting and activating it will activate the alert notification. |

3 Click **Next**.

4 Select if you want to send notification at the organization level or at a specific OVDC level. The list of OVDCs/Organizations is updated based on your selection.

5 From the list of OVDCs/Organizations, drag and drop the OVDC or Organization that should receive a notification, into the left pane.

**Note** You can only add those OVDCs that have email addresses and email outbound instances configured. To configure email addresses and email outbound instances, see Configuring Tenant Email in VMware Aria Operations and Add a Standard Email Plug-In for Outbound Alerts respectively.

The notifications will be sent to all selected OVDCs with configured email addresses and email outbound instances.

6 Click **Next**.

7 Define criteria for the notification rule. From the list of alert definitions, drag and drop the alert definition for which you want to trigger a notification, into the left pane.

**Note** By default, the filter Only Tenant is set to **Yes**.

The notification will be sent when ANY of the selected alert definitions triggers an alert.

8 Click **Create**.

#### What to do next

You can view the created notification from **Configure > Alerts > Tenant Notifications**.

## Configuring Tenant Email in VMware Aria Operations

After creating an email outbound, cloud service providers can assign email addresses to Org-VDCs. These email addresses receive alerts and reports related to a particular Org-VDC.

### Prerequisites

Ensure that you have created an email outbound. For details, see Add a Standard Email Plug-In for Outbound Alerts.

### Procedure

1. To configure tenant email, from the left menu, click **Configure > Alerts**, and then in the right pane, click **VCD Tenant Email Configuration**.
2. Select if you want to configure tenant email at the organization level or at a specific OVDC level. The list of OVDCs/Organizations is updated based on your selection.
3. Select the required OVDCs/Organization, and on the toolbar, click **Configure** to add or edit email configuration.
4. Enter the email address of tenant to whom the email has to be sent.
5. Select an email outbound to which you want to send the email.
6. Click **Save**.

## Managing Chargeback Reports

Cloud service providers can view, schedule, and generate reports in VMware Aria Operations. You can create report templates for your tenants.

A report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your environment.

As a cloud service provider, you can:

- View and access all report templates, for information on accessing reports, see Accessing Report Templates.
- Create report templates for tenants by selecting the **Template for Tenant** option, for more information, see Create a Report Template.
- Generate a report for a specific Org-VDC, for more information, see Generate and Regenerate a Report. You can order the reports by the date and time that they were created, the report name, the object on which the report is generated, the owner, or the status.
- Schedule a report for a specific Org-VDC, for more information, see Schedule a Report. Cloud service providers can generate tenant reports on demand or they can schedule them to be emailed to their tenants.

## Generating Bills in VMware Aria Operations

Chargeback allows you to generate monthly bills to provide an account of the overall expenses used for resources in an organization.

As a cloud service provider, you can charge different OVDCs and generate bills at the Organization/OVDC level. You can provide a bill summary at organization level by separating the costs incurred by the different OVDCs. You can also charge for resources consumed at the organization level rather than the OVDC level.

**Note** There is a service in Chargeback for automatic bill generation. On the first of every month, bills are generated for all the available organization VDCs for the previous month (1st to 30th or 31st). The generated bills are listed under **Chargeback > Bills**.

### Procedure

1 To generate a bill, from the left menu, click **Chargeback > Bills**, and then click **Generate Bill**.

2 Enter the following details.

| Option | Description |
| --- | --- |
| Name | Enter a name for the bill. |
| Description | Describe the bill. |
| Start Date | Select a start date for the bill. |
| End Date | Select an end date for the bill. |
| Advanced Settings | |
| Policy | If you want to generate a bill using a specific policy, then select a policy from the list. Only the policies that have VCD pricing configured will appear in the list.<br>It is not mandatory to select a policy from the list.<br>If you do not select any policy, then the bills will be calculated based on VCD pricing configurations present in the default policy. |

3 Click **Next** to select resources.

4 Select if you want to generate a bill at the organization level or at a specific OVDC level. The list of resources is updated based on your selection.

**Note** If you select **Organization**, then all the OVDCs under it are selected automatically. Deselecting any of the OVDCs will undo the selection at the Organization level.

5 From the list of resources, drag and drop the OVDC or Organization into the left pane.

**Note** You can use filters to filter resources by OVDC Name and Organization.

6 Click **Create**.

### Results

You can view the bill that you generated from **Chargeback > Bills**.

For bills generated at the Organization level, you can view the bill summary of the Organization and associated Organization VDCs. The Cloudian Storage section is available only when cloudian storage is configured and applied as part of the pricing policy.

## Scheduling Bill Generation Using Automation Central

You can create a schedule for bill generation using the Automation Central feature of VMware Aria Operations.

### Prerequisites

Make sure you have VMware Cloud Director configured in VMware Aria Operations to configure the pricing policy for the bill generation. For more information, see the topic, Chapter 5 Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria Operations

### Procedure

1. Click **Automation Central** in the left menu of VMware Aria Operations. Alternatively, click **Chargeback > Bills > Schedule Bills** from the three horizontal dots next to the **GENERATE BILL** button on the **Bills** page.
2. In the Automation Central page, click **Add Job**.

   The **Create New Job** page opens. This page displays the following cards:

   - Reclaim
   - Rightsize
   - Billing. This card is available when you when you have VMware Cloud Director configured in VMware Aria Operations.
   - Actions
3. Click the **Billing** card.
4. In the **Billing Information** step of the wizard, provide the following inputs and then click **Next**:

| Property | Description |
|---|---|
| **Name** | Specify a name for the job. This is displayed in the calendar. |
| **Description** | Provide a description for the job. |

5. In the **Select an Object** step of the wizard, use the toggle button to display either OVDC or Organization objects in the right hand pane. Drag the OVDC or Organization objects to the left hand pane to generate bills for those objects. You can select multiple objects. Click **Next**.

6 In the **Schedule** step of the wizard, provide the following inputs:

| Property | Description |
| --- | --- |
| **Start Date** | Set a start date which is on, or after the current date, for the job to start. Select a date from the calendar displayed on the page.<br>**Note** There will be a delay of up to five minutes for a job to run. The actual start time for an action also depends on the action itself and the number of target objects involved. |
| **Time zone** | For the start time and date to be calculated for the job, select a time zone. Choices are:<br>■ Browser - based on the current location reported by the browser.<br>■ Host - based on the current location of the host machine.<br>■ GMT based time zone - Based on a location calculated as per UTC+0. |
| **Start Time** | Select a start time for the job. The drop-down menu provides options in five minute intervals, starting closest to your current browser reported time. |
| **Recurrence** | Set a recurrence for the job. The choices are:<br>■ **Weekly**. Select the days of the week when you want the job to run, by clicking the day of the week abbreviation.<br>■ **Monthly**. Select the months when you want the job to run, by clicking the month abbreviation. By default, all the months are selected. For the months that you select, you can configure the job to run:<br>  ■ On specific days of the month, by clicking the number. Or, the last day of the month, without specifying the exact date, by clicking **Last**.<br>  ■ On the specific number of the week (first, second, third, fourth or the last) in the month, combined with the specific day of the week.<br>  You can make multiple selections for each of the options in the drop down menu.<br>For the **Daily**, **Weekly**, **Monthly** options, you can set when the job must end based on a date, or the number of occurrences. |
| **End** | Provide a date when the schedule will end. |
| **Bill Start Day** | From the drop down options, provide the start date for the billing period. |
| **Bill End Day** | From the drop down options, provide the end date for the billing period. |
| **Notifications** | In the **Notifications** section, select the **Receive Updates on Job via Email** check box to receive notifications two hours before the job is set to run and after it has been executed. For the email to be sent, you must also select the email outbound plugin from the drop down menu, and enter the email address to which the email must be sent. If you have not created an email outbound plugin, see the topic, Outbound Settings in VMware Aria Operations.<br>**Note** You can send updates via email to only one email ID. |

7 Click **Create** to complete the steps in the wizard and create the job.

### Results

The job to create a bill as per the schedule is created and added to the Automation Central calendar. When you click **Automation Central** in the left menu, you see the calendar of events. You can also switch the tabs to **History**, **Jobs** to see more information about the jobs.

**Note** The **Reports** tab does not display any information related to billing.

## Viewing Chargeback Summary

The **Chargeback Overview** page provides an overview of resources of an organization in VMware Aria Operations.

From the left menu, click **Chargeback > Overview** to view the **Chargeback Overview** page.

| Widget | Description |
| --- | --- |
| Summary | ■ Provides the number of organization VDCs, vApps, virtual machines, and number of running virtual machines that are available in an organization.<br>■ Provides total cost of the entire data center for a given OVDC, and ongoing price in US$. |
| Capacity Overview | Provides an overview of CPU, memory, and storage use of an organization. |
| Organization Details | Provides detailed information about each organization, such as the number of organization VDCs, vApps, VMs, running VMs, cost in US$ charged to the cloud service provider, and ongoing price in US$ charged to the tenant.<br>**Note** Click **Export** to export organization details as an Excel file. |

# Optimizing Capacity and Improving Performance in VMware Aria Operations

6

In VMware Aria Operations, Capacity Optimization is a key feature that focus on efficiently utilizing and managing resources within a virtualized infrastructure, such as a VMware vSphere environment. These features help organizations maintain optimal performance and maximize the utilization of their resources while ensuring that capacity remains within acceptable limits.

## Capacity Optimization

This aspect of VMware Aria Operations involves forecasting and planning for future resource needs based on historical data and current usage patterns. The goal is to ensure that the virtual environment has the necessary resources to accommodate workload growth without sacrificing performance. Capacity optimization involves activities like identifying underutilized or overprovisioned resources, predicting when additional resources will be needed, and suggesting right-sizing of virtual machines and hosts.

## Performance Improvement

Performance improvement focuses on enhancing the efficiency and responsiveness of the virtual environment. It involves monitoring the performance metrics of virtual machines, hosts, storage, and other components to detect performance bottlenecks, latency issues, and other concerns that could impact the overall performance of applications running on the virtualized infrastructure. .

Read the following topics next:

- Capacity Optimization Concepts
- How to View and Assess Capacity
- How to Optimize Capacity and Improve Performance in VMware Aria Operations
- How to Plan for Capacity Changes

# Capacity Optimization Concepts

Capacity optimization in VMware Aria Operations involves various key concepts and practices to ensure optimal resource utilization, plan for growth, and maintain performance within a virtualized environment.

## How Does Capacity Optimization Work in VMware Aria Operations

Capacity Optimization in VMware Aria Operations is achieved using powerful integrated functions - capacity overview, workload balancing and optimization, repurposing of underutilized resources, and what-if predictive scenarios - to reach optimal system performance.

Capacity planners must assess whether physical capacity is sufficient to meet current or forecasted demand. With robust capacity planning and optimization, you can manage your production capacity effectively as your organization addresses changing requirements. The objective of strategic capacity optimization is to reach an optimal level where production capabilities meet ongoing demand.

VMware Aria Operations analytics provide precise tracking, measuring and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery. The system monitors stress thresholds and alerts you before potential issues can affect performance. Multiple pre-set reports are available. You can plan capacity based on historical usage, and run what-if scenarios as your requirements expand.

### Capacity Optimization Using Overview, Reclaim, Workload Optimization and What-If

The Capacity Optimization provides four integrated functions - Overview, Reclaim, Workload Optimization, and What-If Scenarios - that give an overview of the status of all data center activity and trending. You can conduct on-the-spot analysis, including drilling down into further detail on any object to identity possible performance problems or anomalies. You can rebalance and optimize compute resources. The system further identifies underutilized workloads (virtual machines) and calculates the potential cost savings that can accrue when these resources are reclaimed to be deployed more effectively. You can interact with and manipulate data and outcomes based on your requirements.

Use the Capacity Optimization and Reclaim features to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

Workload Optimization provides for moving virtual workloads and their file systems dynamically across datastore clusters within a data center or custom data center. You can potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention triggers an alert and automatically runs an action, a data center performs at optimum.

In addition, the What-If Analysis function- can run scenarios that help determine where additional system resources can be brought online.

**Note** You may see a data center or cluster labeled as optimized when it has few or no days remaining before CPU, memory, or storage is predicted to run out. That is because these are two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

## How Does VMware Aria Operations Calculate and Forecast Capacity

Capacity analytics helps you assess the utilization and capacity remaining in objects across your environment. An evaluation of the historical utilization of resources generates a projection of the future workload. You can plan for infrastructure procurement or migrations based on the projection and avoid the risk of capacity shortage and high infrastructure costs.

Capacity analytics uses the capacity engine to assess historical trends, which include utilization peaks. The engine chooses an appropriate projection model to predict the future workload. The amount of historical data that is considered depends on the amount of historical utilization data.

### Capacity Engine and Calculations

The capacity engine analyzes historical utilization and projects future workload by using real-time predictive capacity analytics, which is based on an industry-standard statistical analysis model of demand behavior. The engine takes the Demand and Usable Capacity metrics as input and generates the output metrics, which are Time Remaining, Capacity Remaining, Recommended Size, and Recommended Total Capacity, as shown in the following figure.

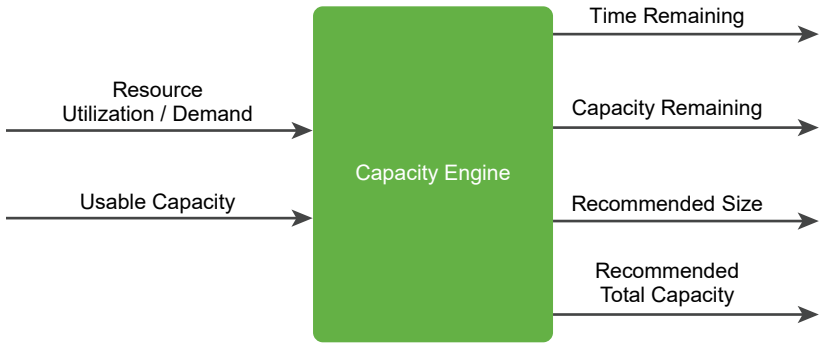


The projection window for the capacity engine is 1 year into the future. The engine consumes data points every 5 minutes to ensure real-time calculation of output metrics.

The capacity engine projects the future workload in a projected utilization range. The range includes an upper bound projection and a lower bound projection. Capacity calculations are based on the time remaining and risk level. The engine considers the upper bound projection for a conservative risk level and the mean of the upper bound projection and lower bound projection for an aggressive risk level.

The capacity engine calculates the time remaining, capacity remaining, recommended size, and recommended total capacity.

**Time Remaining**

The number of days remaining till the projected utilization crosses the threshold for the usable capacity. The usable capacity is the total capacity excluding the HA settings.

**Capacity Remaining**

The largest difference between the usable capacity and the projected utilization between now and 3 days into the future. If the projected utilization is above 100% of the usable capacity, the capacity remaining is 0.

**Recommended Size**

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. The recommended size excludes HA settings.

If the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization 150 days into the future.

VMware Aria Operations caps the recommended size that is generated by the capacity engine to keep the recommendations conservative.

- VMware Aria Operations caps an oversized recommended size at 50% of the currently allocated resources.

  For example, a virtual machine that is configured with 8 vCPUs has never used more than 10% CPU historically. Instead of recommending a reclaim of 7 vCPUs, the recommendation is capped to reclaiming 4 vCPUs.

- VMware Aria Operations caps an undersized recommended size at 100% of the currently allocated resources.

  For example, a virtual machine that is configured with 4 vCPUs has been constantly running very hot historically. Instead of recommending the addition of 8 vCPUs, the recommendation is capped at adding 4 vCPUs.

**Recommended Total Capacity**

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The recommended total capacity includes HA settings.

For example, if the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization including HA values, 150 days into the future.

**Note** Recommended total capacity is not available for objects.

The following figure shows the capacity calculations for a conservative risk level.

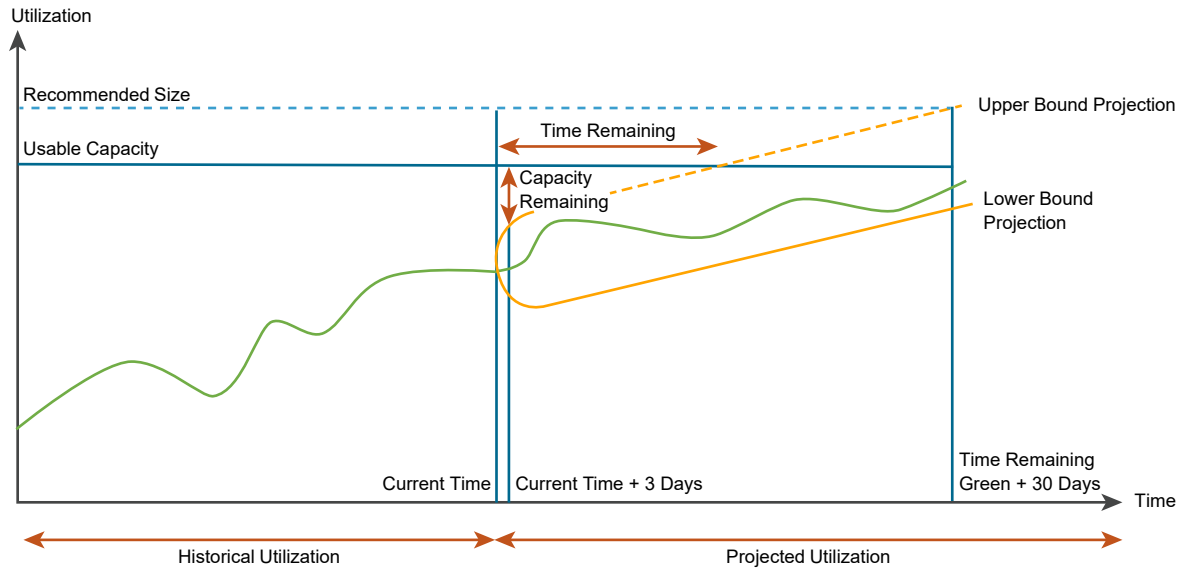


The following figure shows the capacity calculations for an aggressive risk level.

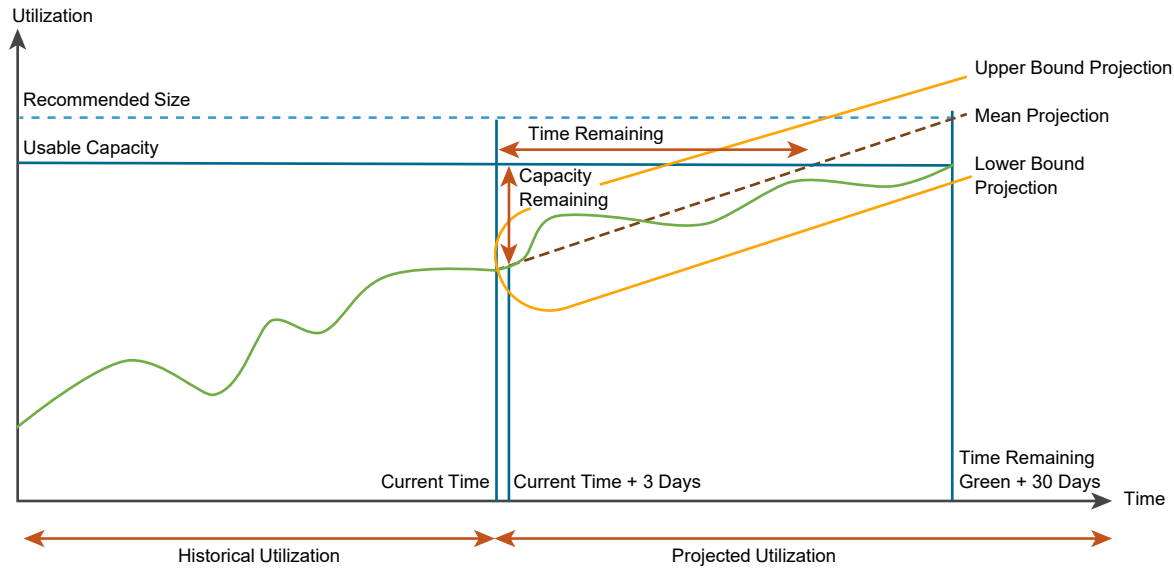


**Note**

- If HA is not enabled in VC then Usable Capacity = Total Capacity. In this case, the Usable Capacity value can be 0 only if there are no hosts in the Cluster.
- If HA is enabled, then Usable Capacity can be 0 in the following cases:
  - There are no hosts in the cluster.
  - HA is configured incorrectly. For example: it can be configured to 100% percent. Please check the HA configuration in vCenter.
  - HA Active host count is less than 2.
  - The host is not HA Active if:
    - Host is in Maintenance Mode.
    - Host is Powered Off.
    - Value of “runtime.dasHostState” property is not equal to “connectedToMaster” or “master”. This can be because of some network issues between the hosts.

## Utilization Peaks

The historical utilization of resources can have peaks, which are periods of maximum utilization. The projection of future workload depends on the types of peaks. According to the frequency of peaks, they can be momentary, sustained, or periodic.

### Momentary Peaks

Short-lived peaks that are a one-time occurrence. The peaks are not significant enough to require additional capacity, so they do not impact capacity planning and projection.

### Sustained Peaks

Peaks that last for a longer time and impact projections. If a sustained peak is not periodic, the impact on the projection lessens over time because of exponential decay.

**Periodic Peaks**

Peaks that exhibit cyclical patterns or waves. The peaks can be hourly, daily, weekly, monthly, during the last day of the month, and so on. The capacity engine also detects multiple overlapping cyclical patterns.

## Projection Models

The capacity engine uses projection models to generate projections. The engine constantly modifies projections and chooses the model that best fits the pattern of historical data. The projection range predicts the general usage pattern that covers 90% of the future data points. Projection models can be linear or periodic.

**Linear Models**

Models that have a steadily increasing or decreasing trend. Multiple linear models run in parallel and the capacity engine chooses the best model.

Examples of linear models are linear regression and autoregressive moving average (ARMA).

**Periodic Models**

Models that discover periodicity of various lengths, such as hours, days, weeks, months, or the last day of the week or month. Periodic models detect square waves that represent batch jobs and handle data streams that contain multiple overlapping periodic patterns. These models ignore random noise.

Examples of periodic models are fast Fourier transforms (FFTs), pulses (edge detection), and wavelets.

## Forecast In Trend Views

Forecasts are generated based on the time range specified in the view settings and are forecasted for the number of days specified in the forecast setting. The forecast is generated based on 3 main algorithms. Change-point detection to find sections of the history with significant changes, linear regression to find linear trends, and cyclical analysis to identify periodic patterns.

## Historical Data Window

The capacity engine captures historical data over a period of time depending on the historical data window. The historical data window that the engine uses is an exponential decay window.

The exponential decay window is a window of unlimited size in which the capacity engine gives more importance to the most recent data points. Beginning from the projection calculation start point, the engine consumes all the historical data points and weighs them exponentially, based on how far back in time they are.

# Allocation and Demand Model in Workload Optimization

VMware Aria Operations uses the demand model and allocation model for Workload Optimization. By default, only the demand model is used. You can turn the allocation model on in the active policy.

## What is Allocation Model?

The allocation model determines how much compute, memory, and storage resources are allocated to object types. You define the allocation values by modifying the policy which is applied to the objects. The allocation values, also known as overcommit ratios, affect performance and cost.

The allocation model works alongside the demand model. Unlike the demand model which always affects the capacity calculations, the allocation model can be turned on or off in the policy setting. You can control the ratio by which VMware Aria Operations overcommits either the CPU, memory, or disk space. By specifying the allocation values in the policy, you can choose whether you want to overcommit your resources or not. Overcommitting helps you measure utilization of resources in a pay-as-you-go model. When you do not overcommit, the utilization of your cluster will never exceed 100%. If your resource utilization is over the allocation ratio that you set, Capacity Remaining becomes zero.

To modify a policy and configure overcommit ratios, see Policy Allocation Model Element in *VMware Aria Operations Configuration Guide.*

## What is Demand Model?

The demand model is a concept used to predict future resource requirements based on historical data and usage patterns. The demand model is a key component of capacity planning and optimization, allowing administrators to anticipate the growth of workloads and infrastructure resources needed to support those workloads over time.

The demand model takes into account various factors, such as historical resource consumption, trends, seasonality, and potential changes in workloads. By analyzing this information, VMware Aria Operations can provide insights into when additional resources (such as CPU, memory, storage, and network capacity) will be needed to meet the expected growth in demand.

## Difference between Allocation Model and Demand Model in Workload Optimization

When clusters in a data center are not optimized, VMware Aria Operations uses the demand model by default to optimize the VMs placed in the cluster. While VMware Aria Operations moves VMs to optimize the cluster, it uses the capacity demand to calculate the best way to optimize the cluster.

When you turn on the Allocation Model in the active policy, VMware Aria Operations optimizes based on the Demand Model, and takes into consideration the overcommit ratios defined by the Allocation Model. As the VMware Aria Operations WLP engine takes into the consideration the overcommit ratios based on the allocation model for CPU, memory and disk space, these resources will not be under stress after optimization due to the over allocation. In any optimization model which is based on allocation the target is always to reduce the stress on all clusters and ensure that after the VM migration the target clusters are not over-allocated.

### How to Enable Allocation Model

From the left menu, click **Configure > Policies**. Select the active policy and click the **Capacity** card. In the Capacity Setting section, click the lock icon next to the Allocation Model for your vCenter object type. The available object types are:

- Cluster compute resource. Set the overcommit ratios for CPU, memory and disk space. You can also consider powered off VMs by selecting the **Activate** checkbox.
- Datastore. Set the overcommit ratio for the vSAN datastore disk space.
- Datastore cluster. Set the overcommit ratio for cluster disk space.

After you have set the overcommit ratios, click **Save**. After setting the overcommit ratios for the Allocation Model in the active policy, go to the Workload Placement page to see the Allocation Model displayed under the **Are your clusters meeting your utilization objective?** section.

## Policy Settings for Capacity

VMware Aria Operations lets you define capacity and workload automation policies on objects and object groups. These policies help you set symptom thresholds, alerts, and recommendations for capacity planning.

To set the capacity settings for your policy, from the left menu, click **Configure > Policies > Policy Definition**, and then click **Add** to add a policy or select the required policy.

For more information on policy settings for capacity, see the topics, Capacity Details and Workload Automation Details.

# How to View and Assess Capacity

VMware Aria Operations helps you asses capacity by predicting future resource needs based on historical data and usage patterns. VMware Aria Operations helps you analyze trends, forecast demand, and allocate resources accordingly to prevent capacity shortages and overprovisioning. To view capacity, you can open the Capacity page or view the Capacity Tab in the object browser.

## Assessing Capacity in the Capacity Page

The Capacity page in VMware Aria Operations helps you assess workload status and how much capacity is remaining in data centers across your environment.

### Where You Find the Capacity Page

Click **Capacity** under **Optimize** in the left menu. Or, from the Launchpad, click **Capacity** under Pillars of Operations. Click **View** in the **Capacity** card.

**Note** Double-click on a data center graphic to display the object details screen for the data center.

For more details on how to use the capacity page to optimize capacity, see the topic, Using the Capacity Page to Asses and Optimize Capacity.

## Viewing Object Capacity in the Capacity Tab

The **Capacity** tab provides Time Remaining and Capacity Remaining data for the selected object. Virtual Machine Remaining data is the sum of Virtual Machine remaining from each Cluster Compute Resource and is available for datastores, datastore clusters, clusters, data centers, CDC, and VC based on the average profile, or when you activate one or more custom profiles in the policy.

### Where You Find the Capacity Tab at an Object Level

From the left menu, click **Environment**, then select a group, custom data center, application, or inventory object. The Object details screen appears. Click the **Capacity** tab.

### Understanding the Capacity Tab

For the selected object, the **Capacity** tab lists two panes with the Time Remaining and Capacity information. These panes display the value of the resources remaining till they run out.

Below the **Time Remaining**, **Capacity** and **Virtual Machine Remaining** panes, the time and capacity utilization metric for CPU, memory, and disk space are displayed in three panes. By default, the most constrained resource is selected. Click **CPU**, **Memory**, or **Disk Space** to change the views to these resources. These panes display the resource information based on the Demand model (default) or Allocation model (if configured).

For datastores, if you have activated Allocation Model and Capacity buffer in the assigned profile, you see the **Disk Space** information based on Allocation and Usage.

#### Time Remaining Pane

When you select the **Time Remaining** pane and click one of the resource types, the utilization graph displays the historical value of the utilization metric and its forecast plotted against time, projecting how swiftly resource utilization is approaching the usable capacity.

Click **RESET** in the **Time Remaining** pane if you want to change the date from when the historical utilization is calculated. By default, it is calculated from the object creation or VMware Aria Operations installation date. Click **RESET** if you want to change projected utilization, for example drop the irrelevant historical data from the calculation.

This affects the capacity calculation for the future trend, which will impact Capacity Remaining, Time Remaining, Recommended Size, and VM Remaining (if it is available for the particular resource, namely, Datastore, Cluster, Datastore Cluster). It will also impact all the resource containers (CPU, Memory, Disk Space, etc.)

For example, if the capacity calculation is using historical data based on the duration when you were provisioning a large number of VMs, but in the recent past, there was no provisioning done. Then, VMware Aria Operations may not project a trend that is based on the recent data when there was no provisioning, but still will consider historical data when there were a large number of VMs provisioned. This historical data may show an increasing trend. In this case, you can change the capacity computation to start from a date after the VM provisioning is complete.

**Note** If you are unable to see the **RESET** button in the **Time Remaining** pane, make sure you have the **Manage Capacity Calculation** permission under **Administration > Access Control > Roles > Optimize > Capacity**

### Capacity Pane

The **Capacity Remaining** pane indicates the unused capacity of your virtual environment to accommodate new virtual machines. VMware Aria Operations calculates the Capacity Remaining as a percentage of the remaining capacity, compared to the total. Capacity Remaining is calculated as the utilization metric forecast 3 days from now subtracted from the Usable Capacity. VMware Aria Operations calculates the average profile and always computes the virtual machine remaining number based on the average profile. You can change the profile by clicking the + icon above the bar chart. VMware Aria Operations calculates virtual machine remaining numbers when you activate one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.

When you select Capacity and click one of the resource types, a bar chart and a table of values based on the Demand and Allocation model (if configured) appears. The bar chart displays total usable resource, the percentage used, the percentage allocated for high availability and buffer, and the percentage remaining based on the Demand and Allocation models (if configured).

The table displays the following information for each resource type:

- **Total**: The total usable capacity for each resource type based on the Demand model or Allocation model (if configured). The difference in Total capacity and Usable capacity is set in the HA (admission control) that is set in the clusters in vSphere.
- **Usable**: The total usable capacity for each resource type based on the Demand model or Allocation model (if configured).
- **Used**: Approximate value how much utilization do you have now. Shows the forecast value of utilization metric in 3 days from now.

- **Recommended Size**: The Total Capacity that must be available for a green level of Time Remaining. The slider in the policy controls the Time Remaining green zone, and the default value is 150 days.
- **Remaining**: The Capacity Remaining metric value and also the percentage. The value of Capacity Remaining metric is calculated by forecasting the utilization metric 3 days from now and subtracting it from Usable capacity.
- **Buffer**: The percentage of the capacity buffer based on the buffer value that you set in the policy. The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required.
- **High Availability**: The percentage of the high availability based on the high availability buffer.

**Virtual Machine Remaining**

The virtual machine remaining number is based on the average profile. The virtual machine remaining numbers are calculated when you activate one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.

When you click on Virual Machine Remaining, you see the number of virtual machines based on the Average Profile. Select a different policy by clicking **Assign Policy** from the drop-down list on the top right hand side of the page. To activate custom profiles in that policy, or other policies, click the **GO TO CUSTOM PROFILES** link. For more information, see the topic, *Custom Profiles in VMware Aria Operations*.

Only vSAN datastore objects, and not other datastore objects, display disk space based on Buffer and HA.

The **Capacity** tab is a subset of the Capacity optimization capability. For additional details, refer to Assessing Capacity in the Capacity Page.

## Capacity VM Shortfall

VMware Aria Operations has a new metric known as VM shortfall, the VM shortfall value is always positive. The metric counts all the negative VMs Remaining and then turns them into positive. The VM shortfall metric will be available for Clusters, Datastores, and Datastore Cluster objects only. The VM Shortfall metric details is displayed in the Capacity Summary pages next to VM Remaining, in case if VM Remaining is equal to zero.

For all these objects you can create custom profiles directly from the capacity page. In the **Capacity** tab, click the + sign under Virtual Machine remaining. From the **Applicable Profiles** page, click **Add** to add a new profile. You can either add new profile or import from existing object. You can modify the required metric value and click **OK**.

## How Does VM Shortfall Work

When the cluster has more VMs than its usable capacity can handle, the shortfall of capacity is shown by the VM counter. VM remaining will show zero and not negative. The shortfall appears only when the usable capacity is zero.

Whenever there is no capacity in that case, VM shortfall detects the VM profiles in the setup and finds the profile for which we have the maximum number of VMs that are overcommitted.

VM Shortfall Formula

It is possible that when one datastore is connected with multiple clusters, each cluster thinks that the whole diskspace available in the datastore is for itself. We can calculate the VM shortfall as:

```
VM Shortfall = (Used - Usable Capacity) / Profile Sizes
```

**Note** In the above formula the Used Capacity represents the specific value that is displayed in the **Capacity > Capacity Remaining** tab.

If you have two clusters, one cluster has 20 VMs remaining and the other cluster has shortfall of 20 VMs, then the number of VMs remaining for the datacenter will be 20 VMs, however, if you have a shortfall of VMs in the datastore connected to that cluster, then we will see on the datacenter the remaining VMs as zero.

Metrics

The VM shortfall displays metrics under Capacity Analytics Generated and under Profiles. The metrics are:

- Capacity Remaining Profile (Average)
- Capacity Shortfall Profile (Average)

## vSAN HCI Mesh

The vSAN HCI Mesh allows vSAN clusters to remotely mount the datastore of another (remote) vSAN cluster, hence sharing the storage capacity and span its usage to a wider pool of compute resources. HCI Mesh allows multiple vSAN clusters to share their datastores remotely. If the vCenter version is 7.0U1 and greater you can share other cluster datastores and provision VMs to the remotely shared data stores.

A new card is created under **Optimize > Capacity > vSAN HCI Mesh**. The vSAN HCI Mesh card is displayed only when there are eligible clusters which can be part of the vSAN HCI mesh. If the version of the vCenter is less than 7.0U1 for the selected datacenter, this card is not displayed. If the mesh is not configured, click the link, which lets you configure the clusters from vCenter.

The vSAN HCI Mesh displays the following information:

- Total Capacity - Displays the total capacity of all the vSAN datastores connected to the mesh.
- vSAN Mesh - Displays how the datastores are interconnected in the HCI mesh.
- Inner Circle - Represents the local datastore and displays capacity details like, free capacity, used capacity, and total capacity.
- Outer Circle - Represents the remote datastore and displays capacity details like, total mounted capacity, free capacity, and total capacity usage percentage.

- Arrow Heads - The arrow heads points to the clients connected to the server vSAN cluster.

**Datastore View**

Displays the capacity consumption details for the selected datastore. You can also identify which are the client and server clusters.

**Cluster View**

Displays the capacity details of the remote datastores. You can also view the capacity details for the mounted datastores. The Total Usage % displays the total capacity consumed by all the client clusters connected to these mounted datastores including the selected cluster.

**Health Alerts Related to HCI Mesh**

Displays the health alerts related to the HCI Mesh. You can also view alerts for individual datastores, which is part of the HCI mesh.

## Custom Profiles in VMware Aria Operations

A custom profile defines a hypothetical configuration of a virtual machine. Custom profiles help you determine how many instances of that virtual machine can fit in your environment, depending on the capacity remaining and the configuration of the parent object. You activate the custom profile in policies for specific object types. The custom profiles are used to calculate VM remaining when the policy is selected for those object types in the **Capacity** tab of the VMware Aria Operations object browser.

### How Custom Profiles Work

As with default profiles, custom profiles is based on metric values for a virtual machine. You can add as many custom profiles as you need for an object. For example, you might create one custom profile for a virtual machine that has a memory demand model of 2 GB. You create another custom profile that has a memory demand model of 4 GB.

VMware Aria Operations uses custom profiles to calculate the number of instances of those virtual machines that can fit in your environment for specific objects. The number of virtual machines is based on the capacity allocation and demand defined for the object.

After you define the metric values for a custom profile, you configure for which policies and object types the custom profiles are activated.

### How Custom Profiles are Related to Policies

After you create a custom profile, you activate it in a policy for one or more of the following object types:

- Datastores
- Datastore clusters
- Cluster compute resource

The custom profile is then available for all objects of that object type in the **Capacity** tab of the VMware Aria Operations object browser.

To determine how many instances of the virtual machine with a specific configuration as defined by the custom profile you can include in the parent object, go to the **Capacity** tab of the parent object in **Environment > Object Browser**. Select the **Capacity Remaining | Virtual Machine Remaining** pane. Select the policy in which the custom profile is activated, from the top right hand side drop-down list. The custom profiles appear in the VM remaining section and indicate how many instances of the virtual machine can fit in your environment. The overall virtual machine remaining is based on the most constrained profile.

### Where You Find Custom Profiles

In the left menu, click **Configure > Custom Profiles**. A list of all the custom profiles available in VMware Aria Operations is displayed in a table. Click the name of a custom profile from the **Name** column to view the metric values and policies of that custom proile on the right pane.

Table 6-1. Options in the Custom Profiles Page

| Option | Description |
| --- | --- |
| Toolbar options | In the toolbar click **Add** to add a custom profile for a specific object type. Click the **Vertical Ellipses** against a profile to perform the following actions:<br>■ **Delete**. Delete the selected profile.<br>■ **Export**. Export the selected profile.<br>■ **Import**. Import selected profiles. |
| Sorting options | Filter the list to display profiles that match the filter you create. You can sort by name, description, object type, or adapter type. |
| Filter and Search | Use the advanced filter and search box to look for custom profiles. Click the drop-down icon to search by name, description, object type, adapter type. |

## Custom Profiles Add and Edit Workspace

You can add a custom profile for a cluster compute resource, datastore cluster, or datastore, to determine how many VMs of the specific object type can fit in your environment. In the Custom Profiles workspace, you create a custom profile for the object and define the metrics for the virtual machines.

### Where You Create or Edit a Custom Profile

To create a custom profile, click **Configure > Custom Profiles** in the left pane. The Custom Profiles page displays all the available custom profiles in a table. You have the option to add custom profiles via the **ADD** button, or filter and search for profiles using the search box. To see the properties of a custom profile, click the custom profile name in the table. The custom profile details are displayed on the right hand side pane.

Click the **Vertical Ellipses** next to a profile name to edit, delete, clone, or export the profile. Alternatively, to delete or export the more than one profile, select the checkbox next to the profile name and click the **Vertical Ellipses** next to the **ADD** button, and perform an action. You can also import custom profiles after clicking the **Vertical Ellipses** next to the **ADD** button.

### Add or Edit Custom Profiles Wizard

When you click the **ADD** button, or **Edit** next to the **Vertical Ellipses**, you get a three step wizard. Click the **Next** button after every step to proceed with the wizard.

#### Table 6-2. Profile Step

| Option | Description |
| --- | --- |
| Name | Descriptive name of the custom profile. This field is mandatory. |
| Description | Meaningful description for the custom profile. Provide specific information that other users must know about this profile. |
| Profile Type | The profile type for which the custom profile is created. Currently, you can create a custom profile only for a virtual machine. The Capacity Remaining is calculated for this object type. |

#### Table 6-3. Metrics Step

| Option | Description |
| --- | --- |
| Define Values for Profile | Populate the value and unit for the capacity metrics. You can also copy the metric values for the custom profile from an existing object. When you click the **Copy from Existing Object** button, the **Select an Object** pane opens on the right hand side. Browse to the object whose metric values you want to copy, or use the advanced search and filter to search by VM Name, vCenter, VM Tag, or Custom Group. Click **Copy** when done. |

#### Table 6-4. Policies Step

| Option | Description |
| --- | --- |
| Checkbox | Select the checkbox against the name of the policy which is displayed in rows for one or more of the following types of objects:<br>■ Cluster Compute Resource<br>■ Datastore<br>■ Datastore Cluster |

Click **CREATE** to create the custom profile or **UPDATE** if you are updating an existing custom profile.

### Clone a Custom Profile

Click the **Vertical Ellipses** next to a profile name to clone an existing profile profile. When you click **Clone**, you get a three step wizard. In the **Profile** step of the wizard, you can provide the name, description, and select the profile type. In the **Metrics** and **Policies** steps of the wizard, the details from the custom profile that you are cloning from, are already populated. You can skip these steps and click the **CREATE** button to create a cloned custom profile with the same values and policies as the original custom profile.

### Custom Profiles Import and Export

Export a custom profile as an XML file and import it again in a different VMware Aria Operations instance to easily transfer custom profile configurations.

#### Where you Import or Export Custom Profiles

In the left navigation menu, go to **Configure** > **Custom Profiles**. Do one of the following:

- To import a custom profile, click the Vertical Ellipses next to the **ADD** button and select Import. The **Import Custom Profile** dialog box opens. Browse to the XML configuration file and select what VMware Aria Operations has to do in case of a conflict. Click **IMPORT** to complete the process.
- To export an existing custom profile, select the box next to the custom profile name. Click the Vertical Ellipses next to the name and select **EXPORT**. The XML configuration file is downloaded to you local system via the web browser.

## Predefined Dashboards for Capacity

VMware Aria Operations provides predefined dashboards to help you monitor and manage your virtualized environments. Capacity dashboards help you analyze the capacity of your environment by providing insights into resource consumption trends, utilization patterns, and potential bottlenecks. They can help you plan for future growth and optimize resource allocation.

For more information on the predefined dashboards available for capacity optimization in VMware Aria Operations, see the topic, Capacity Dashboards.

# How to Optimize Capacity and Improve Performance in VMware Aria Operations

VMware Aria Operations helps you optimize capacity by identifying and reclaiming unused or underutilized resources, such as reclaiming idle CPU or memory from VMs. This practice can help free up resources for other workloads. Rightsizing helps you adjust the resources allocated to virtual machines based on their actual usage. vVMware Aria Operations identifies overprovisioned or underutilized VMs and provides recommendations to optimize their resource allocations.

## Using the Capacity Page to Asses and Optimize Capacity

The capacity page in VMware Aria Operations offers a comprehensive view of resource utilization, allocation, trends, and forecasts, allowing administrators to make informed decisions about resource provisioning, performance optimization, and future growth.

To open the Capacity page, click **Capacity** under **Optimize** in the left menu. Or, from the Launchpad, click **Capacity** under Pillars of Operations. Click **View** in the **Capacity** card.

## How Optimize Capacity Works

The Capacity Optimization and Reclaim features are tightly integrated functions that allows you to assess workload status in data centers across your environment. You can determine time remaining until CPU, memory, or disk space resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the Capacity page, graphical representations of all the data centers and custom data centers in your environment appear. VMware Cloud on AWS data centers has a unique icon to differentiate it from the other data centers.

By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The page refreshes to display the following data:

### Time Remaining

Time Remaining specifies which clusters are most constrained and displays the criticality of the cluster.

### Optimization Recommendations

VMware Aria Operations shows you the number of reclaimable VMs and the associated cost savings. Click **View Reclaimable VMs** to navigate to the **Reclaim** page.

### Cluster Utilization

Cluster Utilization displays an interactive graph that shows time remaining by component. You can explore the demand percentage over time by CPU, memory, and disk space or by the most constrained component. By default, the data displayed is for the Demand model. If you have configured the Allocation model, then you can also see the CPU, memory, and disk space time remaining model based on the overcommit ratios that you have set in the policy.

Click the **Edit** icon to modify the criticality threshold, risk level, and allocation model. These changes affect the selected cluster's policy. Hence, any change that you make here, affects all the clusters under the same policy.

Set the **Show History** and **Show Forecast** variables to create the slice of time in which you want to see time-remaining data. The vertical axis of the graph shows the total capacity being used by the current amount of CPU, memory, or disk space respectively. The bold, black line across the top of the graph depicts the historical value of usable capacity. The horizontal axis is the timeline. Vertical lines in the graph are labeled at the bottom of each line. The first vertical dotted line on the left marks the projection calculation start point. The next line is the current date - now. The third vertical marks the date the resource runs out. If a resource has little time remaining, the current date and the date that time runs out may be the same.

VMware Aria Operations can make recommendations for increasing time remaining based on the data it receives and these recommendations appear at the bottom of the screen. You might see two options: Option 1 shows what you can achieve by reclaiming resources. Option 2 shows the results of adding capacity.

If you choose to reclaim resources, you can run that process immediately by clicking **RECLAIM RESOURCES**. To see the details or choose additional options before running a reclaim action, review the information provided in the **Optimization Recommendations** pane and then click **VIEW RECLAIMABLE VMS** to go to the **Reclaim** page.

Table 6-5. Capacity Optimization Options

| Option | Description |
| --- | --- |
| **Select a datacenter** | Select a data center from the carousel across the top of the page. Information about the datacenter is displayed below. |
| **ALL DATACENTERS \| X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. This option appears if you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. This option appears if you select **ALL DATACENTERS** on the upper right. |
| Sort by: | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>■ Alarm clock graphic - lists data centers/custom data centers by time remaining.<br>■ Dollar sign - lists data centers/custom data centers by potential cost savings.<br>■ Scales graphic - lists data centers/custom data centers by level of optimization. |
| **Select datacenter or ADD NEW CUSTOM DATACENTER** | Options (options appear if you select **ALL DATACENTERS** on the upper right):<br>■ Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object.<br>■ Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| Time Remaining | Appears when you select a data center or custom data center from the top of the screen.<br>Gives overview of cluster status, including how many are at:<br>■ Critical<br>■ Medium<br>■ Normal<br>■ Unknown<br>"Critical" can indicate a resource contention, imbalance, or other stress condition. Thresholds you set in the policies define what is critical. |
| Optimization Recommendations | Lists potential cost savings by reclaiming unused resources.<br>Indicates if workloads can be optimized across clusters.<br>**VIEW RECLAIMABLE VMS** - displays the **Reclaim** screen, where you can research and run potential VM reclamation actions.<br>**VIEW OPTIMIZATION** - displays the **Workload Optimization** screen, where you can optimize workloads based on your policy settings. |

Table 6-5. Capacity Optimization Options (continued)

| Option | Description |
| --- | --- |
| **Cluster Utilization and Time Remaining** | Overall view of cluster health in the selected data center. You can select a cluster from the list to display information about that cluster, or use the options to sort and filter results. The options you select dictate the data displayed in the graph.<br>Sort by:<br>■ Most Constrained: most constrained element<br>■ CPU (allocation or demand)<br>■ Memory (allocation or demand)<br>■ Disk Space (allocation or demand)<br>**Note** Demand model is always on and is the default.<br>**Filter**: search field.<br>**Show History for**: The period before forecasting begins (does not impact the forecast calculation).<br>**Show Forecast For**: The forecast period.<br>**How is the criticality determined?** Displays the criticality threshold you set for this type of object in the Policies Library.<br>**Cluster Time Remaining Settings**: Click the **Edit** icon to edit the default policy for the selected cluster. Change the criticality threshold, risk level, allocation model and capacity buffer. Applying these changes affects all objects in the policy. For more information, see Configuring Policies in the *SaaS Configuration Guide* |
| Time Remaining graph | Data shows current and trending resource usage and pinpoints when a given cluster is projected to run out of CPU, memory, or disk space based on the allocation or demand model (default). VMware Aria Operations takes into account the business hours which you set in the current policy. |
| Legend | Displays a legend of the colours in the time remaining graph. Click on any one of the legends to toggle its display in the chart. |
| Recommendations | Option 1: Reclaim Resources.<br>Shows resources that can be reclaimed to increase time remaining for the selected cluster.<br>**RECLAIM RESOURCES** - displays the **Reclaim** screen, where you can research and run potential VM reclamation actions.<br>Option 2: Add Capacity.<br>Shows resources that can be added to increase time remaining. |

**Note** You might see that a data center or cluster is labeled optimized when it has few or no days remaining before CPU, memory, or disk space is predicted to run out. The seemingly odd assessment is due to optimization and time remaining being two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

# Using Reclaim to Free Up Resources

Use the **Reclaim** feature in VMware Aria Operations to identify underutilized workloads and reclaim resources from across your environment.

## Where You Find Reclaim

Select **Reclaim** under **Optimize** in the left pane. From the Launchpad, select the **Capacity** card under Pillars of Operations and click **View** in the **Reclaim** card.

**Note** Double-click on a data center graphic to display the object details screen for the data center.

## How Reclaim Works

The Capacity Optimization and Reclaim features are tightly integrated functions that allows you to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the **Reclaim** page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The area following refreshes to display details about the selected data center. The **How much you can potentially save** pane reflects potential capacity savings and indicates a possible cost savings once you have reclaimed underused or powered off VMs. The **Total Reclaimable Capacity** pane gives details of the reclaimable percentages for CPU, memory, and disk space.

The table at the bottom of the page provides important information about the VMs that offer the most cost savings. The VMs are listed by **Powered Off VMs**, **Idle VMs**, **Snapshots**, and **Orphaned Disks**. The highest priority heading is at the far left. You can specify what information is included in your reclaim action. For example, when you click a column heading, the table lists, by data center and then by VM, the allocated and reclaimable CPUs and memory, respectively. Then, for example, you can select the box next to one or more VM names and click the **EXCLUDE VM(S)** button to keep those VMs from being included in any reclaim action. You can also select VMs to resize.

## Reclamation Settings

Select the gear icon next to the page heading to customize Reclamation Settings. This affects all data centers. Using the Reclamation Settings, you can exclude, for example all snapshots from being included in the reclaim action - by deselecting the Snapshots check box. Similarly, you can include or exclude powered-off VMs, idle VMs, and orphaned disks. For more information, see Reclamation Settings.

**Note** To provide read-only access to the Reclamation Settings page for a user, configure the user role in the Access Control page (Roles tab) under **Administration > Access > Access Control**. Select the **Manage Global Settings** permissions under **Administration > Management** in the **Permissions** pane to grant access to modify the Reclamation Settings page. Unselect the **Manage Global Settings** permissions to grant read-only access.

## Run a Reclaim Action

Run a reclaim action as follows:

1 In the table headings, **Select** the types of VMs to reclaim.

2 **Click** the name of a listed cluster to show its VM list.

3 **Select** each VM or snapshot you want to reclaim.

4 Click **Delete VM(s)** to reclaim their resources.

### Table 6-6. Reclaim Options

| Option | Description |
|---|---|
| Select a data center. | Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. |
| **ALL DATACENTERS** \| **X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Sort by: | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>■ Alarm clock graphic - list data centers/custom data centers by time remaining.<br>■ Dollar sign - list data centers/custom data centers by potential cost savings.<br>■ Scales graphic - list data centers/custom data centers by level of optimization. |

Table 6-6. Reclaim Options (continued)

| Option | Description |
| --- | --- |
| **Select data center or ADD NEW CUSTOM DATACENTER.** | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>■ Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.<br>■ Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| How much you can potentially save. | Appears when you select a data center or custom data center from the top of the screen. Shows the total calculated potential cost savings when you accept system reclamation recommendations. |
| Total Reclaimable Capacity | Lists potential cost savings for the selected data center when you reclaim unused resources.<br>Resource: CPU, memory, or disk space<br>Reclaimable Capacity: how much capacity is available to reclaim from idle resources<br>% Reclaimable: percentage of total CPU, memory, or storage you can reclaim. |

### Table 6-6. Reclaim Options (continued)

| Option | Description |
| --- | --- |
| Duration older than: | Shows idle or powered off VMs that have been idle or powered off for at least the selected time period: one week, two weeks, or a month. |
| Table of Potential Cost Savings | Tabular representation of the VMs, Idle VMs, Snapshots, and Orphaned disks in the selected data center from which resources can be reclaimed.<br>Click one of the elements - powered off VMs, idle VMs, and so on - to refresh the table with data for that element. The table lists the relevant clusters. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.<br>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.<br>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.<br>**Exclude VM(s)**: The selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.<br>For powered Off VMs:<br>■ **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for powered off VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>■ **DELETE VM(s)**: Deletes the selected VMs.<br>■ **EXCLUDE VM(s)**: Excludes the selected VMs.<br>■ **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br>For idle VMs:<br>■ **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for idle VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>■ **DELETE VM(s)**: Deletes the selected VMs.<br>■ **POWER OFF**: Powers off the selected VMs.<br>■ **EXCLUDE VM(s)**: Excludes the selected VMs.<br>■ **EXPORT ALL**: Exports the list of idle VMs into a CSV file.<br>For Snapshots:<br>■ **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for snapshots. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>■ **DELETE SNAPSHOT(s)**: Deletes the selected snapshots.<br>■ **EXCLUDE VM(s)**: Excludes the selected snapshot.<br>■ **EXPORT ALL**: Exports the list of snapshots into a CSV file. |

Table 6-6. Reclaim Options (continued)

| Option | Description |
| --- | --- |
| | **SHOW\|HIDE EXCLUDED VMS**: Toggle displays or hides the list of VMs you previously excluded.<br>**Note** By default, calculations for reclaimable resources are based on the demand model. But if you turn on the allocation model in the policy settings, the calculations are based on the allocation model.<br>For Orphaned Disks:<br>■ **EXCLUDE DISK(S)**: Exclude the selected disks in the actionable list.<br>■ **EXPORT ALL**: Exports the list of orphaned disks into a CSV file. You cannot reclaim orphaned disks from the UI. Instead, export the list into a CSV file and then reclaim the orphaned disks manually.<br>**Note** VMware Aria Operations reports orphaned VMDKs conservatively. There might be a false positive situation when the used VMDK is reported as orphaned, particularly if the VMDK is located on a datastore which is shared among multiple VCs, while not all the VCs are monitored by VMware Aria Operations.<br>Check the accuracy of the VMDK reported as an orphaned disk, and then perform a reclamation.<br>**SHOW\|HIDE EXCLUDED DISKS**: Toggle displays or hides the list of disks you previously excluded. Excluded disks are not listed in the exported CSV file. |

## Reclamation Settings

Displays information about powered off VMs, idle VMs, snapshots and orphaned disks. This information helps to identify the amount of resources that can be reclaimed and provisioned to other objects in your environment or amount of potential savings that can be done in each month.

The types of VMs are ranked in the order of their importance in a reclamation action. A VM whose attributes match more than one VM type is included with the higher-ranking VM type. Grouping the VMs this way eliminates duplicates during calculations. As an example, powered-off VMs are ranked higher than snapshots, so that a powered-off VM that also has a snapshot appears only in the powered-off VM group.

If you exclude a given type of VM, all VMs matching this type are included with the next lower-ranked group they match. For example, to list all snapshots regardless of whether their corresponding VMs are powered-off or idle, deselect the Powered-off VMs and Idle VMs check boxes.

Further, you can configure how long a given class of VMs must be in the designated state - powered-off, for example, or idle - to be included in the reclamation exercise. You also can choose to hide the cost savings calculation.

| Property | Description |
| --- | --- |
| Show Cost Savings | Controls whether to show Cost savings in the **Optimize > Capacity** page. |
| VM Idleness Criteria | Defines the criteria to identify VMs that remain idle during the defined period of time.<br>You can set the following values:<br>■ MHz of CPU consumption of VMs, and percentage of the time VMs have less than defined MHz of CPU within each day. |
| Show the following types of VMs for reclamation | Configure the types of VMs that are shown for reclamation. You can select the type of VMs which are shown for reclamation, and then configure paramters which those VMs have to meet to be eligible for reclamation. |
| Powerd Off VMs | VMs that have been continuously powered off during the defined period of time.<br>The total storage capacity used is reclaimable. Total storage reclaimable cost is computed by multiplying storage rate with storage utilization. The direct cost of VM is also attributed. |
| Idle VMs | You can configure the following parameters based on which VMware Aria Operations calculates idle VMs:<br>■ How many days the VM has been idle for.<br>■ Number of days before which the VMs were provisioned based on which they are excluded. |
| Snapshots | VM snapshots that have existed for the entire defined period of time.<br>Snapshots of a VM use storage space and such storage is reclaimable. The reclaimable cost is computed by multiplying storage rate with reclaimable storage value. |
| Orphaned Disks | VMDKs on datastores that are not connected to any registered VMs and have not been modified during the defined period of time.<br>Orphaned disks are VMDKs which are associated with a VM which are not in inventory, but still available in a datastore. You can configure the minimum number of days for which VMDKs not related to any existing VM will be reported as orphaned and appear under Orphaned Disks in Reclaim page.<br>**Note** You can navigate to**Cost/Price** section under **Administration > Global Settings** in the left menu, and change the value of the **Orphaned Disks Collection** time. At this time that you set, VMware Aria Operations checks for orphaned VMDKs in vSphere Client instances. The settings for **Cost Calculation** and **Orphaned Disks Collection** are interrelated. The default value for **Cost Calculation** is 9:00 PM, and the default for **Orphaned Disks Collection** is 8:00 PM. It is recommended to schedule **Cost Calculation** after **Orphaned Disks Collection**. |

**Note** If you are unable to make changes in the Reclamation Settings page, your user role in the Access Control page (Roles tab) under **Administration > Access Control** must be modified by an administrator. The **Manage Global Settings** permissions under **Administration > Global Settings** in the **Permissions** pane controls access to the Reclamation Settings page.

## Using Rightsize to Adjust Resource Allocation

Rightsizing in VMware Aria Operations refers to the process of adjusting the resource allocations of virtual machines (VMs) to match their actual workload requirements. It involves optimizing the allocation of CPU, memory, storage, and other resources to ensure that VMs are neither overprovisioned (allocating more resources than necessary) nor underutilized (allocating fewer resources than needed). The goal of rightsizing is to achieve efficient resource utilization, improve performance, and potentially reduce costs Use the Rightsize page to alter the number of CPUs and amount of memory in oversized and undersized virtual machines.

### Where You Find Rightsizing

Select **Rightsize** under **Optimize** in the left pane.

**Note** Click on a data center graphic to display the details for the data center.

### How Rightsizing Works

The Capacity Optimization, Reclaim, and Rightsizing features are tightly integrated functions that allows you to assess workload status and resource usage in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed. With this function, you can change CPU size and memory values for oversized and undersized virtual machines to achieve optimum system performance.

When you open the page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To identify possible oversized and undersized VMs in a data center, click its graphic. The area following refreshes to display details about the selected data center.

"Oversized VMs" displays the number of VMs determined to be oversized based on policies previously set. A chart details suggested reductions in the overall number of CPUs and GBs of memory and shows the percentage of total resources the reductions represent. Similarly, "Undersized VMs" indicates the number of VMs considered to be undersized, with a chart listing suggested increases in CPU and memory.

The table at the bottom of the page provides important information about the VMs. Table headings are Oversized VMs and Undersized VMs. VMs under each heading are grouped by cluster. Click the chevron to the left of a cluster name to list all the oversized or undersized VMs, respectively, in that cluster. You can check the box next to one or more VM names and click the **EXCLUDE VM(S)** button to prevent those VMs from being included in a resizing action. You can also select individual VMs to resize before clicking the **RESIZE VM(S)** button.

## Run a Rightsize Action on Oversized VMs

Run the action as follows:

1 In the table headings, **Select** Oversized VMs.

2 **Select** the boxes next to VMs you want to exclude from the action, if any.

3 Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.

4 **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.

5 Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested reductions for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.

> **Note** Operational Actions must be activated in the vCenter cloud adapter instance.

6 **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

## Run a Rightsize Action on Undersized VMs

Run the action as follows:

1 In the table headings, **Select** Undersized VMs.

2 **Select** the boxes next to VMs you want to exclude from the action, if any.

3 Click **EXCLUDE VM(S)**, if required. In the confirmation dialog box, click **EXCLUDE VM(S)**.

4 **Select** the boxes next to VMs you want to include in the resizing action, or **Select** the box next to VM Name to include all VMs.

5 Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested increases for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.

6 **Select** the box at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

### Table 6-7. Rightsize Options

| Option | Description |
|---|---|
| Select a data center. | Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. |
| **ALL DATACENTERS** \| **X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select **ALL DATACENTERS** on the upper right. |

Table 6-7. Rightsize Options (continued)

| Option | Description |
| --- | --- |
| Sort by: | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>■ Alarm clock graphic - list data centers/custom data centers by time remaining.<br>■ Dollar sign - list data centers/custom data centers by potential cost savings.<br>■ Scales graphic - list data centers/custom data centers by level of optimization. |
| **Select data center or ADD NEW CUSTOM DATACENTER.** | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>■ Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.<br>■ Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| Oversized VMs display | Displays the number of VMs identified as oversized, with suggested reductions for vCPU and memory size. |

Table 6-7. Rightsize Options (continued)

| Option | Description |
| --- | --- |
| Undersized VMs display | Displays the number of VMs identified as undersized, with suggested increases for vCPU and memory size. |
| Table of Oversized and Undersized VMs | Tabular representation of the Oversized and Undersized VMs in the selected data center.<br>Click one of the headings - Oversized VMs or Undersized VMs - to refresh the table with data for that heading. The table lists the relevant VMs. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.<br>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.<br>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.<br>**Exclude VM(s)**: the selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.<br>For Oversized VMs:<br>■ **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more resize actions for oversized VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>■ **RESIZE VM(s)**: The system displays a dialog box with suggestions for reducing vCPUs and memory. Click the edit icons to change resource size.<br>■ **EXCLUDE VM(s)**: Excludes the selected VMs.<br>■ **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br>For Undersized VMs:<br>■ **SCHEDULE ACTION**: Displays a dialog box enabling you to resize VM actions for undersized VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>■ **RESIZE VM(s)**: The system displays a dialog box with suggestions for increasing vCPUs and memory. Click the edit icons to change resource size.<br>■ **EXCLUDE VM(s)**: Excludes the selected VMs.<br>■ **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br>**SHOW\|HIDE EXCLUDED VMS**: toggle displays or hides the list of VMs you previously excluded.<br>**INCLUDE VM(s)**: include the selected VMs in the actionable list. |

## Using Workload Optimization to Improve Performance

Workload Optimization in VMware Aria Operations focuses on intelligently balancing and optimizing workloads across your virtualized infrastructure. The goal is to ensure that resources

are distributed efficiently to maintain optimal performance, prevent resource contention, and make the best use of available hardware resources.

Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.

Workload Optimization further allows you to potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention automatically runs an action, a data center performs at optimum.

Starting with version 8.6, you can run workload optimization on a custom data center which has clusters aross multiple data centers within a single vCenter Server instance. The prequisite for this is that the hosts in the different clusters must be under the same network. This means that the port groups must be the same across the data centers. You can activate cluster level optimization across data center boundaries by enabling the setting in the business intent workspace.

## VMware Aria Automation Integration

When you add an instance to a VMware Aria Automation adapter or solution pack as well as to a vCenter Server adapter instance that is connected to the VMware Aria Automation server, using VMware Aria Automation-managed resources, VMware Aria Operations automatically adds a custom data center for the vCenter Server, using VMware Aria Automation-managed resources.

On the VMware Aria Operations side, to get the day2 chain configured, you must make the following initial configurations:

1. In vCenter Server, **Administration -> Solutions** and then add the VMware vSphere adapter instance for the vCenter Server that is configured as an endpoint in VMware Aria Automation Server.
2. In vCenter Server, **Administration -> Solutions** and then add the VMware Aria Automation adapter instance for the server that will appear in theVMware Aria Operations and VMware Aria Automation integration day2 chain.

VMware Aria Operations can manage workload placement and optimization for the custom data centers that reside in VMware Aria Automation-managed clusters.

However, VMware Aria Operations is not permitted to set tag policies for the custom data center. (At the Workload Optimization screen, the Business Intent window is not operational for VMware Aria Automation custom data centers.) When rebalancing a VMware Aria Automation custom data center, VMware Aria Operations uses all applicable policies and placement principles from both systems: VMware Aria Automation and VMware Aria Operations. For complete information on creating and managing VMware Aria Automation custom data centers that are managed by VMware Aria Operations, see the VMware Aria Automation documentation.

## Configuring Workload Optimization

Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

1 Configure the Workload Automation Details. See Workload Automation Details.

2 If you do not use the AUTOMATE function in the Optimization Recommendation pane at the Workload Automation screen, configure the two Workload Optimization alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the actions calculated by Workload Optimization are run automatically. See Configuring Workload Optimization Alerts

### Prerequisites

Workload Optimization acts on objects associated with the VMware vSphere Solution that connects VMware Aria Operations to one or more vCenter Server instances. The virtual objects in this environment include a vCenter Server, data centers and custom data centers, cluster compute and storage resources, host systems, and virtual machines. Specific requirements:

- Workload Optimization is supported when the optimization candidate clusters are in the same network or port group.
- Optimization candidate clusters should have one of the following datastore configurations:
  - Shared Datastore
  - Datastore Cluster
  - vSAN Datastore
- A vCenter Adapter configured with the actions activated for each vCenter Server instance.
- A vCenter Server instance with at least two datastore clusters. sDRS must be activated and fully automated for datastore clusters configuration only.
- Any non-datastore clusters must have DRS activated and fully automated
- Storage vMotion must be set to ON at Workload Automation Details. The default is On.
- You must have permission to access all objects in the environment.

### Design Considerations

The following rules constrain the possible computer and storage resource moves that can be performed.

**Note** When VMware Aria Operations suggests that you optimize clusters in a data center, the system does not guarantee it can run an optimization action. VMware Aria Operations analytics can determine that optimization is desirable and can create a rebalancing plan. However, the system cannot automatically identify all the architectural constraints that may be present. Such constraints may prevent an optimization action, or cause an action in progress to fail.

- Moving compute and storage resources is allowed only within, not across data centers or custom data centers.
- Storage resources cannot be moved across non-datastore clusters.
- Compute-resource-only moves are permitted through shared storage.
- Virtual machines defined with affinity rules or anti-affinity rules are not to be moved.
- Virtual machines cannot be moved when residing on a local datastore, unless a storage swap exists on the local datastore.
- Virtual machines cannot be moved if they have data residing across multiple datastore clusters. Compute-only moves with similar shared storage are not permitted.
- A virtual machine cannot have data that resides across different storage types. For example, if a virtual machine has a VM disk on a datastore and a second VM disk on a datastore cluster, the virtual machine does not move, even when the datastore is shared with the destination or has swap on it.
- A virtual machine can use RDM so long as the destination datastore cluster can access the RDM LUN.
- A virtual machine can implement VM disks on multiple datastores inside a single datastore cluster.
- Workload Optimization may suggest moving virtual machines that are protected by vSphere Replication or Array Based Replication. You must ensure that all the clusters within a selected data center or custom data center have replication available. You can set up DRS affinity rules on virtual machines that you do not want moving across clusters.

## Workload Placement Page

Workload Placement lets you to optimize virtual machines and storage across clusters to reduce resource contention and maintain optimum system performance.

### Where You Find Workload Placement

Select **Workload Placement** under **Optimize** in the left menu.

### Workload Placement Page Options

In the Workload Optimization page, you see a list of data centers in a carousel, listed under three categories:

- Critical
- Normal
- Unknown

After you select a data center, you see the **ALL DATACENTERS** button on the upper right. Click **ALL DATACENTERS** when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers.

Table 6-8. Workload Placement Page Options

| Option | Description |
| --- | --- |
| View: | Filter results to include data centers, custom data centers, vRA-managed custom data centers, or all three. (Option appears if you select **ALL DATACENTERS** on the upper right.) |
| Group By: | Filter results by criticality (most out of balance data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. (Option appears if you select **ALL DATACENTERS** on the upper right.) |
| Sort By: | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>■ Alarm clock graphic - list data centers/custom data centers by time remaining.<br>■ Dollar sign - list data centers/custom data centers by potential cost savings with capacity optimization.<br>■ Scales graphic - Optimized. |
| Select data center or **ADD NEW CUSTOM DATACENTER** | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>■ Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object.<br>■ Select **ADD NEW CUSTOM DATACENTER** to display a screen that enables you to define a custom data center. |

### Data Center Options

After you select a data center from the carousel, you see the following information and options.

**Note** If you point your cursor to the lower right of a data center graphic, a tooltip may appear to let you know that the data center is using automated optimization.

### Optimization Status Tab

Appears when you select a data center or custom data center from the top of the screen.

#### Table 6-9. The Optimization Recommendation Card

| Option | Description |
| --- | --- |
| Status | ■ Optimized - indicates that workloads are optimized based on the settings you entered in the neighboring Operational Intent window, with no tag violations based on the settings you entered in the Business Intent window.<br>■ Not Optimized - indicates that one of the following conditions is true: workloads are not optimized based on the settings you entered in the neighboring Operational Intent window AND/OR there are tag violations based on the settings you entered in the Business Intent window. In the event of tag violations, the offending tags are listed. |
| OPTIMIZE NOW | Runs optimizing actions based on the settings you entered in your Operational and Business Intent settings. |
| SCHEDULE | Displays a dialog box enabling you to schedule one or more optimization actions. If schedules are currently set for data center or custom data center optimization, a check mark appears next to the data center or custom data center name. |

Table 6-9. The Optimization Recommendation Card (continued)

| Option | Description |
| --- | --- |
| AUTOMATE | Continually seeks optimizing opportunities for data center or custom data center, based on the settings in the neighboring Operational Intent window or Business Intent windows. Scheduled optimizations are turned off while automatic optimization is on. Also, automated alerts are not operational when automatic optimization is on. Once you confirm automation, the system displays message, for example, 1) "Workload Optimization is looking for opportunities to automate," 2) "Your workloads are optimized according to your settings." or 3) "No eligible moves were found within the max number of compatibility checks allowed."<br>**Note** To initiate Automation, you must have privileges for Environment -> Action -> Schedule Optimize Container. |
| TURN OFF AUTOMATION | Stops automatic optimization. Any scheduled optimizations come back online. |

**Note** Sometimes an optimizing action may be recommended, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that recommended optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation can incur stress in the future, then consolidation is not recommended.

**Note** Scheduled jobs created from the Workload Placement page will be shown in the Automation Central page. However, these jobs cannot be edited, cloned or previewed, but can only be activated or deactivated.

Table 6-10. The Operational Intent Card

| Option | Description |
| --- | --- |
| Utilization Objective | Indicates the main attribute of your current automation policy settings. Values are moderate, consolidate, or balance. |
| Edit | Displays the Workload Automation Policy Settings, where you can adjust settings for optimization and cluster headroom. For more information, see the topic, Workload Automation Details. |

### Table 6-11. The Business Intent Card

| Option | Description |
| --- | --- |
| Intent | Allows you to define zones of infrastructure within cluster boundaries. |
| Edit | Displays a workspace where you can select criteria for placement of VMs. |

### Table 6-12. Details for Are your clusters meeting your utilization objective?

| Option | Description |
| --- | --- |
| Are your clusters meeting your utilization objective? | Displays a table which presents data in the following columns based on the Demand model or Allocation model:<br>■ Name<br>■ CPU Workload<br>■ Memory Workload<br>■ DRS Settings<br>■ Migration Threshold<br>■ Violated Tags<br>■ VM Name<br>To activate the Allocation model, set the overcommit ratios in the active policy. For more information, see the topic, Allocation and Demand Model in Workload Optimization.<br>Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. The violated tags shows which clusters or host groups are breaching the business intent. The VM Name column shows the name of the VMs and tag value due to which tag violation is happening.<br>Provides the option to set the DRS automation level for individual objects. |
| VIEW DRS SUMMARY | Select a cluster in the list, then click this link to display a page containing metrics for DRS performance and cluster balance in the selected data center. |
| SET DRS AUTOMATION | Select a cluster in the list, then click this link to set the level of the DRS automation for the cluster. Note that clusters must be fully automated in order for workload optimization alerts to run actions set in the policies. |

## History Tab

Displays a graphical depiction of executed manual and automated optimizations for clusters in the selected data center or custom data center, based on parameters you provide.

Table 6-13. Details for History

| Option | Description |
| --- | --- |
| Selected WLP process drop-down | The optimization action whose details you want to display. |
| Time duration drop-down | Last *n* hours - select the time parameter: last 6, 12, 24 hours or last 7 days. |
| Quick filter | choose a cluster name to search on. |
| Squares graphic | toggle between viewing processes in icon or circle form. |
| Circle | Toggle between viewing processes presented in a circle or on a straight line. |
| Back arrow - reset action. | Reset action. |

If you point your cursor to a specific cluster as displayed on the screen, the details of the cluster appear in a tool tip. Click the note card icon on the lower right of the tool tip to go to the Details screen for the cluster. When displayed in the circle format, rings in the circle indicate how much CPU and how much memory was used at any given time. For example, if memory usage was higher than recommended based on your policy settings, the memory circle appears red.

Note the timeline across the bottom of the screen. When you choose parameters, for example, WLP process name, time parameter and cluster name, indicators appear along the timeline, showing when processes were initiated.

To zero in on a specific event, choose a process from the drop-down menu. You can also click points on the marker floating above the timeline, which causes a descriptive tool tip to appear, then double-click the 'Double-click to zoom' icon on the lower right.

If the event you choose includes an actual movement of VMs, you see a blue ball containing the number of VMs moved and showing the direction of the move and starting and ending clusters.

### Optimization Potential Tab

When you run Workload Optimization, VMware Aria Operations runs compatibility checks and excludes those VMs which have constraints, and only optimizes resources of those VMs which can be moved. If you want to see the total potential of your workload optimization, assuming that all VMs can be moved, click the **CALCULATE OPTIMIZATION POTENTIAL** button in the Optimization Potential tab. Optimization Potential disregards the underlying constraints and recommends moves before the compatibility checks. You can download the report to see more details. Storage DRS must be set to fully automated.

If you want to see what can be realistically optimized, click **OPTIMIZE NOW** in the **Operation Status** tab. After you click **OPTIMIZE NOW**, you can download a report to review incompatibilities.

The optimization potential report helps you understand the difference between the optimization achievable when you run **OPTIMIZE NOW** and the total optimization potential.

See also Example: Run Workload Optimization

## Configuring Workload Optimization Alerts

VMware Aria Operations provides two preconfigured alerts designed to work with the Workload Optimization feature. You must take additional action in the Policies area to turn on the alerts and automate them so that predetermined actions are run when the alerts fire.

The following preconfigured alerts are designed to work with the Workload Optimization feature:

- Data center performance can potentially be optimized in one or more clusters.
- Custom data center performance can potentially be optimized in one or more clusters.

The preconfigured alerts fire only if the AUTOMATE function is not turned on at the Workload Optimization screen which can be accessed by clicking **Optimize > Workload Optimization** on the left menu.

### Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI pages and manage vCenter Server objects.

### Procedure

1. Select **Configure** from the left menu, then **Policies**.
2. Select the policy that includes settings for the relevant data centers and custom data centers, for example, **vSphere Solution's Default Policy**.
3. Click the **Vertical Ellipses** next to **Add** and then **Edit**.
4. Click the **Alerts and Symptoms** tile.
5. Search on "can potentially be optimized" to locate the two alerts you want.
6. The alerts are activated by default/inheritance (see the **State** column).
7. The alerts are not automated by default/inheritance (see the **Automate** column). To automate the alerts, click the menu symbol to the right of the inherited value and select the green check mark.

### Results

Workload Optimization is fully automated for your environment.

### What to do next

To confirm that actions are taken automatically, monitor rebalance activity at the Workload Optimization screen.

## Custom Datacenters in VMware Aria Operations

A custom data center is a user-defined container for a group of objects that includes clusters, hosts, and virtual machines. Custom data centers provide capacity analytics and capacity badge computations based on the objects it contains. You can use custom data centers to forecast and analyze the capacity needs for your environment.

When you create a custom data center, you can include multiple cluster objects that span multiple vCenter Server instances. For example, you might have a production environment that spans multiple clusters, and you must monitor and manage the performance and capacity of the entire production environment.

After you create your custom datacenter, you can select it in the list of custom data centers to display a summary of its health, risk, and efficiency. To access the list of custom data centers, click **Environment > Custom Data Centers**.

This view displays the top alerts for the data center. To examine the capacity remaining for the custom data center, click the **Capacity** tab.

### How Custom Datacenters Work

In vSphere, a data center serves as a container for objects that a vCenter Server instance manages. A custom data center is a container that can include objects from multiple vCenter Server instances.

Custom data centers can contain vCenter Server instances, data centers, clusters, hosts, virtual machines, and datastores. You can add vSphere object types to a custom data center.

When you add an object, the hierarchical children of that object become part of the custom data center. An object can belong to multiple custom data centers.

When you create custom data centers, the system runs capacity analytics on the objects in the custom data center, even if those objects span multiple vCenter Server instances. For example, you might need to examine the capacity analytics data across multiple clusters, and the multiple vCenter Server instances that manage those clusters. You do not have to analyze the capacity of one cluster or one vCenter Server instance at a time. You can create a custom data center, add all the clusters to it, and see the capacity analysis in a single location.

### Custom Datacenters List

You can view the list of custom data centers that exist in your environment, and a summary view of its health, risk, and efficiency. In this view, you can click a custom data center to display the top alerts that the objects in the custom data center triggered.

### Where You Find Custom Datacenters

In the left menu, select **Environment > Custom Data Centers**.

### Table 6-14. Custom Datacenters Toolbar and Grid Options

| Option | Description |
| --- | --- |
| Toolbar options | In the toolbar click **Add** to add a new custom data center. Click the **Vertical Ellipses** against a custom data center name to perform the following actions:<br>■ **Edit**. Modify the custom data center.<br>■ **Delete** . Remove the custom data center.<br>■ **Clone** . Clone the custom data center.<br>■ **Go to Details**. Opens the custom data center page in the Object browser and displays the custom data center summary page |
| Filter | Limit the list of custom data centers to those data centers that match the text that you enter in the **Filter** text box. |
| Data grid | Lists the custom data centers in your environment. Click on the name of a custom data center to view the details.<br>To view a summary of the custom data center health, risk, and efficiency, click the **Vertical Ellipses** and then **Go to Details**. To edit, delete, or clone a custom data center, **Vertical Ellipses** and then the corresponding option. |

## Custom Datacenters Add and Edit Workspace

A custom data center is an object type that provides capacity analytics and capacity badge computations based on the objects it contains. You create a custom data center object and add inventory objects to it.

### Where You Create or Edit a Custom Datacenter

To create a custom data center, in the menu click **Environment > Custom Data Centers**, and click the **ADD** button.

To edit a selected custom data center, click the **Vertical Ellipses** to edit, remove or clone.

Table 6-15. Add and Edit Custom Datacenters Configuration Options

| Step | Option | Description |
|---|---|---|
| 1 - Custom Data Center | Name | Descriptive name of the custom data center. |
| | Description | Meaningful description for the custom data center. Provide specific information that other users must know about this custom data center.<br>Click the **NEXT** button when you are done. |
| 2 - Objects | Objects | This screen in this step is divided into two sections. On the right section, you can select an object from the tree view of your inventory. Drag an object to the left pane to add it to the custom data center. You can search for objects on the right hand side of the screen after selecting an object type. Once you have found the object, drag the object from the right section of the screen to the left section of the screen under **vSphere Hosts and Clusters**.<br>You can add vCenter Server instances, vSphere data centers, vSphere clusters, and ESXi hosts.<br>When you add an object, the hierarchical children of that object become part of the custom data center. An object can belong to multiple custom data centers.<br>To remove an object when editing, click the object in the vSphere Hosts and Clusters panel and click **REMOVE**.<br>Click the **CREATE** when you are done. |

## Workload Automation Policy Settings

Provides options for refining policy settings specifically for Workload Optimization.

### Where You Find Workload Automation Settings

Access this screen through the Policies pages:

1 Select **Configure** from the left menu, then select **Policies**.

2 Click the **Policy Definition** card.

3 Select a policy that you want to modify. Ideally, this should be an active policy. Or, click the **ADD** button to add a new policy.

4 Select the **Workload Automation** card to review the changes, or click **EDIT POLICY** to make changes.

,

Refer to Workload Automation Details.

## Business Intent Workspace

You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. VMware Aria Operations can be configured to leverage tags to define business-related placement constraints. Using tags, you can place VMs on hosts/clusters with matching tags, or exclude VMs and clusters from participating in Workload Placement.

### Where You Find Business Intent

You can open the Workload Placement page from the Optimize option in the left menu. On the Workload Placement page, select a data center or custom data center from the top row, and click **Edit** in the Business Intent window.

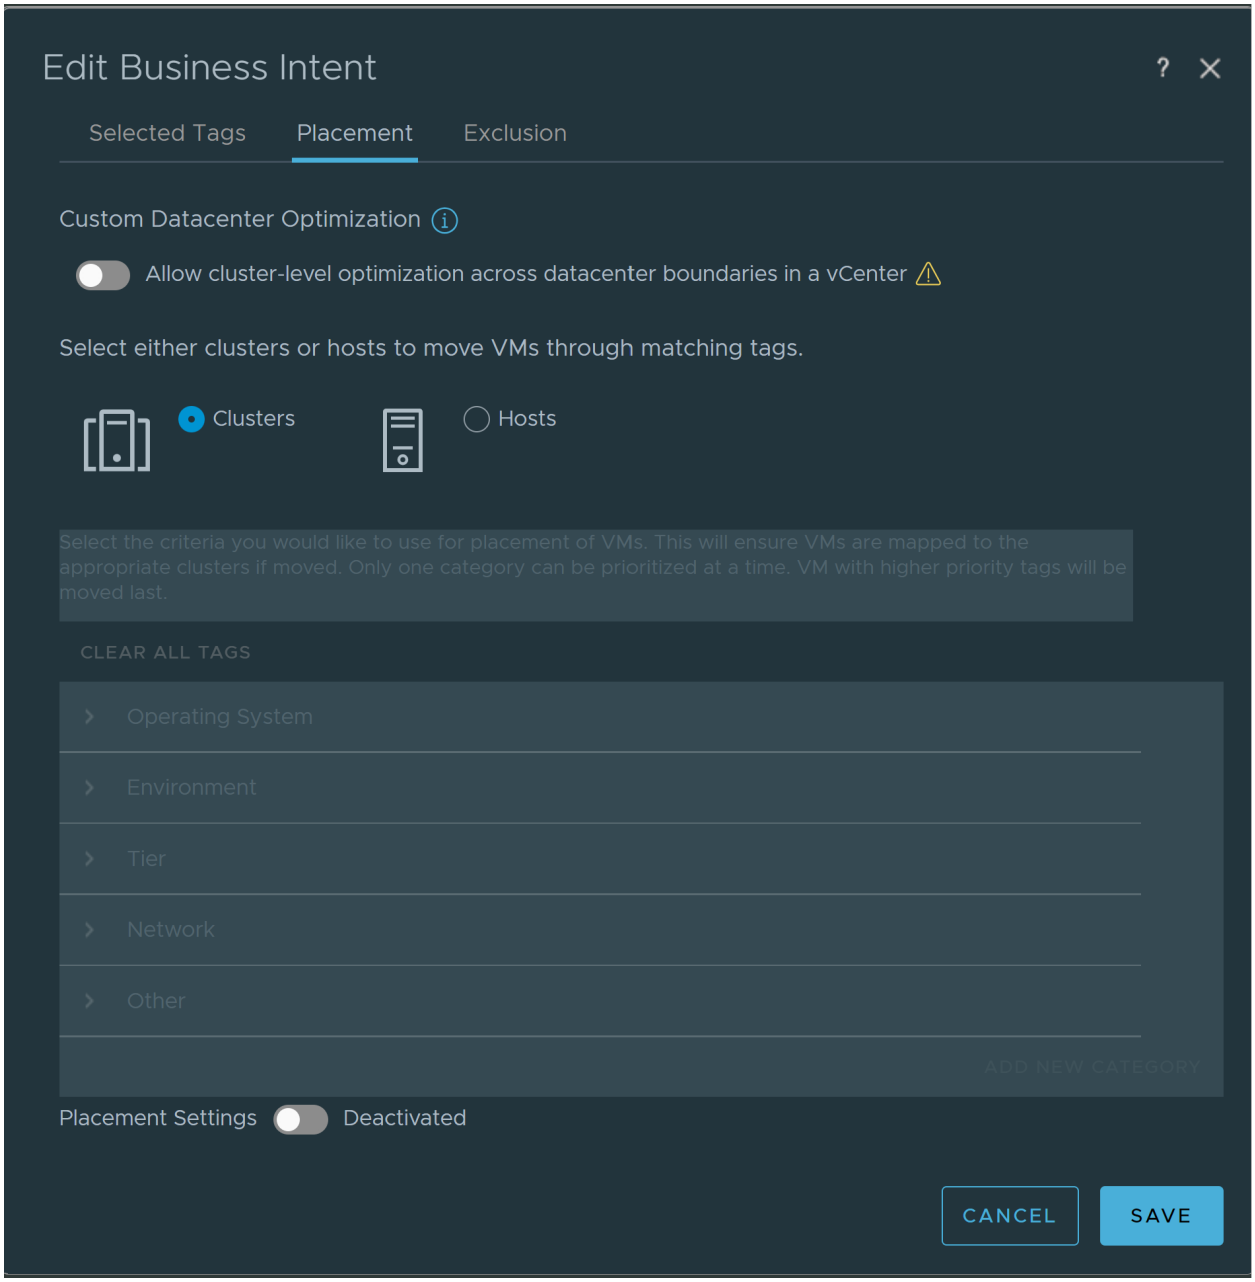


### How the Edit Business Intent UI Works

The Edit Business Intent dialog box has three tabs. The **Placement** tab is where you select the tags which VMware Aria Operations uses for Workload Placement. The **Exclusion** tab is where you specify tags to exclude VMs and Clusters from participating in the Workload Placement. For example, you may have VMs running critical workloads in a large environment. These VMs may be running on specific hosts so that they always get the resources they need for running the

workloads. You can exclude such VMs from Workload Placement if you do not want the VMs from moving from their original location. The clusters which are excluded using the exclusion tags are not displayed in the Workload Placement page, and are not considered in the Workload Placement calculation.

**Note** Tag exclusion was introduced newly in version 8.16.

The tags that you select in the **Placement** and **Exclusion** tabs are displayed in the **Selected Tags** tab.

You can activate cluster level optimization across data center boundaries in a vCenter Server if you have created a custom data center across multiple data centers within a single vCenter Server. You must ensure that the prerequisite networking requirements for movement of VMs are met before running workload optimization on such a custom data center.

To edit Business Intent values, you must have the necessary permissions. When you are logged in with administrative privileges, click Administration on the left menu and go to **Roles** under **Access Control** . Select the name of the role to which you want to provide permissions and then click **Edit** in the permissions section. Select the **Read** and **Write** checkbox under Administration → Configuration → WLP Settings.

### Establishing Business Intent

Tags are implemented in vCenter Server as *key:value* labels that allows operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

To specify tags that are either considered for placement, or excluded from placement, first go to the appropriate tab in the Edit Business Intent dialog box.

Select the radio button for the type of object you want to associate with VMs in this business intent session: Clusters or Hosts. Then, click the **Placement Settings** or **VM Exclusion Settings** toggle button.

The system provides several suggested categories. These categories are only suggestions. You must specify the actual categories in vCenter Server after you expand the section for a suggested category . For example, in section "Tier", you can specify the actual vCenter Server tag category that represents tier semantics, for instance, "service level".

- Operating System
- Environment
- Tier
- Network
- Other

Any actual categories you specify must first be created in vCenter Server.

Then you can associate tagged VMs with clusters or hosts, based on the rules for each type of tagging.

1 Click the chevron to the left of the first suggested category. A **tag category** field appears.

2 Click the drop-down menu indicator and choose a category from the list defined in vCenter Server.

3 Click the drop-down menu indicator in the Tag Name (Optional) field and choose a tag name from the list defined in vCenter Server.

4 Click **Include Tag**. All VMs with that tag are associated with the category.

### Rules for Host-Based Placement

To set host level placement constraints, VMware Aria Operations automatically creates and manages DRS rules. All conflicting user-created DRS rules are deactivated.

These rules include the following:

- Any VM-VM affinity and anti-affinity rules.
- Any VM-Host affinity and anti-affinity rules.

You must check the selection box next to the statement, "I understand that VMware Aria Operations will deactivate all my current and future DRS rules".

## View DRS Summary

The View DRS Summary page provides insight and perspective into the actions DRS is taking to balance a cluster. You can view DRS settings for the cluster and cluster balance metrics, and determine if recent vMotions are DRS- or user-initiated.

### Where You Find the View DRS Summary Page

From the left menu, click **Workload Placement** under Optimize. Then select a cluster name in the "Are your clusters meeting your utilization objective?" sedction. The dimmed View DRS Summary and Set DRS Automation links turn live. Click the link to display the DRS summary information.

**Table 6-16. DRS Summary Values**

| Pane/fields | Value |
|---|---|
| <*cluster name*> | Name of the selected cluster |
| Automation Level | Enabled/Disabled. DRS is running or not. |
| Migration Threshold | Aggressive/Default/Moderate |
| Active Memory Used | False/ *nn*% |
| Cluster Balance | Shows the variations in the DRS cluster balance metric over time as DRS runs. The graph shows how DRS reacts to and clears any cluster imbalance each time it runs. |
| Cluster Imbalance | The range of potential imbalance values, as expressed in vCenter DRS metrics. |

Table 6-16. DRS Summary Values (continued)

| Pane/fields | Value |
|---|---|
| Total Imbalance | The level of imbalance in a cluster, as measured by vCenter DRS metrics. |
| Tolerable Threshold | The upper limit of what is tolerable in cluster imbalance. Designated by a green dotted line, this is a vCenter DRS metric. |
| VM Happiness | A bar graph summarizing the total happy and unhappy VMs in the cluster. For individual VMs, there is a presentation of performance metrics related to its happiness, such as %CPU ready time and memory swapped. |
| Happy VMs | Total of happy VMs are shown in green. Click in the green zone to show a list of these VMs in the Happy/Unhappy VMs pane to the right. |
| Unhappy VMs | Total of unhappy VMs is shown in red. To show a list of these VMs in the Happy/Unhappy VMs pane to the right, click in the red zone . |
| Happy/Unhappy VMs | Lists by name all the VMs in the zone you clicked in the VM Happiness pane. |
| VM Metrics | Shows the trend in VM happiness or unhappiness |
| Recent vMotions | The number of recent vMotions, plotted against time. |
| vMotion Details | Shows the number of DRS-initiated and user (non-DRS) initiated vMotions over time. You can choose which type you want to view. |
| Date/VM | Date of a given vMotion. |
| Source/Destination | Source and destination of moved VMs. |
| Type | DRS-initiated or user initiated. |

## Optimization Schedules

Use the Optimization Schedules page to edit or delete optimization schedules that you set up in the Manage Optimization Schedule Dialog Box at the Workload Optimization main screen.

### Where You Find Optimization Schedules

- **Optimization Schedules** is available in the left menu, **Optimize > Workload Placement**
- At the Workload Placement Page page, select in the data center whose optimization schedule you want to edit or delete. Then click **SCHEDULE** in the Optimization Recommendation pane.

Table 6-17. Optimize Schedules Options

| Option | Description |
| --- | --- |
| Edit icon | Select a schedule from the list, then click the **Edit** icon.<br>The Manage Optimization Schedules appears, with the data for the selected schedule filled in. |
| Delete icon | Select a schedule from the list, then click the **Delete** icon.<br>The selected schedule is deleted and does not run. |

See also Example: Run Workload Optimization

## Manage Optimization Schedules

Allows you to set up a regular schedule for optimizing a selected container.

### Where You Find Manage Optimization Schedules

At the Workload Optimization screen, select **SCHEDULE** from the pane: Optimization Recommendation

| Option | Description |
| --- | --- |
| Schedule Name | Meaningful name for the schedule |
| Time Zone | Choose the time zone for the action |
| Recurrence | Indicate how often you want the optimize action to run. Complex schedules can be defined, for example, select the Monthly option and choose to run the action on Tuesdays and every other Thursday, beginning on the fifth of the month. |
| Start on: | Day to start the optimization schedule. |
| Start at: | Time to start the optimization schedule. |
| Expire after: | Designate a set number of scheduled runs. |
| Expire on: | Designate an exact date for the actions to end. |

See also Example: Schedule a Repeating Optimization Action

## Optimize Placement

A two-page dialog box that provides information about optimizing the workload of a selected container. When you run the optimization action, VMware Aria Operations checks which of the VMs can be moved to a different cluster for better optimization of resources, based on the settings you entered in your Operational and Business Intent settings. You can download a report that provides information about the list of VMs that were included in, and excluded from, the move plan. The report provides reasons as to why some VMs were excluded from the plan.

First page: The current workload ("before," for example, CPU 105%) and projected results ("after," for example storage utilization 45%) for a possible optimizing action.

Second page: The exact moves planned for compute and storage resources.

**Note** It is possible that there is no optimization move plan. Review the report to see why VMware Aria Operations could not provide a move plan.

### Where You Find Optimize Placement

At the Workload Optimization screen, select OPTIMIZE NOW in the Optimization Recommendation pane.

**Table 6-18. Optimize Clusters Options**

| Option | Description |
| --- | --- |
| Compare Cluster Balance | If you are satisfied with the before and after numbers (First page, above), click NEXT. |
| Review Optimization Moves | If you are satisfied with the moves planned (Second page, above), click BEGIN ACTION.<br>**Note** Review the optimization plan report before you click BEGIN ACTION. |
| Download Report | The optimization plan report is in CSV format, and provides the following information:<br>■ Summary of the optimization plan.<br>■ Summary of the moves that make up the optimization plan.<br>■ Issues related to the data center. Resolve these issues before proceeding with the optimization.<br>■ Issues and incompatibilities applicable to specific VMs and their configurations. Resolve these issues, if applicable.<br>■ Failed move attempts applicable to the specific VMs and their target destinations, as determined from the VM move plan. Resolve these issues and incompatibilities. |

See also Example: Run Workload Optimization.

## Using Workload Optimization

Use the Workload Optimization UI pages to monitor optimizing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and run actions directly.

VMware Aria Operations monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use optimization functions to distribute a workload differently in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects of interest.

The following examples demonstrate the primary ways you can use Workload Optimization to keep your data centers balanced and performing their best.

### Example: Run Workload Optimization

As a virtual infrastructure administrator or other IT professional, you use Workload Optimization functions to identify points of resource contention or imbalance. In this example, you manually run an optimization action to consolidate demand.

### Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

### Procedure

1 When you log into VMware Aria Operations, you see the Launchpad page. Click the **Capacity** card and then **View** in the **Workload Optimization** card.

The Workload Optimization page appears. Data centers are grouped by Criticality, with the three troubled data centers appearing in a carousel across the top of the page: DC-Bangalore-18, DC-Bangalore-19, DC-Bangalore-20. A Not Optimized badge appears in the lower right corner of each graphic.

2 If no data center is preselected, select DC-Bangalore-18 from the carousel.

Comprehensive data about the state of the data center follows.

3 Based on the available data, you determine an optimization action is required.

CPU workloads can be consolidated such that a host in Cluster 3 can be freed up.

Table 6-19. Panes and Widgets

| Pane | Contents |
| --- | --- |
| Workload Optimization | Status shows as Not Optimized. A system message says, "You can consolidate workloads to maximize usage and potentially free up 1 host."<br>The message reflects that you have set policies to emphasize consolidation as a goal in optimization moves. The system is saying you can free up a host through consolidation. |
| Settings | The current policy is Consolidate. The system advises: Avoid Performance Issues, Consolidate Workloads. |
| Cluster Workloads | Cluster 1 CPU Workload is 16%. Cluster 2 CPU Workload is 29%. Cluster 3 CPU Workload is 14%. Cluster 4 CPU Workload is 22%. |

4 Click **OPTIMIZE NOW** in the Workload Optimization pane.

The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.

5 If you are satisfied with the projected results of the optimization action, click **NEXT**.

The dialog box updates to show the planned moves.

6 If you need more information about the VMs which are included or excluded in the plan, click **Download Report** to see the optimization plan. You can review the reasons for incompatibilities and why some VMs were excluded from the plan.

7 Optional: If you want to know the total optimization potential of the move, assuming that there were no incompatibilities and all your VMs can be included in the optimization plan, click **Cancel**, and go to the Optimization Potential tab in the Workload Optimization page. Click **Calculate Optimization Potential** to see the total optimization potential of your data center.

8 Review the optimization moves, then click **BEGIN ACTION**.

The system runs the compute and storage resource moves.

#### Results

The optimization action moved compute and storage resources from some clusters to other clusters in the data center, and so freed up a host on one cluster.

**Note** The Workload Optimization page refreshes every five minutes. Depending on when you run an optimization action, the system might not reflect the result for up to five minutes, or longer when longer-running actions extend the processing time.

#### What to do next

To confirm that your optimization action was completed, go to the Recent Tasks page by clicking **Administration** in the left pane. In the Recent Tasks page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

**Note** Sometimes an optimizing action may be suggested, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that suggested optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation might incur stress in the future, then consolidation is not suggested.

### Example: Schedule a Repeating Optimization Action

As a virtual infrastructure administrator or other IT professional, you determine that compute and storage resources in a given data center are volatile and a regularly scheduled optimization action can address the problem.

VMware Aria Operations monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization page. Depending on what appears, you may determine that you must schedule optimization functions to distribute a workload more evenly in a data center or custom data center.

#### Prerequisites

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

### Procedure

1. Select **Workload Placement** under **Optimize** in the left menu
2. From the carousel of data centers across the top of the page, select a data center for which you want to schedule repeated optimization actions.
3. In the Workload Optimization pane, click **SCHEDULE**.
4. Give the schedule a name and choose a time zone.
5. Determine how often you want to repeat the optimization action and click the relevant **radio button** under Recurrence.

   Depending on your selection under Recurrence, additional options appear to the right. In this instance, you choose to repeat the optimization daily.
6. Leave the current date and time.
7. Select the **Repeat every day** radio button.
8. Select the **Expire after** radio button and tick the counter up to 6.
9. Click **Save**.

### Results

The optimization action repeats for six days, then stops.

At the Workload Optimization page, the Scheduled button appears in the upper right of the Workload Optimization pane if optimization actions are scheduled for the selected data center. If you want to edit or delete a schedule, click the **Scheduled** button. The Optimization Schedules page appears, where you can perform those actions.

**Note** If you schedule a number of optimization actions close together, and the optimization plans of two or more actions include overlapping functions, that is, they impact the same set of resources, the system shifts the actions into a queue. As a result, some actions may complete later than expected, with longer running actions and other potential system constraints extending the lag time. Optimization actions that do not overlap can run concurrently.

#### What to do next

To confirm that your optimization action was finished, go to the Recent Tasks screen by selecting **Administration** on the top menu, and clicking **History > Recent Task** in the left pane. In the Recent Tasks screen, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, filter on Event Source and enter the name of the scheduled optimization plan.

**Note** Because real-time data center resource contention is dynamic, the system calculates a new optimization plan each time the scheduled optimization action starts, but before it runs. The system does not run the action if the system determines that the data center container is balanced at this moment. On the Recent Tasks page, the name of the affected data center appears in the Object Name column, and the Message “The optimization of the selected container cannot be improved” appears under Details. Another possibility is that a scheduled optimization plan is attempted, but does not go forward. In this event - which is not the same as a "failed" action - the name of the affected data center also appears in the Object Name column.

# How to Plan for Capacity Changes

Planning for capacity changes in VMware Aria Operations involves simulating hypothetical changes to the environment, such as adding new VMs or hosts, to understand their potential impact on resource utilization. This helps in proactive planning.

## What-If Analysis: Modelling Workload, Capacity, or Migration Planning

Using the What-If tool, you can plan for an increase or decrease in workload or capacity requirements in your virtual infrastructure. To evaluate the demand and supply for capacity on your resources, and to assess the potential risk to your current capacity, you can create scenarios for adding and removing workloads. You can also determine how much capacity you require to make a migration work. You can run one scenario or group scenarios and run them cumulatively.

### Why Create a Scenario

A scenario is a detailed estimation of the resources you must have available in your environment to incorporate upcoming changes. You define scenarios that can potentially add resources to actual data centers. VMware Aria Operations models the scenario and calculates whether your desired workload can fit in the targeted data center. You can save multiple scenarios for comparison or review.

## Committed Scenarios

When you are sure that you need to reserve capacity, you can commit the scenario to have VMware Aria Operations set aside resources for new, upcoming, or planned workloads. A committed scenario is a supposition about how the capacity and load change on your objects when you change the conditions in your virtual infrastructure environment. You do not have to implement the changes that your committed scenario represents. By committing a scenario, you can determine your capacity requirements before you implement the actual changes.

### Why Create a Committed Scenario

In organizations which have separate capacity management and operations teams, committing a scenario helps stakeholders understand the current capacity and upcoming capacity requirements across the board. With committed scenarios, capacity is reserved and this prevents the operations team from performing adhoc resource increase on workloads, while the capacity manager is engaged in resource planning of new projects.

Committed Scenario also helps the team responsible for infrastructure expansion, as it provides actionable insights into future scenarios. In the event capacity becomes limited, it could be accounted for in the expansion.

## Where You Find What-If Analysis and Committed Scenarios

### What-If Analysis

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. On this page, you see options to view or add What-If Analysis scenarios and Committed Scenarios. Click the **What-If Analysis** tile. and in the What-If Analysis page, click the **ADD** button to see a list of seven What-If Analysis panes.

### Committed Scenarios

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click the **Committed Scenarios** pane, and in the Committed Scenarios page, click the **ADD** button to directly create a committed scenario without first creating a What-If scenario.

The overview tab of the What-If Analysis page has seven panes. Each pane lets you run What-If scenarios to optimize capacity based on the following areas:

- Workload Planning: Traditional
- Workload Planning: Hyperconverged
- Infrastructure Planning: Traditional
- Infrastructure Planning: Hyperconverged
- Migration Planning: VMware Cloud
- Migration Planning: Public Cloud
- Datacenter Comparison: Private Cloud

## How What-If Analysis and Committed Scenarios Work

You can run What-If scenarios to see how much capacity will remain after you add or remove VMs or hosts and add hyperconverged infrastructure (HCI) nodes. Migration planning shows you the capacity and cost information after migrating to cloud based infrastructure.

Scenarios that you save for later are displayed as a list when you browse to **Plan > Capacity** and under the **What-If Analysis** or **Committed Scenarios** pages. You can run, edit or delete the saved scenarios. When you run a scenario, whose start date occurs in the past, you get a dialog box asking you if you would like to run the scenario with the current date. If you choose **No**, the scenario is not run. If you choose **Yes**, runs the scenario with the current date as the start date. The end date is not affected.

Use the advanced filter and search box to look for saved scenarios by scenario name, scenario type, datacenter, and cluster. You can select more than one compatible scenarios and run them together. For example, you can create a scenario to remove hosts using the **Physical Infrastructure Planning** pane, because your organization has hardware that will soon become obsolete. You can create another scenario to add hosts to your physical infrastructure to account for new hardware that will replace the obsolete ones. You can run both these scenarios together to see the capacity after removing old hardware and adding new hardware.

You can only combine scenarios that pertain to the same object. Use the filters in the **Plan > Capacity** page to narrow down the list based on scenario name, type, data center, or cluster.

You can select the following combinations of scenarios and run them together:

**Workload Planning and Physical Infrastructure Planning**

- Add VMs
- Remove VMs
- Add Hosts
- Remove Hosts

## The Results Page

The Results page displays the results of running one or more saved scenarios. For Workload and Infrastructure Planning What-If scenarios, the summary page displays the allocation and demand values. If you do not see the allocation values, make sure that the overcommit ratios are activated in the policy. For more information, see the Policy Allocation Model Element topic in the Configuring guide.

To add or remove saved scenarios and run them again cumulatively, click **Edit** in the **Plan > Capacity** page. To commit a scenario and reserve capacity, click the **COMMIT SCENARIO** button. The Create Committed Scenario fly-out opens from the right hand side of the page. Add a name to the scenario you want to commit. Provide and implementation date, and optionally, an end date and click **SAVE**.

# What-If Analysis - Workload Planning: Traditional

You define scenarios that can potentially add workloads to actual data centers. VMware Aria Operations models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center. You can also define scenarios that can potentially remove workloads from data centers. VMware Aria Operations calculates the time remaining and capacity remaining on the cluster when workloads are removed from the cluster.

## Where You Find What-If Analysis - Workload Planning: Traditional

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. Click **Add VMs** or **Remove VMs** in the tile titled Workload Planning: Traditional.

## How What-If Analysis - Workload Planning: Traditional Works

Capacity Optimization allows you to forecast successfully the impact of adding a workload to an application. By trying various scenarios, you can arrive at an optimum configuration. When you add VMs in the Workload Planning: Traditional pane, you can select the exact data center or custom data center where you want to locate the new workload. You can even pick a specific cluster where the workload is to reside.

In selecting the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage. You have the further option to click Advanced Configuration and specify more precise characteristics for your workload.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system lets you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the new workload, enter the start and end date for the period when you want the workload to be active. The default is: starting today and ending one year from today. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria Operations analysis and assessment of your plan.

The system lets you know immediately if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. If you select scenario details, the system displays a graphic depiction of resource use. For each attribute value - vCPU, memory, and storage - the amount by which the workload increases the percentage of total application capacity used is shown against a time line. The graph shows the existing percentage used in blue and the total of existing usage and added usage as a percentage of total capacity in green.

If the proposed workload does not fit, the system announces the outcome and provides the following information:

- How much the added workload reduces the time remaining for the target cluster, for example, from one year to zero.
- The discrepancy between the space available in the target cluster and what the proposed workload requires, for example, 100 GB of memory.
- The cost of the workload on the VMware Hybrid Cloud and on the public cloud.

## About Clouds

When you run a scenario in What-If Analysis, you get a recommendation based on cost relative to workload placement on different clouds. This cost-based recommendation varies for different clouds.

Private Cloud and VMware Cloud on AWS costs are computed based on resource usage levels.

Public clouds, AWS, IBM Cloud, Google Cloud, Microsoft Azure, and user-defined cloud costs are dependent on the selected configuration, that is, for the allocated resources. These public cloud instances are selected based on the close proximity rule, with simulated resource allocation values and in some scenarios, the exact configuration match available in the cloud instance list is not available. Due to this issue, these public cloud costs can be inherently higher in comparison.

## How What-If Analysis - Remove Workload Works

This feature of Capacity Optimization lets you to forecast successfully the impact of removing a workload. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Workload Planning screen, you can select VMs from the concrete cluster data center or from the customer data center from which you want to remove the existing workload.

While removing workloads, you have two options to define the workload:

- Select existing VMs and use their projected utilization to evaluate the impact of removing workloads.
- Configure the workload manually by specifying the vCPUs, memory, storage, and expected use percentage.

Enter the start and end date for the period during which you want the workload to be removed. By default, the start date is today and the end date is one year from today. The end date is left empty by default. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria Operations analysis and assessment of your plan.

### Table 6-20. What-If Analysis Workload Page Options

| Option | Description |
| --- | --- |
| Add/Remove VMs | Click **Add VMs** or **Remove VMs** to create a scenario for adding or removing workload. When clicked, the command displays the Add Workload or Remove Workload screen. |
| Scenario Name | In the heading of the Saved Scenarios table. Selecting the check box next to the name selects all scenarios in the list and turns on the dimmed **Delete** button. |
| Scenario type | Name of the scenario type. Values are Add Workload, Remove Workload, Add Capacity, Remove Capacity, and Migrate. |
| <*scenario_name*> | Name of a saved scenario. Selecting the check box next to a name turns on the dimmed **Run Scenario**,**Edit**, and **Delete** buttons. |
| All Filters | Use the filter to search for a specific scenario by name or type. |
| Show Columns | Click the small button on the lower left to display the Show Columns dialog box. You can select up to four columns to display in the table: Scenario Name, Scenario Type, Date Created, and Scenario Start and End Date. |

## Add or Remove VMs: Traditional

As part of the What-If workload planning for traditional infrastructure, Workload Planning: Traditional is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. You also have an advanced configuration option that lets you define your configuration more precisely.

### Where You Can Add or Remove VMS

At the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional pane.

### Table 6-21. Workload Planning: Traditional Add VMs Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add the workload? Select from the list of existing data centers. You can optionally select the exact cluster where you want the workload to reside. |
| **Application Profile/Configure** | Allows you to configure the virtual compute resource, including vCPU, memory, and storage. |
| **Application Profile/ Import from existing VM** | Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |

### Table 6-21. Workload Planning: Traditional Add VMs Options (continued)

| Option | Description |
| --- | --- |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br>Click **OK** when done. |
| Choose Your Workload:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk space** | With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space. These are your allocation values. |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. VMware Aria Operations calculates demand based on these values. |
| **Annual Projected Growth** | Set the percentage by which you expect your capacity go grow, annually. Click **Advanced Configuration** to set the percentage growth of CPU, Memory, and Disk individually.<br>For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110.<br>The Annual Projected Growth can be set to 0% if no growth is expected. |
| | |
| Number of VMs (optional)/**Quantity** | You can optionally select how many VMs to spread the workload across. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

### Table 6-22. Workload Planning: Traditional Remove VMs Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove the workload? Select from the list of existing data centers. You can optionally choose the exact cluster from where you want to remove the workload. |
| **Application Profile/Configure** | Allows you to configure the virtual compute resource, including vCPU, memory, and storage.<br>After you have configured the scenario, enter the quantity of custom VMs that you want to remove. |

Table 6-22. Workload Planning: Traditional Remove VMs Options (continued)

| Option | Description |
| --- | --- |
| **Application Profile/Import Existing VMs** | Displays the Select VMs dialog box where you can choose one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to remove from your workload.<br>**Note** The recommended limit is 100 VMs as a maximum for workload removal. |
| **Application Profile / Custom**: Choose your workload<br>■ **CPU**<br>■ **Memory**<br>■ **Disk space** | With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank. |
| **Run Scenario** | Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

## Select VMs

Use the **Select VMs** dialog box to choose the VMs whose attributes you want to copy or remove for your Workload Planning: Traditional or Workload Planning: Hyperconverged what-if scenarios.

### Where You Find Select VMs

From the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional or Workload Planning: Hyperconverged pane. When you have entered a **Scenario Name** and **Location**, click the **Import from existing VM/Existing VMs** radio button, then click **Select VMS**. On the left is a selection box that allows you optionally to choose all VMs. To add a VM to the selected list on the right, double-click on the VM name. Following are the rest of your options:

### Select VMs

| Option | Description |
| --- | --- |
| All Filters | Use the advanced filter and search option to look for VMs. Click the drop-down icon to narrow your search based on the available options. |
| Select (nn). | Select the VMs listed on the current page, from which to import, or remove characteristics. |
| Select all (nn) VMS | Click to select all the VMs across all the pages, based on the filters you have set. The number of VMs that you can select by clicking this option is limited to 500 VMs. |

| Option | Description |
| --- | --- |
| Selected | List of VMs you selected from RESULTS. |
| OK | When you have selected the VMs you want, click **OK** to return to the Add Workload or Remove Workload screen, where your selected VMs are listed. |

Under Application Profile, in the Selected VMs table, enter the number of copies of each VM you selected to add or remove in the Quantity column.

## Scenario Summary Page for Workload Planning - Traditional

The scenario summary page for Workload Planning - Traditional displays the workload values after you add or remove VMs in a data center and cluster you specify.

### Summary Section

The summary section displays the inputs that you provided for the Add VMS or Remove VMS scenario. You can change the workload location by selecting a different vCenter and cluster combination. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

### Private Cloud and Public Cloud sections for Add VMS

The Private Cloud: Datacenter, Public Cloud, and Private Cloud sections display information which help you understand where your workload would fit, the associated costs, and the time remaining based on peak CPU, Memory, and Disk Space for demand and allocation model after you add VMs. You can choose to commit the Private Cloud: Datacenter results by clicking the **COMMIT SCENARIO** button. To see more details, click **VIEW DETAILS**.

You can see the cost per month for different private and public clouds. Click **Learn More** for a more detailed comparison based on your workload.

### Scenario Results Section for Remove VMS

This section displays a graphical view of the projected CPU, Memory, and Disk Space values for the demand and allocation models after the VMs are removed. You can understand how much time you have before the demand runs out.

# What-If Analysis - Workload Planning: Hyperconverged and VMC on AWS

You can perform Hyperconverged Infrastructure workload planning by adding or removing VMs to VMware vSAN activated clusters and running What-If scenarios. VMware Aria Operations shows you if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. .

## Where You Find What-If Analysis - Workload Planning: Hyperconverged

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. In the Workload Planning: Hyperconverged tile, click **Add VMS** or **Remove VMS**.

## How What-If Analysis - Workload Planning: Hyperconverged Works

You define scenarios that can potentially add or remove workloads to VMware vSANenvironment. The workload scenarios are based on VMs associated with specific storage policy related factors (such as FTT, RAID).

**Note** When a workload is added based on imported VMs, and the VM is currently in a VMware vSAN-activated cluster, the VMware vSAN policy settings are not applied and the current VM disk space is taken as is.

## Capacity and Cost Planning Support for Virtual Machines - VMC Datacenter

You can now perform capacity planning and cost calculations for a virtual machine (VM) in hyper-converged environment where the VM is part of the Vmware Cloud on Amazon Web Services (VMC) cluster. VMware Aria Operations provides accurate capacity recommendations and cost calculations when you add or remove VMs in hyper-converged environment from VMC data centers.

The cost calculation is based either on bills collected by VMC adapter or based on reference. To know more about VMC costing, see the topic VMware Cloud on AWS Cost Management in in the *VMware Aria Operations Configuration* guide.

## Add or Remove VMs: Hyperconverged

As part of the What-If workload planning for hyperconverged infrastructure, Workload Planning: Hyperconverged is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. The advanced configuration option lets you define your configuration more precisely.

### Where You Find Workload Planning

From the menu, select **Home** and **Optimize Capacity > What-If Analysis** in the left pane. Click **Add VMS** or **Remove VMS** in the **Workload Planning: Hyperconverged** pane.

Table 6-23. Workload Planning: Hyperconverged Add Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add the virtual machines? Select from the list of existing data centers. You can optionally select the exact cluster where you want the virtual machine to reside. |

### Table 6-23. Workload Planning: Hyperconverged Add Options (continued)

| Option | Description |
| --- | --- |
| **Application Profile/Configure** | Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space. |
| **Application Profile/Import Import from existing VM** | Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to incorporate as templates into your workload. |
| Select your workload:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk space** | With the **Configure** radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space. These are your allocation values. |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. VMware Aria Operations calculates demand based on these values. |
| **Annual Projected Growth** | Set the percentage by which you expect your capacity to grow, annually. Click **Advanced Configuration** to set the percentage growth of CPU, Memory, and Disk individually.<br>For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110.<br>The Annual Projected Growth can be set to 0% if no growth is expected. |
| Number of VMs (optional)/**Quantity** | You can optionally select how many VMs to spread the workload across. |
| Additional vSAN configuration | Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

### Table 6-24. Workload Planning: Hyperconverged Remove Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove the VMs? Select from the list of existing data centers. You can optionally select the exact cluster from where you want to remove the workload. |
| **Application Profile/Configure** | Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space.<br>After you have configured the scenario, enter the quantity of custom VMs that you want to remove. |

### Table 6-24. Workload Planning: Hyperconverged Remove Options (continued)

| Option | Description |
| --- | --- |
| **Application Profile/Import Existing VMs** | Displays the **Select VMs** dialog box where you can select one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to remove from your workload.<br>**Note** The recommended limit is 100 VMs as a maximum for workload removal. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br>Click **OK** when done. |
| **Application Profile / Custom**:<br>Choose your workload<br>■ **CPU**<br>■ **Memory**<br>■ **Disk space** | With the **Configure** radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space. |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. |
| Number of VMs (optional)/ **Quantity** | You can optionally select how many VMs to spread the workload across. |
| Additional vSAN configuration | Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank. |
| **Run Scenario** | Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

### Results: Add or Delete VMs to Hyperconverged Infrastructure

The scenario results are displayed when you run the scenario. In Private Cloud Data Center you can view the recommendation which provides details about the number of VMs to be added or removed from the VMware Cloud. You can also view whether the workload fits in your cloud environment and the cost increase or cost saving depending on whether you are adding or removing a VM from the VMware cloud. The Public Cloud tile displays the cost increase or savings across public clouds like Google Cloud, VMware Cloud on AWS, Amazon Web Services, IBM Cloud and others.

## Scenario Summary Page for Workload Planning: Hyperconverged

The scenario summary page for Workload Planning - Hyperconverged displays the workload values after you add or remove VMs to VMware vSAN activated clusters.

### Summary Section

The summary section displays the inputs that you provided for the Add VMS or Remove VMS scenario. You can change the workload location by selecting a different VMware vSAN activated clusters. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

### Private Cloud and Public Cloud sections for Add VMS

The Private Cloud: Datacenter, Public Cloud, and Private Cloud sections display information which help you understand where your workload will fit, the associated costs, and the time remaining based on peak CPU, Memory and Disk Space for demand and allocation model after you add VMs. You can choose to commit the Private Cloud: Datacenter results by clicking the **COMMIT SCENARIO** button. To see more details, click **VIEW DETAILS**.

You can see the cost per month for different private and public clouds. Click **Learn More** for a more detailed comparison based on your workload.

### Scenario Results Section for Remove VMS

This section displays a graphical view of the projected CPU, Memory and Disk Space values for the demand and allocation models after the VMs are removed. You can understand how much time you have before the demand runs out.

# What-If Analysis - Infrastructure Planning: Traditional

You define scenarios that can potentially add capacity to actual data centers or remove capacity from actual data centers. VMware Aria Operations models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center.

## Where You Find Infrastructure Planning: Traditional

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. Click **Add Hosts** or **Remove Hosts** in the Infrastructure Planning: Traditional tile.

## How the What-If Analysis for Infrastructure Planning: Traditional Works

Infrastructure Planning for traditional environments lets you to forecast successfully the impact of adding capacity to your environment or removing capacity from your environment. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Infrastructure Planning: Traditional pane, you can choose where you want to locate the additional capacity or from where you can remove the existing capacity.

In selecting the profile while removing capacity, you can select a profile only from server types that exist in your cluster.

In selecting the profile while adding capacity, you have two options:

- Select a server type from a list of commercially available servers. You can select from a list of 1) server types already in your cluster or 2) all server types approved for purchase.
- Configure a custom server manually by specifying CPU attributes, memory, and cost.

When you have set the profile for the new server, enter the number of servers to purchase or remove and the start and end date for the period when you want the scenario to be active. The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. The system can project scenarios ending up to one year from the current date. By default, the starting date is today and the ending date is one year from today.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria Operations analysis and assessment of your plan.

The system displays immediately the impact on cluster size of the additional or lesser amount of CPU and memory, and shows the total cost of adding or removing the specified capacity. The system also shows whether adding new capacity or removing capacity extends or shrinks the time remaining before CPU or memory runs out.

As well, the system displays a graphic depiction of resource use. For each attribute value - CPU and memory - the amount by which the workload increases or decreases the percentage of total capacity used is shown against a time line.

## Add or Remove Hosts

As part of the What-If analysis for physical infrastructure planning for traditional environments, Infrastructure Planning: Traditional pane is what you use to fill in the details of your What-If scenario. You select where to add or remove hosts, use an existing server type, or configure it yourself (when you add capacity), and establish a time frame.

### Where You Find Physical Infrastructure

At the What-If Analysis screen, click **Add Hosts** or **Remove Hosts** in the Infrastructure Planning: Traditional pane.

#### Table 6-25. Add Hosts Options

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add capacity? Select from the list of existing data centers, then select the cluster where you want one or more servers to reside. |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can select a commercial brand server or configure a custom server.<br>Number of Servers to add: increment the Quantity counter up to the number of servers you want. |

Table 6-25. Add Hosts Options (continued)

| Option | Description |
| --- | --- |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the What-If scenario. |
| **Run Scenario** | Click to run the scenario. The system calculates the cost of the scenario and determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

The system displays immediately the impact on cluster size of the additional CPU and memory, and shows the total cost of adding the specified capacity. The system also shows in graphical form whether adding the new capacity extends the time remaining before CPU or memory runs out.

Table 6-26. Remove Hosts Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario |
| **Location** | From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove one or more servers. |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can select only the server types that exist in your selected cluster.<br>The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the What-If scenario. You can select to keep the end date blank. |
| **Run Scenario** | Click to run the scenario. The system determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

The system displays the time remaining and the impact on CPU and memory with reduced capacity. The system also shows in graphical form whether removing capacity decreases the time remaining before CPU or memory runs out.

You can also see that the cost is based on the original purchase cost.

## Scenario Summary Page for Infrastructure Planning: Traditional

The scenario summary page for Infrastructure Planning - Traditional displays the capacity values after you add or remove hosts to a server in a data center and cluster you specify.

### Summary Section

The summary section displays the inputs that you provided for the Add Hosts or Remove Hosts scenario. You can change the capacity location by selecting a different vCenter and cluster combination, or server. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

### Scenario Results Section

This section displays a graphical view of the projected CPU, Memory and Disk Space capacity values for the demand and allocation models after the hosts are added or removed. You can understand how much time you have before capacity runs out and the costs savings. Click the **COMMIT SCENARIO** button to reserve this capacity.

## What-If-Analysis - Infrastructure Planning: Hyperconverged

You can perform infrastructure planning by adding or removing Hyperconverged Infrastructure (HCI) nodes in vSAN activated clusters and running What-If scenarios. VMware Aria Operations displays the cost, time remaining, and capacity remaining for CPU, memory, and disk space in the scenario results.

### Where You Find What-If Analysis - Hyperconverged Infrastructure

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. Select the Infrastructure Planning: Hyperconverged tile. To run a What-If scenario click **Add HCI Nodes** or **Remove HCI Nodes**.

### How What-If Analysis - Hyperconverged Infrastructure Works

You can add hyperconverged infrastructure to your VMware vSAN activated environment evaluate the increase in HCI capacity and cost. You can add up to 64 hosts per vSAN cluster. This number accounts for existing hosts in the cluster. VMware Aria Operations only lists vSAN and vXRail clusters in the location property. You can select existing server types from these locations and change the number of instances of these servers to add to your scenario.

### Add or Remove HCI Nodes

As part of the what-if analysis for physical infrastructure planning for hyperconverged environments, the Infrastructure Planning: Hyperconverged pane is what you use to fill in the details of your what-if scenario. When you add an HCI node, you can select an existing server type from your vSAN activated data center and change the number of instances of this server to calculate storage, compute capacity, time remaining, and cost. You can run the Remove HCI Nodes scenario to see the capacity changes after you remove HCI nodes from your data center.

#### Where You Find Workload Planning

At the **What-If Analysis** page, click **Add HCI Nodes** or **Remove HCI Nodes** in the **Infrastructure Planning: Hyperconverged** pane.

### Table 6-27. Add HCI Nodes Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario. |
| **Location** | Where do you want to add the HCI node? Select from the list of existing data centers. You must also choose the exact cluster where you want the HCI node to reside. |
| **Server Details** | Allows you to select an existing server type to calculate capacity, time, and storage remaining based on the number of instances of the server. |
| **Number of servers to add** | How many instances of the server do you want to add?<br>**Note** Only 60 new hosts can be added to the specified vSAN cluster as the maximum allowed is 64 hosts. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

### Table 6-28. Remove HCI Nodes Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove the server(s). |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can choose only the server types that exist in your selected cluster.<br>The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the what-if scenario. You can choose to keep the end date blank. |
| **Run Scenario** | Click to run the scenario. The system determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

## Scenario Summary Page for Infrastructure Planning: Hyperconverged

The scenario summary page for Infrastructure Planning - Hyperconverged displays the capacity values after you add or remove Hyperconverged Infrastructure (HCI) nodes in vSAN activated clusters that you specify.

### Summary Section

The summary section displays the inputs that you provided for the Add HCI Nodes or Remove HCI Nodes scenario. You can change the vSAN activated cluster location by selecting a different vCenter and cluster combination, or server. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

### Scenario Results Section

This section displays a graphical view of the projected CPU, Memory and Disk Space capacity values for the demand and allocation models after the HCI nodes are added or removed. You can understand how much time you have before capacity runs out and the costs savings. Click the **COMMIT SCENARIO** button to reserve this capacity.

## What-If-Analysis - Migration Planning: VMware Cloud

What-If-Analysis - Migration Planning, allows you to evaluate your plan for migrating or moving workloads across different VMware clouds. You can compare capacity and cost of workload across Oracle Cloud VMware Solution (OCVS), VMware Cloud on Dell EMC (VMCD), VMware Cloud on Amazon Web Services (AWS), Azure VMware Solution (AVS), and Google Cloud VMware Engine (GCVE). VMware Aria Operations evalautes the migration plan, calculates the cost and capacity requirements, and provides cost estimate for the selected VMC workload.

### Where You Find What-If Analysis - Migration Planning

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. In the Migration Planning: VMware Cloud tile, click **Plan Migration**.

### How What-If Analysis - Migration Planning Works

The What-If-Analysis feature of Capacity Optimization lets you forecast successfully the impact of migrating a workload to the VMware Cloud instance such as Oracle Cloud VMware Solution (OCVS), VMware Cloud on Dell EMC (VMCD), VMware Cloud on Amazon Web Services (AWS), Azure VMware Solution (AVS), and Google Cloud VMware Engine (GCVE). Once you select the Migration Planning screen, choose whether you want to run the scenario against a VMware Cloud on AWS or other type of cloud accounts. For the VMware Cloud, select the region where you want to migrate the workload.

When you have set the profile for the migrating workload, run the scenario to get the analysis and assessment of your plan. You can select one VMware cloud at a time and get the estimate for migration planning cost. Alternatively, you can save the scenario to edit or run later on. A list of saved scenarios is available in the Saved Scenarios tab on the What-If analysis page.

If you selected VMware Cloud on AWS for your scenario, the results list the VMware Cloud on AWS Assessment, with details of the VMware configuration. The result also displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription. In addition, the result displays the resource-use-level cost and monthly purchase cost for one-year and three-year subscriptions.

## About Clouds

The system might provide a recommendation based on the cost of placing the workload on different VMware clouds. This cost-based recommendation varies for different clouds.

For VMware Cloud on AWS, the system displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription, plus those same costs for one-year and three-year subscriptions.

VMware cloud costs are based on the selected configuration, that is, the allocated resources.

## Migration Planning: VMware Cloud

As part of the What-If Analysis function, Migrate is the form you use to fill in the details of your what-if scenario. You choose where to migrate the workload, then select the region.

### Where You Find Migration Planning

At the What-If Analysis screen, click **Plan Migration** in the **Migration Planning: VMware Cloud** tile.

When you run a scenario for What If: Migration for VMware Cloud, VMware Aria Operations might suggest the right cloud Instance suitable for the Workload Configuration selected by you. VMware Aria Operations also calculates the cost for that VMware Cloud's instance and displays the same.

Table 6-29. Migrate Options

| Option | Description |
|---|---|
| **SCENARIO NAME** | Name of your scenario |
| **SELECT CLOUDS** | Where do you want to migrate the workload?<br>Options:<br>■ **VMware Cloud on AWS**<br>■ **Azure VMware Solution (AVS)**<br>■ **Google Cloud VMware Engine (GCVE)**<br>■ **Oracle Cloud VMware Solution (OCVS)**<br>■ **VMware Cloud on Dell EMC (VMCD)**<br>**Note** You can select regions for the above cloud options. |
| **CLUSTER SETTINGS** | Specify the following cluster details:<br>■ Enter the **Instance Type**.<br>■ Enter the **Slack Space** in percentage.<br>■ Enter the **Steady state CPU headroom** in percentage. |
| **APPLICATION PROFILE/Configure** | Using the Application Profile you can configure the virtual compute resources, like vCPU, memory, and storage. |
| Select Your Workload:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk Space** | With the Configure radio button selected, you can size your migrating workload by defining values for vCPU, memory, and storage. |

Table 6-29. Migrate Options (continued)

| Option | Description |
| --- | --- |
| **Expected Utilization** | Specify the expected utilization or Click **ADVANCED CONFIGURATION** and specify the values for the following:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk Space**<br>■ **Disk Space Provisioning** - Select either Thin or Thick. |
| **Annual Projected Growth** | Specify the annual growth rate so that the system adjust the scenario calculations or click **ADVANCED CONFIGURATION** and specify the values for the following:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk Space** |
| **Number of VMs (OPTIONAL)** | You can optionally choose how many VMs to spread the workload across. |
| **Additional vSAN configuration** | Select **Account for Swap Space** to reserve swap space for any unreserved virtual machine memory.<br>■ Select the **Host failures to tolerate value** from the drop-down list.<br>■ Select the **Fault Tolerance Method**, the options are RAID -1 and RAID-5.<br>■ Select the **Dedup** value from the drop down list. |
| **APPLICATION PROFILE/Import from existing VM** | Displays the **Select VMs** button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group.<br>Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br>Click **OK** when done. |
| **RUN SCENARIO** | Click to run the scenario. The system calculates whether it fits into the location you chose. |
| **SAVE** | **SAVE** the scenario. |
| **CANCEL** | **CANCEL** the scenario. |

## VMware Cloud on AWS Assessment - Results

The scenario results are displayed when you run the scenario. You can view the recommendation which provides details about the number of hosts required for the VMware Cloud. You can also view the Total Cost associated with the recommended VMware Cloud for 3 years subscription and the Total Capacity Usage details for CPU, memory and disk space.

For VMware Cloud on AWS Assessment, you can edit the following options.

- **Edit Configuration** - you can edit the change in Reserved Capacity CPU, Reserved Capacity Memory, Fault Tolerance, and RAID Level values and save the values to the original configuration.
- **Change Plan** - you can use the **Choose Plan** option to change your subscription plan, the available options are one-year plan, three-year plan, or Pay-As-You-Go.
- **Edit Discount** - you can use the edit discount option to specify the discount value, the total cost for the subscription is equal to the actual utilization cost minus the discount percentage.

## What-If-Analysis - Migration Planning: Public Cloud

You define scenarios that can potentially migrate workloads to a public cloud. Use this scenario to determine where to move the workloads.VMware Aria Operations models the scenario and calculates the cost and capacity to fit your desired workload.

### Where You Find Migration Planning: Public Cloud

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **PLAN MIGRATION** button in the **Migration Planning: Public Cloud** card.

### How What-If Analysis - Migration Planning Works

This feature of Capacity Optimization lets you to forecast successfully the impact of migrating a workload to a public cloud instance such as Amazon Web Services, IBM Cloud, Microsoft Azure, or Google Cloud. Slect the region where you want to migrate the workload. If the public clouds listed out of the box do not suit your needs, you can also define your own public cloud and upload a rate card.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the migrating workload, run the scenario to get the VMware Aria Operations analysis and assessment of your plan. You can also select up to three public clouds to compare results. Alternatively, you can save the scenario to edit or run later on. A list of saved scenarios is available in the **Saved Scenarios** tab on the What-If analysis page.

For a public cloud target, the system lets you know immediately if the workload proposed for migration fits or does not fit in the suggested location. For example, if you selected AWS and the workload fits, the results list the Amazon Web Services Assessment, with details of the VMware Configuration and the AWS Equivalent. If the proposed workload does not fit, an error message appears: "Unable to identify a matching configuration instance in target location."

## About Clouds

The system might provide a recommendation based on the cost of placing the workload on different clouds. This cost-based recommendation varies for different clouds. You can modify the costs for public clouds by uploading a new rate card.

Public cloud costs are based on the selected configuration, that is, the allocated resources.

The public instance is selected based on the close proximity rule, with simulated resource allocation values. In some scenarios, an exact configuration match is not available in the list. Due to this lack of availability, the public cost can be inherently higher in comparison.

## Migration Planning: Public Cloud

As part of the What-If Analysis function, Migrate is the form you use to fill in the details of your what-if scenario. You choose where to migrate the workload, then select the region.

### Where You Find Migration Planning

At the What-If Analysis screen, click **PLAN MIGRATION** in the Migration Planning: Public Cloud card.

When you run a scenario for What If: Migration for Public Clouds (Not VMC), VMware Aria Operations might suggest the Public Cloud Instance suitable for the Workload Configuration selected by you. VMware Aria Operations also calculates the cost for that Public Cloud's instance and displays the same.

Table 6-30. Migrate Options

| Option | Description |
| --- | --- |
| **SCENARIO NAME** | Name of your scenario |
| **SELECT CLOUDS** | Where do you want to migrate the workload?<br>Options:<br>■ Aamazon Web Services<br>■ IBM Cloud<br>■ Microsoft Azure<br>■ Google Cloud<br>**Note** The cloud providers added in the Add Cloud Provider page are also included in the list.<br>You can select a maximum of three public clouds at a time for comparison. Hold the **Shift** key to select more than one public cloud provider. |
| **ADD CLOUD PROVIDERS** | You can add or edit the cloud providers and also edit the rate card of each individual cloud provider. |
| **APPLICATION PROFILE/Configure** | Using the Application Profile you can configure the virtual compute resources, like vCPU, memory, and storage. |

Table 6-30. Migrate Options (continued)

| Option | Description |
| --- | --- |
| Select Your Workload:<br>■ **CPU**<br>■ **Memory**<br>■ **Disc Space** | With the Configure radio button selected, you can size your migrating workload by defining values for vCPU, memory, and storage. |
| **APPLICATION PROFILE/Import from existing VM** | Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group.<br>Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br>Click **OK** when done. |
| Number of VMs (OPTIONAL)/**Quantity** | You can optionally choose how many VMs to spread the workload across. |
| **RUN SCENARIO** | Click to run the scenario. The system calculates whether it fits into the location you chose. |
| **SAVE** | **SAVE** the scenario. |
| **CANCEL** | **CANCEL** the scenario. |

# What-If Analysis - Data Center Comparison

You can select virtual machines to determine which of the preferred data centers (along with a specific choice of cluster or default cheapest cluster) are best fit from both cost effectiveness and capacity requirements perspective. The comparison helps you to find the right data center to place the workload from cost and capacity perspective.

## Where You Find What-If Analysis – Data Center Comparison

In the left menu, click **Plan > Capacity**. The Capacity Plan page opens. Click **What-If Analysis** and in the What-If Analysis page, click the **ADD** button. Click **COMPARE DATACENTERS** in the Datacenter Comparison: Private Cloud tile.

## How What-If Analysis - data center Comparison Works

This feature of Capacity Optimization lets you to compare cost across data centers within the private cloud environment. After you select the Datacenter Comparison screen, choose one or more data centers to compare the cost and run the scenario. VMware Aria Operations suggests which data center is most cost effective for the selected workload.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying CPU, memory, disk space, expected utilization, and annual projected growth.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for comparing the workload, run the scenario to get the VMware Aria Operations analysis and assessment of your plan. You can select up to three data centers to compare results. Alternatively, you can save the scenario to edit or run later. A list of saved scenarios is available in the Saved Scenarios tab on the What-If analysis page.

Cost varies from one data center to another depending on cost settings, which include cost drivers such as servers, facility, power, labor, license, network, and storage.

The data center comparison feature solves this problem by allowing you to select a data center which suits your requirement, is least expensive, and has adequate capacity.

## Datacenter Comparison

As part of the What-If Analysis function, Compare Datacenters is the form you use to fill in the details of your What-If scenario. Use this scenario to compare cost across data centers within the private cloud environment.

### Where You Find Compare Datacenters

At the **What-If Analysis** page, click **Compare Datacenters** in the pane titled Datacenter Comparison.

Table 6-31. Compare Datacenter Options

| Option | Description |
| --- | --- |
| **Scenario Name** | Name of your scenario. |
| **Select Datacenters** | Select the data centers for which you want to compare the costs. Use the advanced filter search to find a datacenter. |
| **Application Profile/Configure** | Using the Application Profile, you can configure the virtual compute resources, like CPU, memory, disk space, expected utilization, and annual projected growth. |
| Select Your Workload:<br>■ **CPU**<br>■ **Memory**<br>■ **Disk Space**<br>■ **Expected Utilization**<br>■ **Annual Projected Growth** | With the Configure radio button selected, you can size your workload by defining values for CPU, memory, disk space, expected utilization, and annual projected growth. |

Table 6-31. Compare Datacenter Options (continued)

| Option | Description |
| --- | --- |
| **Application Profile/Import from existing VM** | Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group.<br>Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br>Click **OK** when done. |
| Number of VMs (OPTIONAL)/**Quantity** | You can optionally choose how many VMs to spread the workload across. |
| **Implementation date/End Date (optional)** | You can specify the Start Date and End Date to compute the data center infrastructure cost for a specific time period. |
| **Run Scenario** | Click to run the scenario. The system calculates the cost of migration and checks whether the selected workload fits into the location you have chosen. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

## Retain Historical Data of VMs Migrated Using VMware Hybrid Cloud Extension

In VMware Aria Operations you can use VMware Hybrid Cloud Extension (HCX) to perform application migration, workload rebalancing, and business continuity across data centers and clouds. You can also migrate workloads from on-premises data centers to VMware Cloud.

Earlier when you performed HCX bulk migration to migrate workloads from one datacenter to another datacenter, or from one datacenter to VMware Cloud. VMware Aria Operations failed to retain the historic metrics.

Now VMware Aria Operations has implemented a solution which is triggered during HCX migration. The event helps VMware Aria Operations to collect details and manage the target vCenter to identify the workload migration.

After mapping the right attributes of VMs in the source datacenter with the VMs in the destination datacenter, you can verify if VMware Aria Operations is able to retain all the historical metrics. The HCX migration types supported in VMware Aria Operations are:

- Bulk Migration
- vMotion Based Migration (hot & cold)

- Replication Assisted Migration

HCX vMotion

VMware Aria Operations performs the following actions during HCX vMotion.

- Get the target `VCI`, `VM-VC-MOID` from event using the resource key.
- Get the source `VCID`, `VM-VC-MOID` from event using the resource key.
- Maps the correct target `VCID`, `VM-VC-MOID` to the source VM in VMware Aria Operations.

vMotion Generic Scenario

VMware Aria Operations performs the following actions during vMotion.

- Get the target `VCI`, `VM-VC-MOID` from event using the resource key.
- Detects the discovered target VM in VMware Aria Operations , based on the `VCID`, `VM-VC-MOID` attribute.
- For the detected target VMs, get VMs with **VM Entity Instance UUID**, and map the `VCID`, `VM-VC-MOID` for these VMs.
- Search `VCID`, `VM-VC-MOID` in the event message to find the actual vMotion VM.
- Set the correct target `VCID`, `VM-VC-MOID` to the right VM in on source VMware Aria Operations.

**Note** To know more about HCX Migration, see VMware HCX Product Documentation.

# Configuring Cost

7

SDDC costing is out-of-the box with VMware Aria Operations. There is no integration required with vRealize Business for Cloud.

Read the following topics next:


- Cost Overview
- Cost Analysis
- Cost Settings for Financial Accounting Model
- Overview of Cost Drivers
- Cloud Providers Overview
- Billing Enhancements for Horizon Management Pack and Virtual Hosts
- Editing Cost Drivers
- Cluster Cost Overview
- Publish Daily Cost Metrics for Virtual Machines
- Publish Tag Based Cost as Individual Metrics
- Pricing Overview
- Cost Calculation Status Overview
- Costing Enhancements


## Cost Overview

VMware Aria Operations now supports costing for private clouds, public clouds, and VMware Cloud Infrastructure. You can track expenses for a single virtual machine (VM), and how these expenses attribute to the overall cost associated with your private cloud accounts and VMware Cloud Infrastructure accounts.

The Cost Overview home page provides all the details about the costs associated with your VMware Cloud Infrastructure accounts, public cloud accounts, and your private cloud accounts. You can view the Total Cost of Ownership, Potential Savings, and Realized Savings for your VMware Cloud Infrastructure cloud accounts and vSphere Private Cloud accounts, and Total Cost of Ownership for your private cloud accounts.

You can view the cost details for the following private and public cloud accounts in VMware Aria Operations.

- vSphere On-Prem
- VMware Cloud Foundation (VCF)
- VMware Cloud on AWS
- Azure VMware Solutions (AVS)
- Amazon Web Services
- Microsoft Azure
- Google Cloud
- VMware Cloud on Dell EMC (VMC-D)
- Oracle Cloud VMware Solution (OCVS)
- All the values shown in vSphere On-Prem page might not match with the respective metrics at vSphere world level. That is because the metric at vSphere world level has values aggregated for all clouds like, Private, VMC, AVS, VMC-D, GCVE, and OCVS. So the values shown in Overview page are obtained after subtracting the aggregated metric values of VCF clouds from the metric value at vSphere World object.

**Private Cloud - Example: vSphere On-Prem**

The cost component of vSphere On-Prem private cloud account and VCF cloud account are Total Cost of Ownership, Potential Savings, and Realized Savings.

**Total Cost of Ownership** - The total cost of ownership widget displays the cost expenditure by capacity, by cost drivers and by datacenter. You can use the by capacity pie chart to view the compute, storage, and VM direct cost associated with your VMware Cloud Infrastructure cloud accounts. The cost drivers bar graph provides details of the cost drivers associated with your VCF cloud accounts and the horizontal graph for by datacenters provides the expense details of your VMware Cloud Infrastructure cloud accounts for individual datacenters.

**Note** The cost displayed in the Total Cost of Ownership widget might not match with the TCO metric at vSphere world level since it is the Total Aggregated Cost.

All the values shown in vSphere On-Prem widget might not match with the respective metrics at vSphere world level. This is because the metric at vSphere world level has values aggregated for all the clouds, like Private, VMC, AVS,VMC-D, GCVE, and OCVS. Hence the values shown in Overview page are obtained after subtracting the aggregated metric values of VMware Cloud Infrastructure from the metric value at vSphere World object.

**Potential Savings**

The potential savings widget displays the amount of savings you can potentially make from your VMware Cloud Infrastructure cloud accounts and vSphere Private cloud accounts. The pie chart for resources displays the cost savings opportunity across distributed across Idle VMs, Orphaned VMs, Oversized Hosts, Powered Off VMs, Reclaimable Hosts, and VM Snapshots. The horizontal graph for datacenters provides the overall potential cost savings for your cloud infrastructure and potential savings for individual datacenters. To know more about Potential Savings see, Potential Cost Savings Dashboard

**Note** The potential savings option is not available for public cloud accounts.

### Realized Savings

The realized savings widget displays the amount of savings you can potentially make from your VMware Cloud Infrastructure cloud accounts and vSphere Private cloud accounts. The pie chart for resources displays the cost savings opportunity distributed across Idle VMs, Orphaned VMs, Oversized Hosts, Powered Off VMs, Reclaimable Hosts, and VM Snapshots. The horizontal graph for datacenters provides the realized cost saving for your overall cloud infrastructure and realized savings for individual data centers. To know more about Potential Savings see, Realized Cost Savings Dashboard. To know more about reclamation cost savings, see Realized Cost Savings Using Reclamation Suggestion.

**Note** The realized savings option is not available for public cloud accounts.

### Dashboards

The cost dashboards widget lets you compare the cost of VMware Cloud Infrastructure with other public cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently. You can click dashboard links and navigate to the respective dashboard from the cost overview page.

**Note** The data displayed in the dashboards might not be specific to the selected cloud type. The dashboards might contain data from all vSphere instances and VMware cloud instances.

### Public Cloud - Example: Amazon Web Services

The cost component of public cloud account includes the cost of ownership associated with your public cloud account. The cost components of public cloud account are distributed across accounts, regions, and services. You can select individual account and view the cost associated with that account, region wise or service wise. For Google Cloud Platform the services cost component is replaced by product category.

**Note** The data for the public cloud accounts like AWS, Microsoft Azure, and Google Cloud is collected using the Cloud Health adapter, if the data is not displayed for this section, you have to deploy and configure the cloud adapter.

To know the granular cost visibility and to track your expenses of virtual machines accurately in a private cloud, see Overview of Cost Drivers.

To know the expenses related to the CPU, memory, and storage for a single virtual machine (VM), and how they attribute to the overall cost associated with your cloud infrastructure, see VMware Cloud on AWS and VMware Cloud on Dell EMC Cost Management in VMware Aria Operations.

To know more about Reference based costing for Google Cloud VMware Engine (GCVE) and Azure VMware Solution (AVS), see Reference Based Costing for Azure VMware Solution/Oracle Cloud VMware Solution.

## VMware Cloud on AWS and VMware Cloud on Dell EMC Cost Management in VMware Aria Operations

IT teams spend on purchasing infrastructure from Vmware Cloud on AWS (VMC) or VMware Cloud on Dell EMC (VMC-D). Now they can transfer these expenses (CPU, Memory, and Storage) to the application teams using VMC/ VMC-D cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

To use the VMC/ VMC-D costing feature you must set the **Billing Enable** option in **Advance Settings** section of a VMC/ VMC-D adapter to true. If it is set to false, the costing is based on the reference cost.

### VMC/ VMC-D Costing - Points to Remember

- The bill expenses or reference based costs are divided into CPU : memory : storage ratios, you can edit ratio, region, and discount from the Cost Settings for Financial Accounting Model topic.
- The bill expenses are allocated to clusters based on the region to which the cluster belongs.

  **Note** Some of the bill expenses (co-related to component resource objects in VMware Aria Operations) are divided across all the clusters, since at present VMware Aria Operations does not have an understanding of all the types of expenses.

- If the VMC/ VMC-D bills currency format is different from VMware Aria Operations currency format, then the VMC/ VMC-D bills are converted to VMware Aria Operations currency format and published on clusters and VMs. You can find the conversion factor as a property under VMC/ VMC-D Organization resource objects.
- The reference based costs that are picked are always on-demand. For example, add VMC/ VMC-D vCenter directly to VMware Aria Operations. For VMC, if you set the cloud type as VMware on AWS, then the reference costs of US east (N. Virginia) is picked by default. For VMC-D, if you set the cloud type as VMware Cloud on Dell EMC, the reference costs of the EU region is picked by default.

The following are some important points to consider when you select reference based costing and bill based costing.

- In case of reference based costing, we consider the Host as Production host and host type as On Demand, and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.
- When you have some unconfigured SDDCs in the organization, might not list all the hosts in the organization. So, if you use bill based costing which uses the list of hosts to calculate the cost, we might not be able to calculate the correct base rates.
- Expenses from the bills of your VMC/ VMC-D are distributed using a fair allocation algorithm to CPU, memory, and storage at the VM level. For accurate cost numbers, all the SDDCs must be configured in the given Organization.
- Ability to carry out workload planning with VMC/ VMC-D as the destination cloud using the new calculated base rates, based on your bills.

### How Does VMC/ VMC-D Cost Allocation Work

The VMC/ VMC-D cost allocation works as per the following sequence of events defined in VMware Aria Operations.

- Discover inventory of VMC/ VMC-D using vCenter and VMC/ VMC-D adapters.
- Acquire bills for VMC/ VMC-D from VMware Cloud Services Platform (CSP) using the VMC/ VMC-D native adapter.
- Identify the expenses per cluster using approximate values.
- Using the Total Cost Value, determine CPU, Memory, and Storage base rates.
- Apply base rates on VMs for allocation or utilization depending on the capacity model.

## Google Cloud VMware Engine Cost Management

IT teams spend on purchasing infrastructure from Google Cloud VMware Engine. Now they can transfer these expenses (CPU, Memory, and Storage) to the application teams using Google Cloud VMware Engine cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

### Google Cloud VMware Engine - Points to Remember

- To use the Google Cloud VMware Engine costing feature you must set the **Billing Enabled** option in the Advanced Settings section of a Google Cloud VMware Engine adapter to **true**. If it is set to **false**, the costing is based on the reference cost. Similarly, if any project of the Google Cloud VMware Engine adapter instance is configured without the CSP Refresh token, then costing is based on reference cost. For more details about the Billing Enabled and CSP Refresh token options, see Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations.

- The GCVEBill object type is associated with multiple projects of the Google Cloud VMware Engine adapter instance. All projects under the GCVEBill object type need to be configured in VMware Aria Operations to ensure accurate costing.
- If a project does not have any GCVEBill object type associated with the adapter instance, then costing is based on reference costing.
- You can create and use a single service account JSON that is common, similar to a super user account, for all the projects. For more details about configuring single service account JSON, see Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations.
- The bill expenses or reference based costs are divided into CPU : memory : storage ratios, you can edit ratio, region, and discount from the Cost Settings for Financial Accounting Model topic.
- The bill expenses are allocated to clusters based on the region to which the cluster belongs.

  **Note** Some of the bill expenses (co-related to component resource objects in VMware Aria Operations) are divided across all the clusters, since at present VMware Aria Operations does not have an understanding of all the types of expenses.

- If the Google Cloud VMware Engine bills currency format is different from VMware Aria Operations currency format, then the Google Cloud VMware Engine bills are converted to VMware Aria Operations currency format and published on clusters and VMs. You can find the conversion factor as a property under Google Cloud VMware Engine adapter instance resource objects.
- The reference based costs that are picked are always on-demand. For example, add Google Cloud VMware Engine vCenter directly to VMware Aria Operations. For Google Cloud VMware Engine, if you set the cloud type as Google Cloud VMware Engine, then the reference costs of US east (N. Virginia) is picked by default.

The following are some important points to consider when you select reference based costing and bill based costing.

- In case of reference based costing, we consider the Host as Production host and host type as On Demand, and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.
- When you have some unconfigured Private Clouds in the organization, might not list all the hosts in the organization. So, if you use bill based costing which uses the list of hosts to calculate the cost, we might not be able to calculate the correct base rates.
- Expenses from the bills of your Google Cloud VMware Engine are distributed using a fair allocation algorithm to CPU, memory, and storage at the VM level. For accurate cost numbers, all the Private Clouds must be configured in the given Organization.
- Ability to carry out workload planning with Google Cloud VMware Engine as the destination cloud using the new calculated base rates, based on your bills.

### How Does Google Cloud VMware Engine Cost Allocation Work

Google Cloud VMware Engine cost allocation works as per the following sequence of events defined in VMware Aria Operations.

- Discover inventory of Google Cloud VMware Engine using vCenter and Google Cloud VMware Engine adapters.
- Acquire bills for Google Cloud VMware Engine from VMware Cloud Services Platform (CSP) using the Google Cloud VMware Engine native adapter.
- Identify the expenses per cluster using approximate values.
- Using the Total Cost Value, determine CPU, Memory, and Storage base rates.
- Apply base rates on VMs for allocation or utilization depending on the capacity model.

## Reference Based Costing for Azure VMware Solution/Oracle Cloud VMware Solution

IT teams spend on purchasing infrastructure from Azure VMware Solution (AVS) and Oracle Cloud VMware Solution. Now you can transfer these expenses (CPU, Memory, and Storage) to the application teams using reference-based cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory, and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

### Reference Based Costing - Points to Remember

- The reference-based costs are divided into CPU : memory : storage ratios, you can edit ratio,region, and discount from the Cost Settings for Financial Accounting Model topic.
- The reference-based costs are allocated to clusters based on the region to which the cluster belongs.

The following are some important points to consider when you select reference-based costing.

- When you have some unconfigured private clouds in the organization, VMware Aria Operations might not list all the hosts in the organization.
- Ability to carry out workload planning with Azure VMware Solution/Oracle Cloud VMware Solution as the destination cloud using the new calculated base rates, based on your bills.
- For reference-based costing, we consider the Host as Production host and host type as On Demand and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.

# Cost Analysis

You can run cost, price, and VMware Cloud Bills analysis for your custom objects and groups. You can compare cost and price metrics across various custom groups. You can run an analysis and then view a list of the top and bottom 5, 10, and 20 objects based on selected metrics.

Additionally you can also compare the metrics of cost, price, and VMware Cloud bills across objects.

## Analysis of Cost Metrics

Cost is what you spend in currency value on managing and maintaining your infrastructure. Total Cost gives you an idea about the overall cost of ownership of your infrastructure. You can analyze cost metrics like Total Cost, Potential Savings, Realized Savings and so on, for your infrastructure that includes, VMs, hosts, and clusters. You can use lists and compare between selected objects types. For example, you can view the Top 10 VMs within a datacenter based on Total Cost for the last month or you can compare Cost Center A, Cost Center B, and Cost Center C based on Total Cost for the past 6 months.

### How Cost Analysis Helps

Some use cases for cost analysis are as follows:

- Financial reporting helps IT teams justify the expenses they incur to the finance team. It helps track expenses across private clouds, VMware Cloud on AWS, and public clouds.
- To drive accountability for resource usage, you can view the cost of infrastructure at the VM level. You can generate a bill for your usage and drive accountability within the application teams.
- Your infrastructure operations teams may want to quantify savings through reclamation or revoking of resources like idle VMs, snapshots, and so on. In VMware Aria Operations you can see such resources and view the price associated with them. Reclaimable or right sizeable resources have a cost associated that can be seen as Potential Savings.
- Realized Savings covers the cost savings from reclamation opportunities recommended by VMware Aria Operations.

### Metrics and Objects You Can Run Cost Analysis For

Objects for Cost metrics (Total Cost, Realized Saving, Potential Saving, MTD Cost, and Monthly Projected Total Cost):

- vCenter (all objects that have Cost: VM, Host, Cluster, Data Center, Resource Pools, and Datastore)
- CloudHealth (all key objects that have Cost: Services, Accounts, Regions, and so on)
- Horizon (all key objects that have Cost or are introducing Cost in this release – User, Session, VDI Pool, and RDS Farm)
- Container (Business Apps and Custom Groups)

  **Note:** For Custom Groups, it is the summary of Total Cost of all VMs under that custom group.

### Where You Find Cost Analysis

In the left menu, click **Plan > Cost**. The **Cost** page opens on the right. Click **Cost Analysis** from the right side. On the **Cost Analysis** page, you see the **Cost** tile. Click **Analyse** on the **Cost** tile to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

## Analysis of Price Metrics

You can analyze price metrics for your infrastructure. The Total Price metric is based on the pricing policy defined by the user. You can use lists and compare pricing between selected objects types.

### How Price Analysis Helps

Some use cases for price analysis are as follows:

- Typically, when your IT teams want the system to do the metering, but want to charge the application teams using their own Business or Pricing model, they can setup their own rate cards in VMware Aria Operations and in turn use chargeback reports as a billing mechanism for the users to help them understand what they are being charged for.
- Service providers as well as organizations following a chargeback model to drive accountability and to help generate usage/allocation based bills for their application teams or cost centers.
- Chargeback helps IT teams to drive behavior that is better achieved when there is a cost associated with the provisioned resources.

### Metrics and Objects You Can Run Price Analysis For

Objects for Price Metrics (Total Price):

- vCenter (all objects that have Price: VM and Cluster)
- vCloud Director (all key objects that have Price – OrgVdc, Org, vApp, and VM)
- Horizon (all key objects that have Price – User, Session, VDI Pool, and RDS Farm)
- Container (Business Apps and Custom Groups)

  **Note:** For Custom Groups, it is the summary of Total Cost of all VMs under that custom group.

### Where You Find Price Analysis

In the left menu, click **Plan > Cost**. The **Cost** page opens on the right. Click **Cost Analysis** from the right side. On the **Cost Analysis** page, you see the **Price** tile. Click **Analyse** on the **Price** tile to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

# Analysis of VMware Cloud Bills

VMware Cloud bills is the exact value you pay against the VMware Cloud infrastructure through Google Cloud VMware Engine, VMware Cloud on AWS, and VMware Cloud on Dell EMC. Currently these bills are available if the seller on record is VMware. Using this analysis you can view, list, and compare the different bills you have received for the given cloud over a period of time

## Metrics and Objects You Can Run VMware Cloud Bills Analysis For

Objects for VMware Cloud Bills (Total Component Expense, Outstanding Expense, Monthly Commit Expense, Monthly On demand expense, and Monthly Total Expense):

- VMware Cloud on AWS (all bill objects)
- VMware Cloud on Dell EMC (all bill objects)
- Google Cloud VMware Engine (all bill objects)

## Where You Find VMware Bills Analysis

In the left menu, click **Plan > Cost**. The **Cost** page opens on the right. Click **Cost Analysis** from the right side. From the **Cost Analysis** page, you see the **VMware Cloud Bills** tile. To analyze your prices, click **Analyse** on the **VMware Cloud Bills** tile to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

# Cost Analysis - Metrics and Metric Keys

As part of Cost Analysis you can analyze Total Cost, Realized Saving, Potential Saving, MTD Cost, Monthly Projected Total Cost, Total Price, Total Component Expense, Outstanding Expense, Monthly Commit Expense, Monthly OnDemand Expense, and Monthly Total Expense. Metrics are divided into Cost, Price and Bill sections. The tables below map the metric names as per the cost UI with the metric names and the metric keys for each adapter and analysis (cost, price, and bills).

## vCenter Adapter Metrics: Analyzing Cost

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Datacenter** | | |
| Total Cost | Cost\|Monthly Aggregated Total Cost (US$/Month) | cost\|aggrTotalCost |
| Realized Saving | Cost\|Realized Savings\|Total Realized Cost (US$) | cost\|realized_savings\|realizedTotalCost |
| Potential Saving (Current month) whole month | Reclaimable\|Potential Savings (US$) | reclaimable\|cost |
| **Object Type: Cluster Compute Resource** | | |
| Total Cost | Cost\|Monthly Cluster Total Cost (US$/Month) | cost\|totalCost |

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Host System** | | |
| Total Cost | Cost\|Monthly Server Fully Loaded Cost (US$/Month) | cost\|totalLoadedCost |
| MTD Cost | Cost\|MTD Server Total Cost (US$) | cost\|totalMTDCost |
| **Object Type: Virtual Machine** | | |
| Total Cost | Cost\|Effective MTD Cost (US$) | cost\|effectiveTotalCost |
| Monthly Projected Total Cost (current month) whole month | Cost\|Monthly Effective Projected Total Cost (US$/Month) | cost\|effectiveProjectedTotalCost |
| Potential Saving | Cost\|Potential Savings (US$) | cost\|reclaimableCost |
| **Object Type: Datastore** | | |
| Total Cost | Monthly Total Cost | cost\|totalCost |

## vCenter Adapter Metrics: Analyzing Price

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Virtual Machine** | | |
| Total Price | Summary\|Metering\|Total price | summary\|metering\|value |
| **Object Type: Cluster Compute Resource** | | |
| Total Price | Summary\|Metering\|Total price | summary\|metering\|value |

## VMware Cloud on AWS, VMware Cloud on Dell EMC, and Google Cloud VMware Engine metrics : Analyzing Bills

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Component** | | |
| Total Component Expense | Cost\|Total Component Expense | cost\|total_component_expense |
| **Object Type: VmcBill** | | |
| Outstanding Expense | Cost\|Outstanding Expense | cost\|outstanding_expense |
| Monthly Commit Expense | Cost\|Monthly Commit Expense | cost\|monthly_commit_expense |
| Monthly OnDemand Expense | Cost\|Monthly OnDemand Expense | cost\|monthly_on_demand_expense |
| Monthly Total Expense | Cost\|Monthly Total Expense | cost\|monthly_total_expense |

## Horizon Metrics: Analyzing Price

| Metric Name Per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Horizon RDS Farm** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type:Horizon RDS Host** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon User** | | |
| Total Price | Summary\|Total Price | summary\|vdiTotalPrice |
| **Object Type: Horizon VDI Application Session** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon VDI Desktop Session** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon VDI Pool** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |

## Horizon Metrics : Analyzing Cost

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Horizon RDS Farm** | | |
| Total Cost | Cost \| Monthly Total Cost | farmRDSHostCost\| totalRDSHostCostPerMonth |
| Monthly Projected Total Cost | Cost \| Monthly Projected Cost | farmRDSHostCost\| totalProjectedRDSHostCostPerMonth |
| **Object Type: Horizon RDS Host** | | |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | rdsHostCost\| rdsHostMonthlyEffectiveProjectedTotalCost |
| MTD Cost | Cost\|MTD Total Cost | rdsHostCost\|rdsHostMTDTotalCost |
| **Object Type: Horizon User** | | |
| Total Cost | Cost\|VDI\|Monthly Total Cost | userCost\|vdiUserCost\| totalUserCostPerMonth |
| Monthly Projected Total Cost | Cost\|VDI\|Monthly Projected Cost | userCost\|vdiUserCost\| totalProjectedUserCostPerMonth |
| **Object Type: Horizon VDI Application Session** | | |

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | vdiApplicationSessionCost\| vdiAppSessionMonthlyEffectiveProjectedTotalCost |
| MTD Cost | Cost\|MTD Total Cost | vdiApplicationSessionCost\| vdiAppSessionMTDTotalCost |
| **Object Type: Horizon VDI Desktop Session** | | |
| Monthly Projected Total Cost | Cost\|Monthly Projected Cost | vdiSessionCost\| totalProjectedVDISessionCostPerMonth |
| Total Cost | Cost\|Monthly Total Cost | vdiSessionCost\| totalVDISessionCostPerMonth |
| **Object Type: Horizon VDI Pool** | | |
| Monthly Projected Total Cost | Cost\|Monthly Projected Cost | poolDesktopCost\| totalProjectedDesktopCostPerMonth |
| Total Cost | Cost\|Monthly Total Cost | poolDesktopCost\| totalDesktopCostPerMonth |

### Cloud Director Adapter: Analyzing Price

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Organization** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |
| **Object Type: Organization VDC** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |
| **Object Type: vApp** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |

## New or Edit Cost Analysis

As part of analyzing cost, New Analysis is the page you use to fill in the details for your cost analysis scenario. You can select a metric, an object type, and the time range for which you want to analyze cost. Cost analysis are rendered using lists or bar graphs.

### Where You Can Create or Edit Cost Analysis

In the left menu, click **Plan > Cost**. The **Cost** page opens on the right. Click **Cost Analysis** from the right side. From the **Cost Analysis** page, based on what you want to analyze, click the **Analyse** button from the **Cost**, **Price**, or **VMware Bills** tile.

## Table 7-1. New or Edit Analysis Options

| Option | Description |
| --- | --- |
| Analysis Name | Provide a name for your cost, price, or VMware bill analysis. |
| Metrics | Based on the type of analysis you select, such as Cost, Price, or VMware Bills analysis, the metric options change. You select the analysis on the **Cost Analysis > New Analysis** tab<br>**Cost Analysis**<br>Select a metric from the list based on which you can analyze and compare costs. Here is the list:<br>■ Total Cost<br>■ Realized Saving<br>■ Potential Saving<br>■ MTD Cost<br>■ Monthly Projected Total Cost<br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| | **Price Analysis**<br>Select the Total Price metric based on which you can analyze and compare prices. The Total Price metric is based on the pricing policy defined by the user.<br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| | **VMware Cloud Bills Analysis**<br>Select a metric from the list based on which you can analyze and compare expenses. Here is the list:<br>■ Total Component Expense<br>■ Outstanding Expense<br>■ Monthly Commit Expense<br>■ Monthly OnDemand Expense<br>■ Monthly Total Expense<br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| Object/Group Type | Select an object or group type for which you want to calculate cost. Object types are filtered and displayed based on your metric selection. The objects for which the metric is applicable are displayed. For example, you can find out the Total Cost of a VM by selecting the Total Cost as a metric and VM as an object type.<br>You can also select a Group from the Container drop-down option. |
| Time Range | Select the time range for which you want to analyze and calculate costs. The time range displayed depends on the metric and object type you select. |

Table 7-1. New or Edit Analysis Options (continued)

| Option | Description |
| --- | --- |
| Analysis Scope | If you selected an object from **vCenter** in the Object/ Group Type field:<br>■ Select or enter the specific object type to narrow down your scope of analysis. You can also click the **Select** option to select a specific object type based on which the analysis is created.<br>If you selected a group from **Container** in the Object/ Group Type field:<br>■ This field is not activated. When you run the analysis, you will see a list of custom groups with the selected type. For example, Department, Function, Location, etc. |
| Run | Click to run the analysis. You see two kinds of results, a list or a bar chart. For more information, see the **Results: List or Compare** section below this table. |
| Save | Click to save the analysis. You can view saved the analysis from the **Saved Analysis** tab. For more information, see Saved Analyses. |

## Results: List or Compare

### Lists

The list results are displayed when you **Run** the Analysis. The list allows you to view the top and bottom 5, 10, and 20 results of your analysis. You can see a list or a bar graph of the selected object type defined in your analysis and the associated metrics for each object type. For example, if you selected Total Cost as a metric and VM as an object type, your list will display the all the VMs and their associated Total Cost.

### Compare

You can also compare costs, prices, or expenses by selecting specific objects from your defined analysis. Click **Select Objects** to select specific objects for comparison. You can see a list or a bar graph of the selected object from your defined analysis and the associated metrics for each object.

# Saved Analyses

The analyses that you save are listed in the **Saved Analysis** tab.

## Where You Find Saved Analyses

In the left menu, click **Plan > Cost**. The **Cost** page opens on the right. Click **Cost Analysis** from the right side. From the **Cost Analysis** page, click the **Saved Analyses** tab.

## Saved Analyses Options

| Option | Description |
| --- | --- |
| Run | Select and run a single analysis. |
| Horizontal Ellipsis > Delete | Select and delete multiple analyses. |
| Vertical Ellipsis > Run | Select and run the analysis. |
| Vertical Ellipsis > Edit | Select and edit the analysis. For more information, see New or Edit Cost Analysis. |
| Vertical Ellipsis > Delete | Select and delete the analysis. |

# Cost Settings for Financial Accounting Model

You can configure Server Hardware cost driver and resource utilization parameters to calculate the accurate cost and improve the efficiency of your environment.

Cost Drivers analyze the resources and the performance of your virtual environment. Based on the values you define, Cost Drivers can identify reclamation opportunities and can provide recommendations to reduce wastage of resources and cost.

## How to Set Your Depreciation Model and Years for vCenter

You can set your depreciation model and years using the following steps.

1 From the left menu, click **Configure** and then, click **Cost Drivers**.

2 Click **Settings**.

3 Select vCenter as **Infrastructure Type** from the drop-down menu.

4 In **Cost Settings - Financial Accounting Model** page, select the **Depreciation Years** between two and five.

5 Select the **Depreciation Model** as per your requirement and click **Save**.

## Editing Cost Ratio for VMware Cloud on AWS, Azure VMware Solution, Google Cloud VMware Engine, Oracle Cloud VMware Solution, and VMware Cloud on Dell EMC

You can set or modify the cost ratio for public cloud account using the following steps.

1 From the left menu, click **Configure** and then, click **Cost Drivers**.

2 Click **Settings**.

3 Select the public cloud account of your choice from the **Infrastructure Type** drop-down menu.

4 Select the required organization from the **Organization** drop-down menu.

5 Enter the cost ratio for **CPU**, **Memory**, and **Storage**.

6 Enter the Discount percentage.

**Note** Discount is applied only when billing is deactivated on VMC on AWS/Azure VMware Solution/ VMware Cloud on Dell EMC/Google Cloud VMware Engine/Oracle Cloud VMware Solution adapter and it is applied on the reference cost (list price) per hour applicable for that region.

7 Select the **Region** and click **Save**.

**Note** This value is considered when VMware Cloud vCenters are configured to VMware Aria Operations using vCenter adapters only.

The default discount % value is zero. You can set or edit the discount % for all the organizations in any of your VMware Cloud environments or you can set or edit the discount % for specific organization in any of your VMware Cloud environments. You can run cost calculation and check whether the discount % is reflected in the Monthly CPU Base Rate, Monthly Memory Base Rate, and Monthly Storage Base Rate metrics.

## Configuring Depreciation Preferences

To compute the amortized cost of the Server Hardware cost driver, you can configure the depreciation method and the depreciation period. Cost Drivers supports two yearly depreciation methods and you can set the depreciation period from two to five years.

**Note** Cost Drivers calculates the yearly depreciation values and then divides the value by 12 to arrive at the monthly depreciation.

| Method | Calculation |
|---|---|
| Straight line | `Yearly straight line depreciation = [(original cost - accumulated depreciation) / number of remaining depreciation years]` |
| Max of Double or Straight | `Yearly max of Double or Straight = Maximum (yearly depreciation of double declining balance method, yearly depreciation of straight line method)`<br>`Yearly depreciation of double declining method= [(original cost - accumulated depreciation) * depreciation rate].`<br>`Depreciation rate = 2 / number of depreciation years.`<br>**Note** `Double declining depreciation for the last year = original cost - accumulated depreciation` |

## Example: Example for Straight Line Depreciation Method

| Year | Original Cost | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|
| Year 1 | 10000 | 0 | `[(10000-0)/5] = 2000` |
| Year 2 | 10000 | 2000 | `[(10000-2000)/4] = 2000` |

| Year | Original Cost | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|
| Year 3 | 10000 | 4000 | `[(10000-2000)/3] = 2000` |
| Year 4 | 10000 | 6000 | `[(10000-2000)/2] = 2000` |
| Year 5 | 10000 | 8000 | `[(10000-2000)/1] = 2000` |

## Example: Example for Max of Double and Straight Line Depreciation Method

| Year | Original Cost | Depreciation Rate | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|---|
| Year 1 | 10000 | 0.4 | 0 | `Maximum([(10000-0)*0.4],[(10000-0)/5]) = Maximum(4000, 2000) = 4000`<br>which is 333.33 per month. |
| Year 2 | 10000 | 0.4 | 4000 | `Maximum([(10000-4000)*0.4],[(10000-4000)/4]) = Maximum (2400, 1500) = 2400`<br>which is 200 per month. |
| Year 3 | 10000 | 0.4 | 6400 | `Maximum([(10000-6400)*0.4],[(10000-6400)/3]) = Maximum (1440, 1200) = 1440`<br>which is 120 per month. |
| Year 4 | 10000 | 0.4 | 7840 | `Maximum([(10000-7840)*0.4],[(10000-7840)/2]) = Maximum (864, 1080) = 1080`<br>which is 90 per month. |
| Year 5 | 10000 | 0.4 | 8920 | `Maximum([(10000-8920)*0.4],[(10000-8920)/1]) = Maximum (432, 1080) = 1080`<br>which is 90 per month. |

# Overview of Cost Drivers

Cost Drivers are the aspect that contributes to the expense of your business operations. Cost drivers provide a link between a pool of costs. To provide a granular cost visibility and to track your expenses of virtual machines accurately in a private cloud, VMware Aria Operations has identified eight key cost drivers. You can see the total projected expense on your private cloud accounts for the current month and the trend of cost over time.

You can now set a total cost for the License, Labor, Network, Maintenance, and facilities cost drivers in VMware Aria Operations:

**Note** The total cost set by you is distributed across resources in the data center. For example, if you set the total cost for the RHEL license, the cost is divided across all the hosts and VMs which use the RHEL license.

According to the industry standard, VMware Aria Operations maintains a reference cost for these cost drivers. This reference cost helps you for calculating the cost of your setup, but might not be accurate. For example, you might have received some special discounts during a bulk purchase or you might have an ELA with VMware that might not match the socket-based pricing available in the reference database. To get accurate values, you can modify the reference cost of cost drivers in VMware Aria Operations, which overrides the values in the reference database. Based on your inputs, VMware Aria Operations recalculates the total amount for the private cloud expenses. After you add a private cloud into VMware Aria Operations, VMware Aria Operations automatically discovers one or more vCenter Servers that are part of your Private Cloud. In addition, it also retrieves the inventory details from each vCenter Server. The details include:

- Associated clusters: Count and names
- ESXi hosts: Count, model, configuration, and so on.
- Datastores: Count, storage, type, capacity
- VMs: Count, OS type, tags, configuration, utilization

Based on these configuration and utilizations of inventory, and the available reference cost, VMware Aria Operations calculates the estimated monthly cost of each cost driver. The total cost of your private cloud is the sum of all these cost driver expenses.

You can modify the expense of your data center. These costs can be in terms of the percentage value or unit rate, and might not always be in terms of the overall cost. Based on your inputs, the final amount of expense is calculated. If you do not provide inputs regarding expenses, the default values are taken from the reference database.

To know more about Reference Cost, see Reference Based Costing for Azure VMware Solution/ Oracle Cloud VMware Solution.

You can see the projected cost of private cloud for the current month and the trend of total cost over time. For all the expenses, cost drivers in VMware Aria Operations display the monthly trend of the cost variations, the actual expense, and a chart that represents the actual expense and the reference cost of the expense.

**Note** If the vCenter Server was added from more than six months, the trend displays the total cost for the last six months only. Otherwise, the trend displays the total cost from the month the vCenter Server was added into VMware Aria Operations.

Infrastructure Type

You have the option to select the infrastructure type as either vCenter or VMC on AWS, based on your selection the cost drivers are displayed on the Cost Drivers page. You can add or edit the cost drivers as per your requirement.

**Note** You can edit the cost driver values either in All Datacenter mode or Specific Datacenter mode. Ensure that you download and upload the cost driver configuration file in the same mode (either All DC mode or Specific DC mode).

For the vCenter infrastructure type the following private cloud cost drivers are applicable.

- Server Hardware : Traditional
- Server Hardware : Hyper - Converged
- Storage
- License
- Applications
- Maintenance
- Labor
- Network
- Facilities
- Additional Cost

For the VMC on AWS infrastructure type the following private cloud cost drivers are applicable.

- License
- Additional Cost

All other costs from VMC that are not directly attributed to specific hosts like load balancer, Tax, and other costs are grouped under additional cost driver and equally distributed among all hosts.

Export and Import Cost Drivers

Except for additional cost drivers, you can export or import the remaining cost drivers associated with your private cloud. With this functionality you can edit the cost driver values from the excel sheet instead of editing them from the user interface. You have an option to select all cost drivers and export them or you can select individual cost drivers and export them.

**Note** The import and export functionality is applicable only to vCenter cost drivers, the functionality is not available for VMware Cloud on Amazon Web Services.

## Table 7-2. Expense Types

| Cost Drivers | Description |
| --- | --- |
| Select Datacenter | The Select Datacenter option allows you to choose the data center for which the cost driver changes are applicable.<br>**Note** You can select a specific data center and modify the cost driver values of that data center, or you can modify the cost drivers and apply the changes to all the data centers. |
| Export | Click export to export the cost details for all the cost drivers. You can select individual cost drivers and export them also. |
| Import | Click import to browse and upload the updated cost driver configuration file (xls/csv). The import cost driver file should have the same template as that of the exported file.<br>**Note** You might be prompted with error messages if the uploaded file has errors. You can ignore the error or you can download the log file. You can click ignore error to omit the incorrect values and include the correct ones. |
| **Server Hardware : Traditional** | The Server Hardware cost driver tracks all the expenses for purchasing of hardware servers that are part of vCenter Servers. You see the server cost based on CPU age and server cost details.<br>**Note** You can now select an individual server from the server group and specify the unique cost for each individual server. |
| **Server Hardware : Hyper-Converged** | The Server Hardware : Hyper-Converged cost driver, tracks the expenses associated with hyper converged infrastructure components. The Server Hardware : Hyper-Converged cost driver includes expenses for the Hyper Converged servers like vSAN activated servers and vXRail. The expense provided is for both compute and storage.<br>**Note** The customizations that were performed for vSAN server costing under Server Hardware : Traditional in the earlier versions will not be carried forward to 7.5 as the vSAN activated servers will fall under Server Hardware : Hyper-Converged servers now. |
| **Storage** | You can calculate the storage cost at the level of a datastore based on the tag category information collected from vCenter Server. You see the storage total distribution based on category and the uncategorised cost details.<br>**Note** The vSAN datastores are not displayed as part of this cost driver page. |
| **License** | You see the licenses cost distribution for the operating systems cost and VMware license of your cloud environment.<br>**Note** For Non-ESX physical servers, VMware license is not applicable. |
| **Maintenance** | You see the maintenance cost distribution for the server hardware and operating system maintenance. You can track your total expense with hardware and operating system vendors. |
| **Labor** | You see the labor cost distribution for the servers, virtual infrastructure, and operating systems. You can view the total administrative cost for managing physical servers, operating systems and virtual machines. You can track all expenses spent on human resources to manage the data centers.<br>**Note**<br>■ Labor cost includes expenses on backup appliance virtual machine (VDP virtual appliance).<br>■ For physical servers, operating system labor cost and servers labor costs are applicable, virtual infrastructure cost is not considered. |

Table 7-2. Expense Types (continued)

| Cost Drivers | Description |
| --- | --- |
| **Network** | You see the networks costs by NIC type. You can track a network expense based on different types of NICs attached to the ESX server. You can view the total cost of physical network infrastructure that includes the internet bandwidth, and is estimated by count and type of network ports on the ESXi Servers.<br>**Note** For physical servers, the network details are not captured. So, the network cost is considered as zero. |
| **Facilities** | You see the cost distribution for the facilities such as real estate costs, such as rent or cost of data center buildings, power, cooling, racks, and associated facility management labor cost. You can point to the chart to see the cost details for each facility type. |
| **Additional Cost** | You can see the additional expenses such as backup and restore, high availability, management, licensing, VMware software licensing. |
| **Application Cost** | You can see the cost of different application services you are running in your environment compared to your overall expenses. Some examples of application cost are, cost of running SQL server cluster and cost of running Anti-virus on VMs. |

You can select a data center to view the information specific to the data center.

## Import or Export Cost Drivers

The cost driver editing process has been enhanced to support export and import of existing cost driver configurations. You can download (export) the existing cost driver configurations as an xls/csv file, edit the cost drivers and import the updated file back to the system. You must ensure that the import cost driver file should have the same template as that of the exported file.

### Where to find the Import or Export Option

From the left menu, click **Configure** > **Cost Drivers** > **Cost Drivers** tab. In the **Cost Drivers** tab, select either **Import** or **Export**.

The import and export functionality is applicable only for vCenter cost drivers, the functionality is not available for VMware Cloud on Amazon Web Services.

Using the import and export option, you can perform the following actions on the cost drivers:

- Export or import the cost driver configuration file.
- Read and edit the cost driver configuration file.
- Validate the updated cost driver configuration file and report errors.
- Identify the error from the log file and correct the errors.

You are prompted with error messages if the uploaded file has errors. You can correct the errors and upload the file, or you can ignore the errors, the system still allows you to upload the file.

# Cloud Providers Overview

By default, you can see that Amazon Web Services (AWS), Google Cloud, IBM Cloud, Microsoft Azure, VMware Cloud on AWS, Azure VMware Solution, Google Cloud VMware Engine VMware Cloud on Dell EMC, and Oracle Cloud VMware Solution are included in VMware Aria Operations. You can also add your own cloud provider by using a standard VMware Aria Operations template.

You can configure the new cloud provider as per the standard VMware Aria Operations template and perform a migration scenario. The VMware Aria Operations template contains data points for vCPU, CPU, RAM, OS, region, plan term, location, and built-in instance storage, you must provide these values when you add cloud providers. The result of the migration scenario helps you assess the cost savings achieved using your cloud provider against the default cloud providers.

You can edit the rate card for new cloud providers and default cloud providers. However, you cannot delete the default cloud providers.

## Add or Edit Cloud Provider

You can use the Add Cloud Provider workspace to add or edit a cloud provider. You can edit the cloud provider rate card for default cloud providers and add the new cloud provider.

### Procedure

1. From the left menu, click **Configure** and then, click **Cost Drivers** > **Cloud Providers**.
2. To add or edit the cloud provider, click **Add** icon or **Edit** option from the vertical ellipsis menu.
3. Enter or edit the **Cloud Provider Name**.
4. Select the cloud provider logo and click **Upload Logo**.
5. Click **Next**.
6. Click **Download Template** to edit the **Instance Configuration** and **Pricing Configuration** values.

   **Note** When you edit a cloud provider the Download Template link is replaced with Download Existing Rate Card. You can add, edit, or delete the rows with the instance names. You can fill in the instance settings values and the mandatory price configuration values as mentioned in the template.

7. Select the updated template and click **Upload Rate Card**.
8. Click **Validate**.

   **Note** VMware Aria Operations validates the rate card and reports success or failure. If errors are reported, you can correct the errors and proceed further.

9. Click **Finish**.

Results

The new cloud provider is now part of the VMware Aria Operations cloud provider list.

# Billing Enhancements for Horizon Management Pack and Virtual Hosts

The cost calculation of VMware Aria Operations has been enhanced to include the end point objects of Horizon Management Pack and virtual hosts. Earlier, the cost calculation was based on the metrics collected for each end point object.

The cost calculation for the end point objects is now based on the following criteria:

- Each Virtual Desktop Infrastructure Virtual Machine ( (VDI VM)) is counted as 0.25 Operating System Instance (OSI)
- Each Remote Desktop Service Host (RDS Host) is counted as 0.25 Operating System Instance
- One Operating System Instance for each Connection Server
- Virtual Hosts (ESXi hosted on a VM) is not counted against license usage
- VMs hosting the virtual hosts are counted against license usage

There are no VDI VM objects discovered by Horizon MP. VMware Aria Operations reports the number of VDI VMs in the bill. The number of VDI VMs appear under Virtual Machine node of the vCenter MP.

A new property `Is Horizon Managed` is introduced which is published by vCenter Adapter on VMs, Hosts and Clusters. VDI VMs are identified by this property (if the property is true for the VM, then the VM is considered as a VDI VM). If the property is true at VM level, then the same property is published for the cluster and all the hosts under that cluster, indicating that those hosts and clusters are Horizon Managed objects.

**Note** VMs and Hosts are considered as Horizon Objects only when Horizon Management pack is installed. If Horizon Management pack is not installed the objects are considered as Non-Horizon Objects irrespective of the object properties published by vCenter Adapter.

## How to Identify the Virtual Host

You can identify the virtual hosts by the following property.

- Hardware |Vendor = "VMware, Inc"

# Editing Cost Drivers

You can manually edit monthly cost of all the eight expense types from the current month onwards.

The configuration used for cost drivers determines how VMware Aria Operations calculates and displays the cost.

## Editing Server Hardware : Traditional

You can view, add, edit, or delete the cost of each server group, based on their configuration and the purchase date of a batch server running in your cloud environment. You can also specify the server cost for individual servers in a server group. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

### Procedure

1 From the left menu, click **Configure** and then, click **Cost Drivers**.

2 In the Cost Drivers tab, click **Server Hardware : Traditional**.

**Note** You can customize the default value of cost per server and specify exclusive values for other servers in the list.

For example, if you have a system that has eight servers you can modify the default reference value from $1000 to $800 for eight servers. You can also select two servers from the list and customize their value as $600. So, any new server that is added to the system will have the default value as $800.

3 Select the required edit mode for changing the server hardware cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 Click any server from the list of **Server Group Description**.

The cost drivers groups all server hardware from all data centers in your inventory based on their hardware configuration.

| Category | Description |
|---|---|
| Server Group Description | Displays the name of the server in your inventory. |
| Number of Servers | Displays the total number of servers of any particular hardware configuration in your inventory. |
| Monthly Cost | Displays the average monthly cost for server. This value is calculated as a weighted average of prices of purchased and leased batches. |

**Note** If the vCenter does not set the server vendor model, then that server is listed under the group Others in VMware Aria Operations for **Server Hardware : Traditional** cost driver and **Server Hardware : Hyperconverged** cost driver.

5 After selecting a server group, you can manually enter the required fields.

   a Enter the Purchase Type and Cost Per Server.

   **Note** You can use the **+ ADD COST PER SERVER** option to create multiple server batches and set the cost for a specific server in a server group.

   b Click **Save**.

## Editing Server Hardware: Hyper-Converged

You can view, add, edit, or delete the cost of Hyper converged Infrastructure (HCI) component in your server group. You can specify the cost per server and compute percentage exclusively for the HCI servers. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **Server Hardware : Hyper-Converged**.

3 Select the required edit mode for changing the server hardware cost.

   - **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   - **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

   **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 Click any server from the list of **Server Group Description**.

   The cost drivers groups all server hardware from all data centers in your inventory based on their hardware configuration.

   **Note** If the vCenter does not set the server vendor model, then that server is listed under the group Others in VMware Aria Operations for **Server Hardware : Traditional** cost driver and **Server Hardware : Hyperconverged** cost driver.

| Category | Description |
|---|---|
| Server Group Description | Displays the name of servers falling under vSAN clusters and vXrail servers in your inventory. |
| Number of Servers | Displays the total number of servers of any particular hardware configuration in your inventory. |

| Category | Description |
| --- | --- |
| Monthly Cost | Displays the average monthly cost for server. This value is calculated as a weighted average of prices of purchased and leased batches. |

**Note** You can edit the Compute Pct column to adjust the storage rate of the vSAN datastores. You can use the same percentage to determine the cost.

5 After selecting a server group, you can manually enter the required fields.

a Enter Purchase Type, Cost Per Server, and Compute Percentage.

**Note** You can use the **+ ADD COST PER SERVER** option to create multiple server batches and to customize the cost per server.

b Click **Save**.

## Edit Monthly Cost of Storage

The storage hardware is categorized according to the datastore tag category. You can edit the monthly cost per storage GB for the datastores based on their storage category (using tags) and storage type (NAS, SAN, Fiber Channel, or Block).

### Prerequisites

To edit the cost based on the storage category, you must create tags and apply them to the datastores on the vCenter Server user interface. For more information, see the VMware vSphere Documentation.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **Storage**.

3 (Optional) Select a tag category.

Assume that you have two tag categories (for example, Profile and Tiers) with three tags in each category, you can select either Profile or Tiers from **Tag Category** to categorize the datastores based on tags.

| Category | Description |
| --- | --- |
| Edit Mode | You can select the storage cost to be applicable for all the data centers or a specific data center.<br>■ **Edit for All Data Centers** mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.<br>■ **Edit for specific Data Center** mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost. |
| Select Data center | You can select the data center for which you want to change the storage cost. This field is applicable only for specific data centers. |
| Tag Category | ■ **Category** displays the tag categories for datastores and also the tags associated with the category. |

| Category | Description |
| --- | --- |
| Datastores | Displays the total number of datastores for a specific category or type. You can click the datastore value to see the list of datastores and its details such as monthly cost, total GB for each datastore. |
| Total Storage (GB) | Displays the total storage for a specific category or type. |
| Monthly Cost Per GB | Displays the monthly cost per GB for a specific category or type. You can edit this value for defining the monthly cost per GB for datastores. |
| Monthly Cost | Displays the total monthly cost for a specific category or type. |

4 Click **Save**.

## Edit Monthly Cost of License

You can edit the total operating system licensing cost and VMware license cost of your cloud environment. You can now set a total fixed cost for the license in VMware Aria Operations. The total license cost is divided across all the hosts present in the data center. You can edit the license cost by either selecting the ELA charging policy or selecting the per socket value.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **License**.

3 Select the required edit mode for changing the license cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 Click **Save**.

The Cost drivers display all the licenses in your cloud environment.

| Category | Description |
| --- | --- |
| Name | Displays the category of the operating system. If the operating system is not Windows or Linux, cost drivers categorize the operating system under **Other Operating Systems**.<br>**Note** Two new cost components, Monthly cost of VMware vSAN Per Socket and Monthly cost of VMware vSAN SnS have been included for the vSAN cost calculation. The default values for these components are based on the reference database values.<br>The licensing cost for the Windows operating system falls under one of the following categories:<br>Per Core License, applicable for<br>■ Windows Server 2016<br>■ Windows Server 2019<br>Per Socket License, applicable for<br>■ Windows NT 4.0<br>■ Windows Server 2003<br>■ Windows Server 2008<br>■ Windows Server 2012<br>Per Instance License, applicable for<br>■ Windows XP<br>■ Windows Vista<br>■ Windows 98<br>■ Windows 95<br>■ Windows 8<br>■ Windows 7<br>■ Windows 3.1<br>■ Windows 2000<br>■ Windows 10 |
| VMs | Displays the number of virtual machines that are running on the specific operating system. |
| Sockets | Displays the number of sockets on which the specific operating system is running. |
| Charged by | Displays whether a cost is charged by socket or ELA.<br>**Note** The Charged By column can be edited to mention that the cost is charged by socket, core, instance, or ELA. |
| Total Cost | Displays the total cost of the specific operating system. |

5 Click **Save**.

### Results

According to your inputs, VMware Aria Operations calculates and displays the total cost and updates the Charged by column with the option that you have selected.

## Customizing License Assignment

You can customize the licensing cost associated with your host using the custom license assignment option. Based on your requirement you can add or delete different operating system licenses to your host. With the custom license assignment option, you can increase or decrease the licensing cost associated with your host.

Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **License**.

3 Select the required edit mode for changing the monthly license cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 To customize the license cost for a specific server, click **Customize License Assignment**.

5 Select the host for which you want to customize the license cost and click **Assign**.

6 From the drop-down menu, select the operation system and click **OK**.

   The new operating system is listed under the Current Assignment column.

7 To remove an existing operating system from the host, under **Current Assignment** click X icon next to the operating system.

   The license cost of the removed operating system is reduced from the total cost.

8 Click **Save**.

9 Navigate to the **Cost Calculation Status** tab and click **Run**.

Results

The license cost is updated for the host, the * sign next to the host indicates that the license cost for the host has changed.

| Category | Description |
|---|---|
| Server | You can select the server for which you want to customize the license cost. |
| Current Assignment | Displays the current operating systems associated with the host. |
| Default Assignment | Displays the default operating systems associated with the host. |
| Filter | Filters the hosts based on the operating system type. |
| Reset | Resets the license cost of the host to the default value. |

## Edit Monthly Cost of Maintenance

You can edit the monthly cost of maintaining your cloud environment. Maintenance cost is categorized into hardware maintenance cost and operating system maintenance cost. Hardware

maintenance cost is calculated as a percentage of the purchase cost of servers. Operating system maintenance cost is calculated as a percentage of the Windows licensing costs. You can snow specify a total fixed cost for maintenance in VMware Aria Operations. The total maintenance cost is divided across all the hosts present in the data center.

Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **Maintenance**.

3 Select the required edit mode for changing the monthly maintenance cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 To customize the maintenance cost for a specific server, click **Edit For Individual Servers**.

5 Click **+Add Cost Per Server**.

6 From the **Select Server's for customization** drop-down select the required server and click **OK**.

7 Specify the Server Hardware Percentage and OS Percentage and click **Save**.

View the change in maintenance cost after you have run the cost calculation cycle.

## Edit Monthly Cost of Labor

You can edit the monthly cost of labor for your cloud environment. You can set a total fixed cost for labor in VMware Aria Operations. The total labor cost is divided across all the hosts present in the data center. The labor cost is combination of the total cost of the server administrator, virtual infrastructure administrator, and the operating system administrator.

Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Driver tab, click **Labor**.

3 Select the required edit mode for changing the monthly labor cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 Edit the monthly labor cost.

- Edit the detailed cost of labor.
- Edit the total monthly labor cost for servers, virtual infrastructure, and operating system.

5 To customize the labor cost for a specific server, click **Server** and then click **Edit For Individual Servers**.

6 Click **+Add Cost Per Server**.

7 From the **Select Server's for customization** drop-down select the required server and click **OK**.

8 Specify the Monthly hours of labor per hour, Labor hourly rate, and click **Save**.

The monthly labor cost is displayed.

| Category | Description |
|---|---|
| Category | Displays the categories of labor cost, servers, virtual infrastructure, and operating system |
| Calculated by | Displays whether the cost is calculated hourly or monthly. |
| Total Monthly Cost | Displays the total monthly cost of the particular category |
| Reference Cost | Displays the reference cost for the category from the cost drivers database |

### Results

The total monthly cost is updated. The hourly rate option or the monthly cost option that you select is updated in the **Calculated by** column.

## Edit Monthly Cost of the Network

You can edit the monthly cost for each Network Interface Controller (NIC) type or can edit the total cost of all the networking expenses associated with the cloud. You can now set a total fixed cost for network resources in VMware Aria Operations. The total network cost is divided across all the hosts present in the data center.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Driver tab, click **Network**.

3 Select the required edit mode for changing the monthly network cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.

- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **Note** When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4 Edit the monthly cost of network.

- Modify the values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and the 100 Gigabit NIC.
- Modify the total monthly cost of all network expenses associated with the cloud.

5 To customize the network cost for a specific server, click **Edit For Individual Servers**.

6 Click **+Add Cost Per Server**.

7 From the **Select Server's for customization** drop-down select the required server and click **OK**.

8 Specify values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and 100 Gigabit NIC and click **Save**.

View the change in network cost after you have run the cost calculation cycle.

## Edit Monthly Cost of Facilities

For your cloud environment, you can specify the total monthly cost of facilities or edit the facilities cost for real estate, power, and cooling requirements. You can now set the total fixed cost for facilities in VMware Aria Operations. The total facilities cost is divided across all the hosts present in the data center.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Driver tab, click **Facilities**.

3 Select the required edit mode for changing the monthly facilities cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

4 (Optional) Select the data center from the drop-down menu.

**Note** If you select Edit for specific data center as the edit mode, then the select data center option is activated.

5 Edit the monthly facilities cost.

- Modify the cost of rent or real estate per rack unit and modify the monthly cost of power and cooling per kilowatt-hour.

- Modify the total monthly cost of facilities.

6 To customize the facilities cost for a specific server, click **Edit For Individual Servers**.

7 Click **+Add Cost Per Server**.

8 From the **Select Server's for customization** drop-down select the required server and click **Ok**.

9 Specify the Cost Per Kilowatt and Real Estate Cost Per Rack Unit and click **Save**.

View the change in network cost after you have run the cost calculation cycle.

## Editing Additional Costs

The additional cost lets you add any additional or extra expense that is not covered by other expenses categorized by VMware Aria Operations. No reference value is present for this expense.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Driver tab, click **Additional Costs**.

3 Enter or select the cost type for the expenses.

**Note** Additional cost driver allows you to assign costs at Host, vCenter, VM, cluster, or data center level. For example, if you want to keep a cluster protected using the disaster recovery services, which involves an additional cost of $5000, you can do that by editing the additional cost driver.

4 Select the **Entity Type** and **Entity Selection**.

The **Entity Count** gets updated.

5 Enter the **Monthly Cost per entity** .

The **Total Cost per month** gets computed automatically.

6 Click **Save**.

**Note** After you update the Additional Cost configuration, you must reload the page manually to view the updated values.

## Edit Application Cost

VMware Aria Operations allows you to edit the application cost of an application present in your cloud environment. You can only modify the cost associated with the application, as all the other attributes are predefined.

### Prerequisites

Create applications in VMware Aria Operations.

Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cost Drivers tab, click **Applications**.

3 Click the edit icon next to the application cost you want to edit.

**Note** You can now specify the cost of packaged applications that are discovered by the Service Discovery Management Pack. Earlier the option to specify the application cost was available only for business applications defined by the user.

4 Modify the cost of the application.

5 Click **Save**.

# Cluster Cost Overview

VMware Aria Operations calculates the base rates of CPU and memory so that they can be used for the virtual machine cost computation. Base rates are determined for each cluster, which are homogeneous provisioning groups. As a result, base rates might change across clusters, but are the same within a cluster.

1 VMware Aria Operations first arrives at the fully loaded cost of the cluster from the cost drivers. After the cost of a cluster is determined, this cost is split into CPU and memory costs based on the industry standard cost ratios for the different models of the server.

2 The CPU base rate is first computed by dividing the CPU cost of the cluster by the CPU capacity of the cluster. CPU base rate is then prorated by dividing the CPU base rate by expected CPU use percentage to arrive at a true base rate for charging the virtual machines.

3 The memory base rate is first computed by dividing the memory cost of the cluster by the memory capacity of the cluster. Memory base rate is then prorated by dividing the memory base rate by expected memory use percentage to arrive at true base rate for charging the virtual machines.

4 You can either provide the expected CPU and memory use or you can use the actual CPU and memory usage values.

5 The base rates are monthly rates.

| Cluster Cost Elements | Calculation |
|---|---|
| Total Compute Cost | Total Compute Cost = (Total Infrastructure cost, which is a sum of all cost drivers) – (Storage cost) – (Direct VM cost, which is sum of OS labor, VM labor and any Windows Desktop licenses). |
| Expected CPU and Memory use | Expected CPU and Memory use = These percentages are arrived based on historical actual use of clusters. |
| Per GHz CPU base rate | Per GHz CPU base rate = (Cost attributed to CPU out of Total compute cost) / (Expected CPU Utilization * Cluster CPU Capacity in gHZ). |

| Cluster Cost Elements | Calculation |
|---|---|
| Per GB RAM base rate | Per GB RAM base rate = (Cost attributed to RAM out of Total compute cost) / (Expected Memory Utilization * Cluster RAM Capacity in GB). |
| Average CPU Utilization | Average CPU Utilization = (Cost attributed to CPU utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster). |
| Average Memory Utilization | Average Memory Utilization = (Cost attributed to Memory utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster). |
| Expected CPU Utilization | The utilization percentage level of CPU that the cluster is expected to operate.<br>**Note** When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value. |
| Expected Memory Utilization | The utilization percentage level of Memory that the cluster is expected to operate.<br>**Note** When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value. |

## Cluster Cost Computation with Allocation Model

You can now use the allocation model to compute the cost of clusters in VMware Aria Operations, earlier the cluster cost computation was based on the cluster utilization. When you perform cost computation using the allocation model, you can set the over commit ratio for CPU, RAM, and storage.

**Note** The allocation ratio can be set at both cluster level and datastore cluster level. You can also mention the storage base rate, which will displayed at the datastore level.

### Table 7-3. Cluster Base Rate Computation with Allocation Model

| Base Rate | Formula |
|---|---|
| vCPU Base Rate | vCPU base rate = B1 = (Cost attributed to CPU) / (Number of vCPUs in a cluster) |
| RAM Base Rate | RAM base rate = B2 = (Cost attributed to RAM) / Number of vRAMs in a cluster)<br>**Note** The cost computation is based on Over Commit ratio. If the Over Commit ratio is 1:4, and total cores in cluster are 6, then vCPU count = 24, in case if the allocated vCPU exceeds this targeted number, then the maximum value is selected. |

Table 7-4. Virtual Machine Cost Computation with Allocation Model

| Cost | Formula |
| --- | --- |
| Virtual Machine Cost | Virtual machine cost = (Number of vCPU allocated x B1 of cluster it belongs to) + Number of vRAMs allocated x B2 of cluster it belongs to) + storage cost + direct cost.<br>**Note** Storage allocated represents the Storage Base Rate based on allocation. |

## Editing Cluster Cost Calculation Methods

You can edit the cluster cost calculation method based on your business requirement. The cost of a cluster is derived from cost drivers. Virtual machine cost is calculated by multiplying base rates with the utilization of the VMs.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers**.

2 In the Cluster Cost tab, click **CHANGE**.

The Cluster Cost Calculation Methods dialog box is displayed.

3 Select any one of the Cluster Cost Calculation methods.

| Option | Description |
| --- | --- |
| **Cluster Usable Capacity After HA and Buffer** | The cluster cost calculated total capacity minus resources needed for High Availability (HA) and the capacity buffer setting.<br>Base rates are calculated based on the total cost of the cluster and Usable Capacity after HA and Buffer. Virtual machine costs are calculated from these base rates. Things to note:<br>■ A lower buffer reduces the base rates and causes the virtual machines to become cheaper.<br>■ A higher buffer increases base rates and causes the virtual machines to become more expensive.<br>■ Base rates and virtual machine costs do not change with the utilization of the cluster.<br>■ The difference between Usable Capacity after HA and Buffer and actual utilization is used to compute unallocated costs. |
| **Cluster Actual Utilization** | To calculate the base rates using the month to date average utilization of the cluster resources, select this option.<br>Base rates are calculated based on the total cost of the cluster and average utilization. Virtual machine costs are calculated from these base rates. Things to note:<br>■ Lower utilization level causes base rates to be high and virtual machines also become more expensive.<br>■ Higher utilization level causes base rates to be lower and virtual machines to become cheaper.<br>■ Base rates and virtual machine costs can change frequently based on the utilization of the cluster.<br>■ Unallocated cost of the cluster is near to zero.<br>■ The costs for unused resources are distributed across all virtual machines based on their actual utilization within the cluster. |

4 Click **SAVE**.

# Publish Daily Cost Metrics for Virtual Machines

In VMware Aria Operations, you can now publish daily cost metrics for all virtual machines. The daily cost metric of a virtual machine is the sum of daily cost of CPU, memory, storage, and additional cost associated with the virtual machine. Daily cost metrics provide granular details of the costs associated with the virtual machine.

## Formula to Calculate the Daily Cost and Monthly Cost of Virtual Machines

You can calculate the daily cost associated with a virtual machine using the following formula.

| Virtual Machine Cost Elements | Calculation |
| --- | --- |
| Daily Total Cost of Virtual Machine | Daily total cost of virtual machine = Sum of Daily cost of (CPU + memory + storage + additional cost) |

The change in daily cost metrics also changes the way you calculate the effective month to date cost of a virtual machine. You can use the following formula to calculate Effective Month to Date cost for a virtual machine.

| Virtual Machine Cost Elements for a Month | Calculation |
| --- | --- |
| Effective MTD Cost of VM | Sum of CPU daily cost from the beginning of the month until now + Sum of memory daily cost from the beginning of the month until now + Sum of storage daily cost from the beginning of the month until now + Sum of additional daily cost from the beginning of the month until now |

## How to View the Daily Cost Metrics of a Virtual Machine

To view the daily cost metrics of a virtual machine, from the menu, select **Administrator** and then in the left pane select **Inventory** > **vCenter Adapter** , select the specific **Virtual Machine**, and click the **Metrics** tab.

# Publish Tag Based Cost as Individual Metrics

You can publish tag-based additional cost as individual metrics using VMware Aria Operations.To publish tag-based additional costs as individual metrics, you must first activate the Tag based Costing Metrics at the Global Settings level. If you activate the tag-based costing metrics at a VM level, then each of the tag-based cost is considered as an independent instance metric on the virtual machine.

## How to Activate Tag Based Costing Metrics

From the left menu, click **Administration** and then click **Global Settings**. Select **Cost/Price** and navigate to the Tag based Costing Metrics. Move the toggle button to the right to activate the Tag based costing metrics. To know how to set the Additional Cost at a VM level, see Editing Additional Costs.

After you assign the cost for VMs with tags or custom properties, you must run the Cost Calculation and verify whether the additional cost is reflected in the costing metrics. The additional cost metrics can be viewed from the following location for the selected virtual machine **Object Browser > VM > Cost > Daily Tags and Custom Properties Cost**.

To verify whether the tag-based cost metrics update is reflected in your cost, run the cost calculation and check the tag-based cost metrics.

# Pricing Overview

You can create pricing cards in VMware Aria Operations to calculate the price associated with your virtual infrastructure. You can assign pricing cards to vCenters or Clusters, depending on the pricing strategy determined by VMware Aria Operations administrator. The pricing cards help you to set the price for each resource present in your virtual environment.

You can customize the pricing card as per your requirement. VMware Aria Operations has two types of pricing cards, rate-based pricing card and cost-based pricing card. After configuring a pricing card, you can assign it to one or more vCenters or Clusters as determined by the pricing strategy.

**Note** From this release, you can continue to create pricing cards or you can migrate vCenter pricing cards to vCenter pricing policies. For details about creating a new pricing card, see Add New Pricing Card. For details about migrating vCenter Pricing cards to vCenter pricing policies, see Migration of vCenter Pricing Cards to vCenter Pricing Policies.

## How Is Price Calculated

In rate-based pricing policy VMware Aria Operations calculates the virtual infrastructure price based on the rate card defined by you. For rate-based pricing policy VMware Aria Operationslets you define cost elements as per your requirements.

The server recalculates the price every 24-hours, the price calculation for the new pricing cards is done in the next VMware Aria Operations price calculation cycle.

## Hierarchy of Pricing Policy

The assignment of policy in VMware Aria Operationswill be for Clusters and vCenters. The price is calculated for virtual machines, then it is aggregated and rolled up to vCenter. If there are two policies, a default policy for vCenter and another policy for Cluster, then the price calculation is based on the cluster policy for all the resources under the cluster. After that the cluster cost is rolled up to vCenter.

When a virtual machine is under VMware Aria Automation hierarchy and vCenter hierarchy, then the pricing is calculated based on the VMware Aria Automation hierarchy and the virtual machine is removed from the vCenter resources and included under VMware Aria Automation resources.

## Pricing Support for VMware Cloud on AWS Resources

You can create a pricing policy in VMware Aria Operationsand assign it to VMware Cloud on AWS (VMC) resources, however you can only use the rate-based pricing policy for VMC-related objects.

**Note** When you assign Cost-Based Policy for VMC resources, the policy is not applied, and price calculated for the policy is reported as zero.

## Add New Pricing Card

You can add and assign new pricing card to vCenter and Clusters in VMware Aria Operations. The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

### Procedure

1 From the left menu, click **Configure** and then click **Cost Drivers** > **Pricing**.

2 Click **New Pricing Card** and configure the details of the pricing card.

#### Table 7-5. Pricing Card Configuration

| Parameter | Description |
|---|---|
| Name and Description | 1 Enter a name and description for your pricing card.<br>2 Optional: Select Default for Unassigned Workloads.<br>3 Click **Next**.<br>Default pricing card applies to all vCenter resources which do not have a direct cost policy assigned to them. |
| Basic Charges | Select the type of pricing card. Follow the steps for cost-based pricing card.<br>1 Enter the Cost Factor for the following.<br>a CPU Cost<br>b Memory Cost<br>c Storage Cost<br>d Additional Cost<br>2 Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**.<br>3 Select how to charge for the resources, the options are **Always** or **Only When Powered On**.<br>4 Click **Next**.<br>**Note** Cost - The cost is defined in . If selected, a multiplication factor is required. For example, if you select 1.1 as a factor, the cost is multiplied by 1.1 resulting in a 10% increase to the calculated cost. The price equation using cost is: <cost> x <multiplication factor> = Price<br>Follow the steps for rate-based pricing card.<br>1 Enter the CPU Rate in MHz per vCPU.<br>2 Enter the Memory Rate per GB.<br>3 Enter Storage Rate per GB.<br>4 Select the ChargingPeriod for all the values.<br>5 Select the Charge On Power State for all the values. |
| Guest OSes | 1 Enter the Guest OS Name.<br>2 Enter the base rate.<br>3 Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**. |

Table 7-5. Pricing Card Configuration (continued)

| Parameter | Description |
| --- | --- |
| Tags | Enter the Tag name and Tag Value. Define the charging method and base rate.<br>■ Recurring - enter a base rate and define recurring interval as the charge period. The absolute rate value is required and it is added to the overall price.<br>■ One time - define the one-time base rate charge. The absolute value is required and it is added as a one time price.<br>■ Rate Factor - A multiplication factor is required that is applied to the select charge category.<br>Select how to charge the Tag based on powered on state. |
| Overall Charges | You can define overall charges to VMs that match this policy.<br>1 Enter the VM setup charges.<br>2 Enter the Recurring charge and select the time period from the drop-down menu. |
| Assignments | You can assign the new pricing card to vCenters and Clusters.<br>1 Select the vCenter or Cluster to which you want to apply the pricing card.<br>2 Click **Add** and Click **Finish**. |

Results

The new pricing card details are displayed in the Pricing tab.

# Migration of vCenter Pricing Cards to vCenter Pricing Policies

When you upgrade to this version and later versions of VMware Aria Operations, and you have vCenter Pricing Cards, you can migrate them to VMware Aria Operations policies to utilize the full capability of VMware Aria Operations policies.

## Steps to Migrate vCenter Pricing Cards to vCenter Pricing Policies

1 From the left menu, click **Configure > Cost Drivers**, and then click the **Pricing** tab.

2 From the text box, click the **Actions** drop-down and then click **Migrate Now**.

After pricing card migration is complete, the following scenarios are available:

- **Pricing Policies Migration:** All the pricing rate cards will be converted to separate policies under the default policy. The converted policies will have the same name as the rate cards with the rate card UUID added as a suffix.
- **Pricing Cards Assignment Migration:**
  - If there is just one active policy before migration, then all the rate card assignments will be transferred to the newly created policies.

- If there is more than one active policy you will have to manually reassign policies to the relevant objects after the migration. It is recommended to proceed with policy assignments immediately after the migration to ensure continuous pricing calculations.

  The rate cards will no longer be available. The new updated capability will be accessible under each policy under Policy Definition in VMware Aria Operations.

- Default rate card configuration will be migrated to the policy. However, it will not be set as the default. You can achieve this by either making it the default policy or by utilizing VMware Aria Operations Custom Groups capabilities.

After migration of the vCenter pricing cards, you can view the **VC Pricing Migration Audit** logs from **Migration > Audit**, and then click the **VC Pricing Migration Audit** tab. For more details, see VC Pricing Migration Audit.

# Cost Calculation Status Overview

You can check the ongoing status of manually triggered cost calculation process.

Cost calculation by default, occurs daily and whenever there is a change in the inventory or cost drivers values. You can trigger the cost calculation manually so that changes in the inventory and cost driver values reflect accordingly on the VM cost without having to wait there for any failures in the cost calculation process. It also shows default schedules time for next cost calculation process.

**Note** To run the manual cost calculation, from the left menu click **Administration** > **Cost Calculation** and click **Run**.

# Costing Enhancements

In VMware Aria Operations, a new global property Cluster Utilization Ceiling Factor is introduced. Using Cluster Utilization Ceiling Factor, you can specify the ceiling value and calculate the base rate for a cluster.

You can use the ceiling factor only if the base rate cost calculation is done using Cluster Actual Utilization method. After you set the ceiling factor value, the Actual Utilization of the cluster is rounded off to the next available multiple of the ceiling value. When ceiling value is 0, Expected Utilization is equal to actual utilization. When ceiling value is 20, it is not considered as special case, actual utilization is rounded off to the next multiple.

**Note** The ceiling value range is from 0 to 20. If the number is out of this range, the default value of five is used as the ceiling number.

## How to Set the Cluster Base Rate Calculation Method

To change the Cluster Base Rate Calculation method, you must go to **Configure** > **Cost Drivers** > **Cluster Cost** page. Click **Change** next to the Cluster Base Rate calculation method and select Cluster Actual Utilization.

## Where to Find Cluster Utilization Ceiling Factor

To set the ceiling value for a cluster, you must go to **Administration** > **Global Settings** > **Cost/Price**>**Cluster Utilization Ceiling Factor**. Enter the ceiling value between 0 and 20 and click **Save**.

To view the change is cost metrics, run the Cost Calculation Status and select a cluster .

If the Actual Utilization of the cluster for CPU is 30 % and Memory is 45%, and the ceiling value specified is 10, then

- Cluster Expected CPU Utilization (%) = 40
- Cluster Memory Expected Utilization (%) = 50

Actual Cluster Utilization is rounded off to the ceiling value.

If you set the Cluster Utilization Ceiling Factor to either 0 or 20, then the value of Expected Memory Utilization changes to the next number. For example, if you set the ceiling factor to 0 then, the expected utilization value changes to 1.

## Support to Roll up Name Space Cost Metrics

The cost metrics of Point of Delivery (Pod) virtual machines (VMs) has been enhanced to support the following scenarios:

- Cost metrics of Pod VMs are rolled up to the Name Space and Guest Cluster level.
- All the cost metrics of VMs, Pods, and guest cluster which are present under Name Space are rolled up to Name Space and Guest Cluster level.

| Old Cost Metrics | Rolled up Cost Metrics |
|---|---|
| Effective MTD Total Cost | Aggregate Additional Daily Cost |
| Deleted VM Daily Cost | Aggregate Deleted VM Daily Cost |
| Daily CPU Cost | Aggregate CPU Daily Cost |
| Daily Memory Cost | Aggregate Memory Daily Cost |
| Daily Storage Cost | Aggregate Storage Daily Cost |
| Daily Additional Cost | Aggregate Additional Daily Cost |

## Reclaimable Hosts Cost Metric

You can use the cost metrics at cluster level to identify the clusters with reclaimable hosts and the potential cost savings from reclaiming these hosts. To know the cost associated with all the reclaimable hosts in a cluster, check the value of Total Host Reclaimable Host Cost metric.

### How to View Reclaimable Host Cost

To view the reclaimable host cost, go to **Environment** > **Object Browser** >**All Objects** > **vCenter Adapter** > **Cluster Compute Resource** > **Cost**.

You can also view the total host reclaimable cost using **Environment** > **Object Browser** > **All Objects** > **vCenter Adapter** > **vSphere World** > **Metrics** > **Cost**.

**Note** If the cluster does not have reclaimable hosts, then the cost metric associated with the reclaimable host is not displayed.

## Realized Cost Savings Using Reclamation Suggestion

In VMware Aria Operations you can track the cost savings using reclamation suggestions. Using the reclamation option, you can view the cost, capacity, and allocation metrics related to individual data centers. The metrics provides an estimate of the potential savings achieved through VMware Aria Operations.

You can track the realized cost savings and actual capacity reclaimed for data centers, in the following scenarios.

- Reclaim the cost for Idle VMs by deleting the VM.
- Reclaim the cost for Powered off VMs by deleting the VM.
- Reclaim the cost for Idle VMs by powering off the VM.
- Reclaim the cost for snapshots VMs by deleting the snapshot.
- Reclaim the cost for orphaned disks by deleting the orphaned disk space.
- Reclaim the cost by removing vCPU and Memory from an oversized VM.
- Reclaim the cost by removing a host from the vCenter.

## Costing for Oversized VM and Undersized VM

Rightsizing is defined as changing the amount of resources allocated to a VM based on the Recommended Size for a VM. Recommended Size is the maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. If the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization 150 days into the future. While rightsizing a VM can reclaim capacity, the change in allocation may not equal the amount of reclaimable capacity.

### Quantifying the Effect on Capacity Due to Rightsizing

#### Demand Model

- **Reclaimable CPU Usage (GHz)**: If an oversized VM's CPU usage is 100MHz before rightsizing, removing vCPU's will not change its CPU usage and it should still be at 100MHz. This means there is no reclaimable capacity associated with overallocation of vCPUs. Reclaimable CPU Usage for oversized VM's will always be 0 MHz.

- **Reclaimable Memory Consumed (GB)**: An oversized VM can have reclaimable memory only if consumed memory is greater than the new recommended size of the VM. The reclaimable memory capacity is the difference between consumed memory and recommended size.
- **Increased CPU Usage (GHz)**: CPU usage of an undersized VM is expected to be the current CPU Demand. The difference between CPU Demand and CPU Usage is the expected increase in capacity utilized after rightsizing.
- **Increased Memory Consumed (GB)**: It can be expected for consumed memory to increase by the same amount of memory recommended to add to an undersized VM.

**Allocation Model**

In case of allocation model, you can directly pick the recommendation provided which is given as a part of the metric groups **Summary|Oversized** and **Summary|Undersized**.

## Potential Cost Savings Calculation Detail

- **Oversized CPU Utilization**: $0 since Reclaimable CPU Usage (GHz) is always 0.
- **Oversized Memory Utilization**: Reclaimable Memory Consumed (GB) * Cluster Memory Base Rate.
- **Oversized CPU Allocation**: vCPU(s) to Remove * Allocation Cluster CPU Base Rate.
- **Oversized Memory Allocation**: Memory to Remove * Allocation Cluster Memory Base Rate.

## Potential Cost Increase Calculation Detail

- **Undersized CPU Utilization**: Increased CPU Usage (GHz) * Cluster CPU Base Rate.
- **Undersized Memory Utilization**: Increased Memory Consumed (GB) * Cluster Memory Base Rate.
- **Undersized CPU Allocation**: vCPU(s) to Add * Allocation Cluster CPU Base Rate.
- **Undersized Memory Allocation**: Memory to Add * Allocation Cluster Memory Base Rate.

The rightsizing value calculated here is available as part of

- **Potential Savings** metric (For VM) for Oversized VMs.
- **Potential Increase** metric (For VM) for Undersized VMs.

**Note** Reclaimable memory consumed, Increased CPU Usage, and Increased memory consumed are metrics that are available for reference under **Summary|Oversized** metrics and **Summary|Undersized** metrics respectively.

# Viewing and Configuring Compliance

8

VMware Aria Operations helps you visualize the compliance of the objects in your inventory by displaying compliance benchmarks which are measured against a set of standard rules, regulatory best practices, or custom alert definitions. You can discover non-compliant objects and the associated compliance violations and take remedial measures.

Read the following topics next:

- Measuring Compliance of Objects
- How To Configure Compliance Benchmarks

## Measuring Compliance of Objects

VMware Aria Operations continuously monitors your infrastructure to ensure that it remains in compliance. Compliance is an ongoing process. VMware Aria Operations evaluates the collected data against the defined policies. It assigns a compliance score to each object or group based on how well they adhere to the policies. The compliance score is typically represented as a percentage.

### How Compliance Benchmarks Work

Compliance is used to monitor the vCenter Server instances, hosts, virtual machines, distributed port groups, and distributed switches in your environment to ensure that the settings on your objects meet the defined standards. Compliance benchmarks display score cards that help you proactively detect compliance problems in VMware Aria Operations. The compliance benchmarks are measured against a set of standard rules, regulatory best practices, or custom alert definitions.

All the compliance standards in VMware Aria Operations, including any standards that you define, are based on alert definitions. Only alert definitions of the Compliance subtype are counted. Custom score cards can monitor user-defined alerts. The alerts and symptom definitions are based on the properties and metrics of the underlying object.

When VMware Aria Operations detects non-compliance with a policy, it can generate alerts or notifications. Depending on the severity of the non-compliance, you can configure automated remediation actions to bring the object back into compliance.

You can manage all compliance related tasks from the **Optimize > Compliance** page. The data sources are displayed in a carousel on the top of the page. To see a compliance score card, you must first configure the data source that VMware Aria Operations can monitor, and then activate the benchmarks for those types of data sources. When you activate a benchmark for a data source, you select an applicable policy. VMware Aria Operations then activates the appropriate alert definitions in the policy to measure compliance.

## Data Sources for Calculating Compliance

VMware Aria Operations can measure compliance from different data sources to ensure that your virtualized environment adheres to predefined policies and standards. These data sources provide the information necessary to evaluate the compliance of your infrastructure. The data sources that VMware Aria Operations can use for compliance measurement are as follows:

- VMware Self-Managed Cloud (SDDC) environment, including DC and Edge environments
- VMware Managed Cloud (VMC SDDC) environment
- VMware Cloud Foundation Domains
- VMware Cloud on Dell EMC SDDC
- Oracle Cloud VMware Solution SDDC
- Azure VMware Solution
- Google Cloud VMware Engine

- Private Cloud

**Note** For VMware Aria Operations to measure compliance against these data sources, you must first configure them. See the relevant topic in the Chapter 4 Integrating Data Sources with VMware Aria Operations chapter of the *Configuring VMware Aria Operations* guide.

**Note** In version 8.16, the symptom set for the FISMA Security Standards, DISA Security Standards and vSphere compliance pack were updated to generate an alert when the following conditions are false:

- Block Override Allowed should be true
- Port Config Reset at Disconnect should be true

**Note** In version 8.16, a bug involving DVPG symptoms on Promiscuous mode, MAC address changes, Forged transmits not considering the uplink/non-uplink state was fixed.

The symptoms on DVPG properties `allow_promiscuous`, `forged_transmits` and `mac_changes` are defined in the following compliance packs:

- CIS Security Standards
- DISA Security Standards
- FISMA Security Standards
- HIPAA
- ISO Security Standards
- PCI DSS Compliance Standards
- vSphere Security Configuration Guide

Compliance benchmarks on VMware Cloud on AWS, VMware Cloud Foundation, VMware Cloud on Dell EMC, Oracle Cloud VMware Solution, Azure VMware Solution, and Google Cloud VMware Engine are applicable only on customer VMs that you have deployed in the respective data centers.

You can automate the remediation of some of the alerts by installing the Management Pack for VMware Aria Automation Orchestrator. See the management pack documentation in the *VMware Aria Operations for Integrations Product Documentation* for more details.

## Compliance Benchmarks

### VMware SDDC and Benchmarks

Displays score cards based on alerts which are measured against the latest hardening guides:

- vSphere Security Configuration Guide
- vSAN Security Configuration Guide
- NSX Security Configuration Guide

For more details, see VMware SDDC Benchmark Details.

**VMware Cloud Foundation Benchmarks**

Displays score cards based on alerts which are measured in VMware Cloud Foundation domains based on the following audit guides:

- VCF 4.2 Audit Guide
- VCF 4.3 Audit Guide
- VCF 4.4 Audit Guide
- VCF 4.5 Audit Guide

The alerts are based on objects in your VMware Cloud Foundation environment.

For more details, see the topic, VMware Cloud Foundation Benchmarks based on VMware Cloud Foundation Compliance Kits.

**Custom Benchmarks**

Displays benchmarks that you define. Use compliance alerts from vSphere and regulatory management packs, or define your own alerts to monitor. You can define up to five custom score cards. You can import custom benchmarks from other instances of VMware Aria Operations.

**Regulatory Benchmarks**

Displays benchmarks for industry standard regulatory compliance requirements. You can install compliance packs for the following regulatory standards:

- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry Data Security Standard (PCI DSS) Compliance Standards
- CIS Security Standards
- Defense Information Systems Agency (DISA) Security Standards
- The Federal Information Security Management Act (FISMA) Security Standards
- International Organization for Standardization (ISO) Security Standards

For more details, see Regulatory Benchmark Details.

For instructions on installing these compliance packs, see Install a Regulatory Benchmark.

## VMware SDDC Benchmark Details

VMware SDDC (Software-Defined Data Center) security benchmarks are guidelines or standards that provide recommendations and best practices for securing VMware's Software-Defined Data Center infrastructure. These benchmarks help organizations assess and enhance the security of their virtualized data center environments.

## vSphere Security Configuration Guide

For details of the conditions that allow automated compliance assessment and the list of controls that can be used to perform manual checks based on the VMware vSphere 7 Security Configuration Guide Update 3 and VMware vSphere 8 Security Configuration Guide, see KB 88721.

The vSphere Security Configuration Guide is a comprehensive document provided by VMware that outlines best practices and recommendations for securing VMware vSphere, which is a virtualization platform used for creating and managing virtualized data centers.

The guide covers various aspects of vSphere security and provides configuration guidelines to help administrators protect their vSphere infrastructure from potential threats and vulnerabilities. It addresses security considerations for different components of the vSphere environment, including ESXi hosts, vCenter Server, virtual machines, networking, and storage.

The vSphere Security Configuration Guide typically covers the following areas:

- ESXi Host Hardening: It provides recommendations for securing the ESXi hosts, including configuring authentication and authorization settings, disabling unnecessary services, enabling security features like Secure Boot and Trusted Platform Module (TPM), and implementing network security measures.
- vCenter Server Hardening: The guide offers best practices for securing the vCenter Server, which is a centralized management platform for vSphere. It covers securing the vCenter Server installation, securing access and authentication, configuring roles and permissions, and enabling auditing and logging.
- Virtual Machine Security: It provides recommendations for securing virtual machines (VMs), such as implementing secure VM configurations, using virtual machine encryption, enabling guest operating system security features, and protecting VMs from unauthorized access.
- Networking Security: The guide offers guidelines for securing the vSphere networking infrastructure, including configuring virtual switches, VLANs, and firewalls, implementing network segmentation, and securing network communication between vSphere components.
- Storage Security: It covers best practices for securing the vSphere storage environment, including securing access to storage devices, implementing storage encryption, and protecting virtual machine data.
- Monitoring and Auditing: The guide emphasizes the importance of monitoring and auditing vSphere components, and provides recommendations for enabling and configuring logging, using security information and event management (SIEM) tools, and monitoring for suspicious activities.

The vSphere Security Configuration Guide is regularly updated by VMware to incorporate the latest security recommendations and considerations. It serves as a valuable resource for administrators seeking to secure their vSphere infrastructure and protect their virtualized environments from potential security risks.

## vSAN Security Configuration Guide

The vSAN Security Configuration Guide is a document provided by VMware that offers guidance and best practices for securing VMware vSAN (Virtual SAN), which is a software-defined storage solution integrated into the vSphere environment.

The guide focuses on helping administrators secure the vSAN infrastructure and protect the data stored within it. It provides recommendations for various security considerations related to vSAN, including authentication, encryption, network security, and access controls.

Here are some key areas covered in the vSAN Security Configuration Guide:

- Authentication and Authorization: The guide provides recommendations for configuring strong authentication mechanisms for vSAN components, such as vCenter Server and ESXi hosts. It suggests using secure protocols, enforcing secure password policies, and implementing multi-factor authentication where possible. It also covers role-based access control (RBAC) and permissions management.
- Encryption: vSAN supports encryption at rest, which ensures that data stored on vSAN disks is protected. The guide offers guidance on configuring vSAN encryption, including selecting the appropriate encryption algorithm, managing key management servers (KMS), and enabling encryption for data at rest.
- Network Security: It provides recommendations for securing network communication within the vSAN environment. This includes using secure protocols, enabling network encryption (VMkernel Encryption), and configuring firewall rules to control network traffic between vSAN components.
- Auditing and Logging: The guide emphasizes the importance of monitoring and auditing vSAN operations. It offers recommendations for enabling and configuring logging, implementing log management solutions, and monitoring vSAN events for potential security incidents.
- Secure Configuration Settings: It provides a list of vSAN-specific configuration settings that can enhance security. This includes recommendations for enabling features like Data-at-Rest Encryption, Secure Erasure, and ensuring proper network configurations.
- Compliance and Hardening: The guide addresses compliance considerations and provides information on how vSAN aligns with various security frameworks, such as the Center for Internet Security (CIS) benchmarks. It also offers recommendations for hardening vSAN components, including ESXi hosts and vCenter Server.

The vSAN Security Configuration Guide is regularly updated by VMware to incorporate the latest security best practices and recommendations. It serves as a valuable resource for administrators who want to ensure the security of their vSAN deployments and safeguard their storage infrastructure and data.

## NSX Security Configuration Guide

The VMware Aria Operations Compliance Pack for NSX-T is updated to support the following standards:

- NSX-T 3.2 Security Configuration Guide
- NSX-T 3.1 Security Configuration Guide
- NSX-T 3.0 Security Configuration Guide
- **Note** NSX-T is renamed to NSX.

For more details, see KB 93136.

The NSX Security Configuration Guide is a comprehensive document provided by VMware that offers guidance and best practices for securing VMware NSX, which is the next-generation network and security platform for software-defined networking.

The NSX Security Configuration Guide focuses on helping administrators secure the NSX infrastructure and protect network traffic, ensuring the confidentiality, integrity, and availability of data. It provides recommendations for various security considerations related to NSX, including authentication, authorization, network security, encryption, and compliance.

Here are some key areas covered in the NSX Security Configuration Guide:

- NSX Component Hardening: The guide provides recommendations for securing NSX components such as NSX Manager, NSX Controllers, and NSX Edge nodes. It includes guidelines for configuring strong authentication, implementing role-based access control (RBAC), and securing administrative access to NSX components.
- Network Security: It covers best practices for implementing network security within NSX. This includes recommendations for creating and managing security groups, implementing distributed firewalling, configuring security policies, and enforcing microsegmentation to control network traffic.
- Authentication and Authorization: The guide offers recommendations for configuring strong authentication mechanisms for NSX, such as integrating with identity providers, enabling multi-factor authentication, and defining access controls based on user roles and permissions.
- Secure Network Communication: It provides guidance for securing network communication within the NSX environment. This includes configuring Transport Layer Security (TLS) encryption for communication between NSX components, securing communication between NSX and vCenter Server, and protecting east-west and north-south traffic.
- Virtual Machine Security: The guide provides recommendations for securing virtual machines (VMs) deployed in NSX environments. This includes implementing security groups, defining security policies, and leveraging NSX's integration with third-party security solutions for advanced VM protection.

- Logging and Auditing: It emphasizes the importance of logging and auditing in NSX environments. The guide offers recommendations for enabling and configuring logging, integrating NSX with logging and monitoring solutions, and monitoring for security events and anomalies.
- Compliance and Hardening: The guide addresses compliance considerations and provides information on how NSX aligns with various security frameworks and regulations. It also offers recommendations for hardening NSX components and ensuring compliance with security policies.

The NSX Security Configuration Guide is regularly updated by VMware to incorporate the latest security best practices and recommendations. It serves as a valuable resource for administrators who want to ensure the security of their NSX deployments, implement effective network security controls, and protect their software-defined networking infrastructure.

## Regulatory Benchmark Details

Regulatory compliance benchmarks are standards or guidelines that help organizations measure and assess their level of compliance with applicable laws, regulations, and industry standards.

### Health Insurance Portability and Accountability Act (HIPAA)

The Health Insurance Portability and Accountability Act (HIPAA) is a U.S. federal law enacted in 1996. HIPAA establishes standards and regulations to protect the privacy, security, and confidentiality of individuals' personal health information (PHI) and electronic health records (EHRs) in the healthcare industry.

The HIPAA Privacy Rule and Security Rule are two key components of the HIPAA standard:

- HIPAA Privacy Rule: The Privacy Rule sets standards for the use and disclosure of PHI by covered entities, which include healthcare providers, health plans, and healthcare clearinghouses. It grants individuals certain rights over their health information, such as the right to access, request amendments, and obtain an accounting of disclosures. Covered entities are required to implement safeguards to protect PHI, provide patients with notice of privacy practices, and obtain written authorization for certain uses and disclosures of PHI.
- HIPAA Security Rule: The Security Rule establishes security standards for protecting electronic PHI (ePHI). It requires covered entities and their business associates to implement administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and availability of ePHI. These safeguards include risk assessments, access controls, encryption, audit controls, disaster recovery plans, and employee training on security awareness.

The HIPAA standard applies to healthcare providers, health plans, healthcare clearinghouses, and their business associates who handle PHI or ePHI. Compliance with HIPAA regulations is mandatory, and non-compliance can result in significant penalties, including financial fines and potential criminal charges.

HIPAA also includes provisions related to the electronic exchange of health information and establishes the Health Information Technology for Economic and Clinical Health (HITECH) Act, which promotes the adoption and meaningful use of electronic health records.

It's important to note that my knowledge cutoff is in September 2021, so it's advisable to consult the latest information and updates from official sources such as the U.S. Department of Health and Human Services (HHS) for the most current and accurate information regarding HIPAA standards and regulations.

## Payment Card Industry Data Security Standard (PCI DSS) Compliance Standards

The Payment Card Industry Data Security Standard (PCI DSS) is a set of compliance standards established by the major payment card brands, including Visa, Mastercard, American Express, Discover, and JCB. PCI DSS aims to ensure the security of cardholder data and protect against fraud and unauthorized access within the payment card industry.

The PCI DSS compliance standards consist of twelve high-level requirements, organized into six control objectives. These requirements outline security measures that organizations handling payment card data must implement:

- Build and Maintain a Secure Network
- Protect Cardholder Data
- Maintain a Vulnerability Management Program.
- Maintain an Information Security Policy

PCI DSS compliance requirements vary based on the organization's level of involvement in payment card transactions, classified as levels 1 to 4. Level 1 merchants and service providers with the highest transaction volumes have more stringent requirements and undergo annual on-site audits by a Qualified Security Assessor (QSA). Level 2, 3, and 4 merchants may have different validation requirements, ranging from self-assessment questionnaires (SAQs) to external vulnerability scans.

Compliance with PCI DSS is necessary for entities involved in payment card processing, including merchants, service providers, and payment processors. Non-compliance may result in penalties, fines, increased transaction fees, or restrictions on card acceptance.

It's important to note that while this information provides an overview of the PCI DSS compliance standards, specific requirements and guidance can evolve over time. Therefore, it is recommended to consult the official PCI Security Standards Council (PCI SSC) website and the latest PCI DSS documentation for the most up-to-date information and requirements.

## CIS (Center for Internet Security) Security Standards

The VMware Aria Operations Compliance Pack for CIS is updated to support the following benchmarks:

- CIS_VMware_ESXi_6.7_Benchmark_V1.3.0
- CIS_VMware_ESXi_7.0_Benchmark_V1.2.0

For more details, see the KB 93135.

The CIS (Center for Internet Security) Security Standards are a set of best practices and guidelines for securing computer systems and networks. The CIS organization is a non-profit entity that collaborates with experts from various industries to develop and promote consensus-based security configurations and benchmarks.

The CIS Security Standards include two primary components:

- CIS Controls: The CIS Controls are a set of 20 security actions that organizations can take to mitigate the most common and impactful cyber threats. These controls are prioritized based on their effectiveness in reducing risk. They cover various security domains, including asset management, access control, incident response, network security, and security awareness training. The CIS Controls are regularly updated to address emerging threats and evolving technology landscapes.
- CIS Benchmarks: CIS Benchmarks provide detailed configuration guidelines for securing specific technology platforms and systems. These benchmarks outline recommended settings and configurations for operating systems, applications, and network devices to ensure security and reduce vulnerabilities. CIS Benchmarks are created through a consensus-driven process involving input from cybersecurity experts, vendors, and practitioners.

CIS Security Standards are known for their practical and actionable nature, providing step-by-step instructions and specific configuration recommendations. They are widely adopted across industries and are used as a reference by organizations to assess, improve, and maintain the security of their IT systems and networks.

The CIS organization regularly updates its security standards and benchmarks to address emerging threats, technology advancements, and changes in regulatory requirements. The CIS Security Standards are available to the public, and organizations can leverage them as a valuable resource for enhancing their cybersecurity posture and reducing the risk of cyberattacks.

## Defense Information Systems Agency (DISA) Security Standards

The Defense Information Systems Agency (DISA) establishes and provides security standards and guidelines for the U.S. Department of Defense (DoD) and its information systems. DISA is responsible for ensuring the secure operation and defense of DoD's global information infrastructure.

DISA has developed several security standards and guidelines to protect sensitive information and ensure the integrity, availability, and confidentiality of DoD systems. Some of the key security standards and guidelines provided by DISA include:

- Security Technical Implementation Guides (STIGs): STIGs are a set of guidelines and configuration standards for various operating systems, applications, and network devices. They provide detailed instructions on how to secure and configure these systems to meet DoD security requirements. STIGs cover a wide range of technologies, including Windows, Linux, Cisco devices, databases, and web servers.

- Security Requirements Guides (SRGs): SRGs are comprehensive documents that outline security requirements for specific technology platforms, systems, or applications. They provide guidance on how to secure and configure systems in accordance with DoD security policies. SRGs address various security domains, including access control, identification and authentication, audit and accountability, and encryption.
- Security Technical Implementation Guides (STIGs) Viewer: DISA provides a STIG Viewer tool that helps organizations assess and implement STIG recommendations. The STIG Viewer automates the process of checking system configurations against STIG requirements, allowing organizations to identify and remediate security vulnerabilities more efficiently.
- Information Assurance Vulnerability Management (IAVM): DISA maintains the IAVM program, which identifies and manages vulnerabilities in DoD systems. IAVM alerts provide timely information about security vulnerabilities and patches. Organizations within the DoD are required to promptly apply these patches to mitigate potential risks.
- DoD Cybersecurity Discipline Implementation Plan (CDIP): The CDIP outlines the implementation and management of cybersecurity practices within the DoD. It provides guidelines and best practices for managing risks, protecting systems, responding to incidents, and fostering a culture of cybersecurity awareness.

DISA's security standards and guidelines play a critical role in ensuring the security and resilience of DoD systems and information assets. They are constantly updated and refined to address emerging threats and align with evolving cybersecurity practices. Organizations within the DoD are expected to adhere to these standards to maintain the security of their systems and networks.

## The Federal Information Security Management Act (FISMA) Security Standards

The Federal Information Security Management Act (FISMA) is a U.S. federal law enacted in 2002. FISMA establishes a framework for securing information systems and managing cybersecurity risks within federal government agencies and their contractors. FISMA requires federal agencies to develop, implement, and maintain information security programs to protect sensitive government information.

While FISMA itself does not provide detailed security standards, it sets requirements for federal agencies to follow certain security guidelines and standards, including those established by the National Institute of Standards and Technology (NIST). NIST Special Publication (SP) 800-53, titled "Security and Privacy Controls for Federal Information Systems and Organizations," is a key document referenced under FISMA.

NIST SP 800-53 provides a catalog of security controls that federal agencies must implement to protect their information systems. The controls cover various areas, including access control, incident response, configuration management, encryption, network security, and security assessment and authorization. The controls are categorized into families and are tailored to address specific security requirements.

FISMA requires federal agencies to develop and maintain a risk-based approach to information security. This involves conducting risk assessments, implementing security controls based on the identified risks, periodically testing and evaluating the effectiveness of these controls, and ensuring continuous monitoring of information systems.

Under FISMA, federal agencies are also required to undergo annual security assessments, including independent audits, to evaluate the effectiveness of their information security programs and controls. The results of these assessments are reported to the Office of Management and Budget (OMB) and Congress.

FISMA compliance is crucial for federal agencies to demonstrate their commitment to protecting government information and ensuring the security of their information systems. It helps establish a standardized approach to information security across federal government entities and aligns with other security frameworks and standards, such as the NIST Cybersecurity Framework and NIST Risk Management Framework.

It's important to note that FISMA requirements may evolve over time, and agencies should refer to the latest guidance provided by NIST and other authoritative sources to ensure compliance with FISMA security standards.

## International Organization for Standardization (ISO) Security Standards

The International Organization for Standardization (ISO) is an independent, non-governmental international standardization body that develops and publishes international standards across various industries. ISO has also created a series of security standards specifically related to information security management systems (ISMS). The most well-known among them is ISO/IEC 27001.

ISO/IEC 27001: The ISO/IEC 27001 standard specifies the requirements for establishing, implementing, maintaining, and continuously improving an ISMS within the context of an organization. It provides a systematic and risk-based approach to managing the security of sensitive information. The standard covers areas such as risk assessment, information security policies, asset management, access control, incident management, and compliance. ISO/IEC 27001 is widely adopted by organizations globally and serves as a benchmark for information security management.

ISO/IEC 27002: ISO/IEC 27002 (formerly known as ISO/IEC 17799) is a code of practice for information security controls. It offers guidance and recommendations for implementing security controls and safeguards based on the best practices of information security management. ISO/IEC 27002 covers a broad range of security areas, including organizational security, human resource security, physical and environmental security, communications and operations management, and compliance.

ISO/IEC 27005: ISO/IEC 27005 provides guidelines for conducting risk assessments in the context of information security. It offers a structured approach for identifying, analyzing, evaluating, and treating information security risks. ISO/IEC 27005 helps organizations assess the potential impact of risks, determine risk tolerance, and make informed decisions on implementing appropriate security controls.

ISO/IEC 27017 and ISO/IEC 27018: These standards specifically focus on cloud security. ISO/IEC 27017 provides guidelines for implementing information security controls in cloud computing environments, while ISO/IEC 27018 offers guidance for protecting personal data in the cloud and addresses privacy concerns related to cloud services.

ISO/IEC 27701: This standard is an extension to ISO/IEC 27001 and provides guidelines for implementing a Privacy Information Management System (PIMS). ISO/IEC 27701 helps organizations establish and maintain controls to protect personal data and comply with privacy regulations, such as the General Data Protection Regulation (GDPR).

ISO security standards provide a framework for organizations to establish effective information security management practices. Compliance with these standards demonstrates a commitment to securing sensitive information, managing risks, and implementing robust security controls. Organizations can seek ISO certification through a formal audit process conducted by accredited certification bodies to validate their adherence to ISO security standards.

## VMware Cloud Foundation Benchmarks based on VMware Cloud Foundation Compliance Kits

Security and Compliance for VMware Cloud Foundation provides guidance for auditors who would want to evaluate their VMware Cloud Foundation environment. You can leverage SDDC benchmarks and regulatory frameworks available under VMware Cloud Foundation to be able to assess individual SDDC components based on needed benchmarks or regulatory frameworks of your choice.

### VCF Audit Guide

Compliance Kit for VMware Cloud Foundation is a solution that builds on top of VMware Cloud Foundation and leverages security fundamentals. The kit addresses the top ten most frequently requested compliance standards, regulations, and frameworks.

The compliance kit is designed and validated to tailor security configurations without impacting the ability of VMware Cloud Foundation to meet its design objectives. The kit can assist organizations to secure information systems in a compliance context.

The VCF Audit Guide is part of the Compliance Kit for VMware Cloud Foundation and can be used to evaluate both default and non-default configurations. For more information, see the topic, Compliance Kit for VMware Cloud Foundation in the *VMware Cloud Foundation Product Documentation*.

The .xlsx file, that is based on the VCF Audit Guide, containing a collection of all compliance conditions required for assessing compliance of VCF environments is available in the KB article: KB94848.

Versions of the supported VMware Cloud Foundation Compliance kits are:

- VCF 4.2 Audit Guide
- VCF 4.3 Audit Guide
- VCF 4.4 Audit Guide

- VCF 4.5 Audit Guide

VMware Aria Operations supports the following products for VCF Benchmarks:

- ESXi,
- SDDC Manager,
- vCenter Server,
- vSAN
- NSX

Within these products, VMware Aria Operations can measure the compliance of the following resources:

- VMware Distributed Virtual Switch
- VMware Adapter Instance
- Virtual Machine
- Host System
- Logical Router
- NSX
- NSX Manager Service
- Logical Switch
- Management Cluster
- VCF Adapter Instance
- vSAN
- DC
- Cluster

**Note** Assessment of objects starting from version VMware Cloud Foundation 4.5 and above is done based on VCF 4.5 Audit Guide.

**Note** You must configure the VMware Cloud Foundation data source before you can configure the VCF benchmarks. For information, see the topic, VMware Cloud Foundation in the *Configuring VMware Aria Operations* guide.

## Compliance Score Cards

The compliance page in VMware Aria Operations displays score cards for each type of benchmark. A score card is a compliance visualization term.

## What is a Compliance Score Card

Score cards in the Compliance landing page display the number of non-compliant objects, and the total number of objects affected by each hardening guide and the compliance score which is counted as the ratio of compliant objects to total number of objects assessed by the given benchmark, represented in percentage. In addition, you can see the breakdown of the total number of objects that are compliant and non-compliant. You can click a score card to see more details, including alerts that were triggered based on the compliance standards.

The compliance score card of an object is counted as the smallest rounded off integer (100*(total number of symptoms triggered on an object/total number of symptoms)).

The compliance score for the object is based on the most critical of the violated standards. The score card displays 100 when all objects are compliant. When an object is non-compliant, the number of non-compliant symptoms are displayed in red and the total number of symptoms in grey.

**Note** The compliance score for a user with limited object visibility is the same as for a user with complete object visibility. This is because the compliance score is calculated for all objects, irrespective of whether the user has access to the object or not.

## Where You Find Compliance Score Cards

You can view score cards for each of the different types of benchmarks in the Compliance page available under **Optimize > Compliance**.

You can view score cards for objects in the **Environment > Object Browser > Compliance** tab.

### Compliance Page

In the **Optimize > Compliance** summary page, VMware Aria Operations monitors compliance for VMware SDDC benchmarks.

VMware Aria Operations displays compliance score cards in the following sections:

- VMware SDDC Benchmarks
- Custom Benchmarks
- Regulatory Benchmarks

### Compliance Tab

In the **Environment > Object Browser > Compliance** tab, VMware Aria Operations displays score cards for the benchmarks that include the current objects in their calculations, based on the alert definitions and policies associated with that benchmark. The score cards display the total number of rules and the number non-compliant (violated) rules based on symptoms for each hardening guide.

## Score Cards in the Compliance Page

In the **Optimize > Compliance** page, you can view scores for benchmarks that you have activated. Click a score card to view more information.

### Table 8-1. Compliance Page Score Card Options

| Item | Description |
|---|---|
| Score card for the configured hardening guides, custom benchmark and management packs | Displays the compliance score, total compliant and non-compliant objects for the compliance standards you have configured. |
| Object Breakdown | Displays the number of compliant and non-compliant objects for the following types of objects:<br>■ vCenter<br>■ ESXi Host<br>■ Virtual Machine<br>■ Distributed Port Group<br>■ Distributed Virtual Switch<br>■ vSAN Cache Disk<br>■ vSAN Capacity Disk<br>■ vSAN Cluster<br>■ NSX Manager<br>■ NSX-V EDGE<br>■ NSX-V Logical Router<br>■ NSX-V Manager<br>■ NSX-V Routing Edge Service |
| Compliance Alert List | A list of alerts, grouped by time by default. You can either remove the grouping of the alerts, or group by criticality, definition, and object type.<br>The alerts which caused the compliance violation are displayed in a table. You can sort the table by the following columns:<br>■ Alert ID<br>■ Criticality<br>■ Alert<br>■ Triggered On<br>■ Updated On<br>Click **Actions** to edit the policies applied to the configuration guide.<br>Click an alert to view more details. The **Environment > Object Browser > Alert** tab opens.<br>Use the advanced quick filter to search for alerts. Click the drop-down arrow for more options to narrow your search. |

## Compliance Alerts

You use the compliance score card as an investigative tool when you evaluate the state of objects in your environment, or when you research the root cause of a problem. If the score card indicates a problem, you can view the alerts to see details about the violation. Violated rules are based on the symptoms defined in the compliance alert.

The compliance alerts, which have the subtype named Compliance, include one or more symptoms that represent the compliance rules. Compliance alerts that are triggered appear on the **Environment > Object > Compliance** tab as violations to the standard, and the triggered symptoms appear as violated rules. The rules are the alert symptoms, and the symptom configuration identifies the incorrect value or configuration. If a rule symptom is triggered for any of the alerts in the standard, the triggered rule violates the standard and affects the score that appears on the **Environment > Object > Compliance** tab.

Table 8-2. Compliance Tab Alert Display

| Item | Description |
| --- | --- |
| Score card for the configured hardening guides | Displays the score card value, total number of rules, and number of non-compliance rules for the compliance standards you have configured. |
| Active Compliance Alerts | If you click the score card, the rules for the score card appear. When a symptom is triggered, the rule is considered to be violated. View the list of rules in the following tabs:<br>■ Violated Rules. Displays only the triggered symptoms. Click a symptom to view more information.<br>■ All Rules. Displays triggered and untriggered symptoms. |

# How To Configure Compliance Benchmarks

Configure VMware SDDC, custom, and regulatory benchmarks from the Compliance page. Unlike previous releases, you can now enable alert definitions in one of the active policies, from the Compliance page directly.

## Activate VMware SDDC Benchmarks

You can activate the VMware SDDC Benchmark to monitor objects for violation of vSphere Security Configuration Guide, vSAN Security Configuration Guide, NSX Security Configuration Guide (SDDC only). The score cards in the VMware SDDC Benchmark warn you when compliance alerts trigger on your vCenter Server instance, NSX-V objects, NSX objects, vSAN objects, ESXi hosts, virtual machines, distributed port groups, or distributed virtual switches.

### Prerequisites

You must configure the data source for which you are enabling the VMware SDDC benchmarks before you proceed. For information on how to configure a data source, see the relevant topic in

the Chapter 4 Integrating Data Sources with VMware Aria Operations chapter of the *Configuring VMware Aria Operations* guide.

### Procedure

1 From the left menu, click **Optimize > Compliance** to access the compliance page.

2 Select a data source from the carousel on the top of the page.

3 In the VMware SDDC Benchmarks section, click **Enable** under the vSphere Security Configuration Guide, vSAN Security Configuration Guide or NSX Security Configuration Guide card.

**Note** To activate the NSX Security Configuration guide, you must first install the NSX for vSphere, or the NSX-T solution. .

The **Enable Policies** dialog box opens.

4 Select the policy that you want to modify. When there are child policies, you can select a child policy and unselect a parent policy. VMware Aria Operations modifies the selected policy and activates the alert definitions associated with the current scorecard.

5 Click **Enable** to confirm your selection.

### Results

VMware Aria Operations starts to assess the objects based on the policy that you selected. To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

## Activate VCF Benchmarks

You can activate the VMware Cloud Foundation benchmarks to audit the compliance of the objects in the VMware Cloud Foundation stack. The products currently included for assessment are ESXi, SDDC Manager, vCenter Server, vSAN and NSX

### Prerequisites

You must configure the VMware Cloud Foundation data source before you proceed. For information, see the topic, VMware Cloud Foundation in the *Configuring VMware Aria Operations* guide.

### Procedure

1 From the left menu, click **Optimize > Compliance** to access the compliance page.

2 Select VMware Cloud Foundation from the carousel on the top of the page.

3 In the VCF Benchmarks section, click **Enable** in the **VCF Compliance based on the VCF Audit Guides** card.

The **Enable Policies** dialog box opens.

4 Select the policy that you want to modify. When there are child policies, you can select a child policy and unselect a parent policy. VMware Aria Operations modifies the selected policy and activates the alert definitions associated with the current scorecard.

5 Click **Enable** to confirm your selection.

#### Results

VMware Aria Operations starts to assess the objects based on the policy that you selected. To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

## Create a New Custom Benchmark

You can create a custom compliance benchmark to ensure that objects comply with compliance alerts available in VMware Aria Operations, or custom compliance alert definitions. When a compliance alert is triggered on your vCenter instance, hosts, virtual machines, distributed port groups, or distributed switches, you investigate the compliance violation. You can add up to five custom compliance score cards.

#### Prerequisites

To create a custom benchmark based on industry standard regulatory compliance requirements, you must first download and install the compliance management packs.

#### Procedure

1 From the left menu, click **Optimize > Compliance** to access the compliance page.

2 In the Custom Benchmarks section, click **Add Custom Compliance**.

The **Add Custom Compliance** dialog box opens.

3 Select **Create a New Custom Benchmark**.

a In the Name and Description step, provide a name and description for the custom benchmark and click **Next**.

b In the Alert Definitions step, select the compliance alerts that you want to add to this custom compliance benchmark and click **Next**.

c In the Policies step, select the policies to activate compliance and click **Finish**.

#### Results

The custom compliance which monitors alert definitions that you selected is available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit**.

## Import or Export a Custom Benchmark

You can export custom benchmarks from any VMware Aria Operations instance and import it to another instance. Reusing custom benchmarks saves you time and effort. You can modify

an imported custom benchmark. Exported files are in the XML format. The XML file contains information about alert groups, alerts, and filters.

### Prerequisites

You must first export a XML file with the custom benchmarks from another instance of VMware Aria Operations before importing the XML file to another instance.

### Procedure

1. From the left menu, click **Optimize > Compliance** to access the compliance page.
2. To import a custom benchmark, select either the SDDC or the VMC SDDC tab depending on where your objects are present.
3. In the Custom Benchmarks section, click **Add Custom Compliance**.

   The **Add Custom Compliance** dialog box opens.
4. Select **Import An Existing Custom Benchmark**.
   a. In the Import Compliance Score card dialog box, select the Score card Definition XML file from your local computer. If the XML file contains cloned alerts from the VMware Aria Operations instance that was used to export the file, the cloned alerts are also imported.
   b. VMware Aria Operations displays a message to indicate if the XML file was successfully imported.
   c. If you see a message which indicates that there is a conflict between the data in the XML file and the custom benchmarks already defined, make a selection on how to handle a conflict.
   d. Click **Done**.
5. To export an existing custom benchmark, click the score card to select the benchmark and select **Export** from the **Actions** menu.

### Results

The imported compliance benchmarks are available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit** from the **Actions** menu after clicking the score card.

## Install a Regulatory Benchmark

To enforce and report on the compliance of your vSphere objects, you activate the compliance pack that contains the policies for regulatory standards. Then, you select the policy to activate the appropriate regulatory alerts for your virtual machines.

### Procedure

1. From the left menu, click **Optimize > Compliance** to access the compliance page.

   The compliance packs for the regulatory standards are displayed under the Regulatory Benchmark section.

2. To install any regulatory benchmark, click **Activate From Repository** on the required compliance pack.

   You are redirected to the **Management Pack Details** page.

3. Click **Activate** to complete the installation.

4. To activate the compliance pack policies, navigate to the **Compliance** homepage and click **Enable** on the installed compliance pack.

   The **Enable Policies** window opens.

5. Select the policies that you want to activate and click **Enable** to complete the process.

### Results

VMware Aria Operations starts to assess the objects based on the regulatory benchmark that you installed.

# Configuring Green Score to Track Sustainability

9

VMware Green Score is used to measure sustainability in digital operations through data center virtualization. Configure Green Score in VMware Aria Operations to track where you are in your sustainability journey and how it is progressing over time.

You can configure green score at the Organization level and for physical data centers.

Read the following topics next:

- Configuring Green Score Sustainability Data at the Organization Level
- Configuring Green Score Sustainability Data for Physical Data Centers

## Configuring Green Score Sustainability Data at the Organization Level

Virtualization plays a crucial role in data center consolidation and thereby helps in the reduction of hardware footprint in the data centers. Virtualization helps in the saving of power consumption and floor space in data centers which improves the overall efficiency and curbs carbon emissions that result from IT infrastructure growth. Use this page to configure green score data in VMware Aria Operations. Using green score, organizations can track where they are in their sustainability journey and how they are progressing over time.

**Note** Ensure that the following permissions are assigned:

- Home > Change Organization Settings
- Home > View Sustainability Page

### Configuring Green Score for the First Time

From the **Home** page, click **Launchpad > Sustainability**. From the **Sustainability** page, click **Configure** from the Green Score panel. Enter the details in the **Organization** tab.

### Table 9-1. Organization Details

| Option | Description |
| --- | --- |
| **Green Score Component Input - Virtualization Ratio(%)** | |
| Compute Virtualization | Enter the percentage of compute services that are virtual. Default value is 80%. The default values are based on the VMware estimate of the most common numbers. |
| Network Virtualization | Enter the percentage of network services that are virtual. Default value is 30%. The default values are based on the VMware estimate of the most common numbers. |
| Storage Virtualization | Enter the percentage of storage technologies that are virtual. Default value is 40%. The default values are based on the VMware estimate of the most common numbers. |
| Desktop Virtualization | Enter the percentage of desktops that are virtual. Default value is 20%. The default values are based on the VMware estimate of the most common numbers. |
| **Power Sources (%)** | |
| Add Power Source | Click this option to add the share of power sources that you use. The total should add up to 100%.<br>The power sources are coal, oil, natural gas, hydroelectric, solar PV, wind, biomass, and nuclear.<br>The default value for coal is 79.2% and the default value for hydroelectric is 20.8%. The default values are based on the global averages from the Greenhouse Gas (GHG) Emissions Intensity of Electricity Generation methods.<br>The actual global average value is 709g/kWh. |
| **Average Age of Hardware in years (as of today)** | |
| Servers (years) | Enter the average age of the servers in your organization. Default age is two years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Storage (years) | Enter the average age of storage in your organization. Default age is two years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Network (years) | Enter the average age of network in your organization. Default age is three years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Desktop (years) | Enter the average age of the desktops in your organization. Default age is five years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |

After you click **Save**, the Sustainability data is calculated. Calculation of the sustainability data may take some time.

## Green Score Details

Green Score is calculated based on five components, which are: Workload Efficiency, Utilization of Physical resources, the extent of Virtualization, Power Source, and Hardware Efficiency. Each of these components has varying levels of impact on the carbon emissions within a data center and hence has a different weightage when contributing to a common Green Score.

- **Workload Efficiency**: Contributes to 22.5% weightage Efficiency = Wastage / (Wastage + Non Wastage).
- **Resource Utilization**: Contributes to 12.5% weightage based on the utilization of hardware from the absolute capacity without considering buffers and HA.
- **Virtualization**: Contributes to 15% weightage.
- **Power Source**: Contributes to 37.5% weightage. Scoring is based on the carbon intensity of the power source used.
- **Hardware Efficiency**: Contributes to 12.5% weightage, based on the fact that newer generation hardware will be more energy efficient.

Table 9-2. Green Score Details and Dependencies

| Name | Description |
| --- | --- |
| Greenscore Meter | Indicates how sustainable a company is. A higher score indicates a greater green score. Green score is calculated based on the different values calculated for workload efficiency, resource utilization, virtualization, power source, and hardware efficiency. |
| Greenscore Trend | Indicates the green score trend of the company. You can view daily, weekly, and monthly trends. |
| **Clean Demand** | |

## Table 9-2. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Workload Efficiency | Measures the wastage in the data center and calculates a score based on various waste sources. The workload efficiency value is calculated based on different metrics, such as powered off VMs, oversized VMs, idle VMs, snapshots, and orphaned disks. The metrics are as follows:<br>Total Consumption CPU = (vCenter) CPU\|vCPUs Allocated on all Powered On Consumers<br>Total Consumption Memory (GB) = (vCenter) Memory\| Host Usage (KB)<br>Total Consumption Disk (TB) = (vCenter) Disk Space\|Total Provisioned Disk Space (GB)<br>Oversized VM\|vCPU = (VM) Summary\|Oversized\|Virtual CPUs and Oversized VM\|Memory (GB) = (VM) Summary\| Oversized\|Memory (KB) - These metrics are used to calculate workload efficiency only for powered on VMs. This means that vCPUs and memory of powered off and oversized VMs are excluded from the calculation.<br>Idle VM\|Disk (TB) = (Cluster) Reclaimable\|Idle VMs\|Disk Space (GB)<br>Idle VM\|vCPU = (Cluster) Reclaimable\|Idle VMs\|CPU (vCPUs)<br>Idle VM\|Memory (GB) = (Cluster) Reclaimable\|Idle VMs\| Memory (KB)<br>Powered Off VM\|Disk (TB) = (Cluster) Reclaimable\| Powered Off VMs\|Disk Space (GB)<br>Snapshot\|Disk (TB) = (Cluster) Reclaimable\|VM Snapshots\| Disk Space (GB)<br>Orphaned VMDK\|Disk (TB) = (Datastore) Reclaimable\| Orphaned Disks\|Disk Space (GB) |
| **Lean Operations** | |

## Table 9-2. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Resource Utilization | Provides a value based on server and storage utilization. The score is based on the utilization of hardware from absolute capacity without considering buffers and HA. Here are the metrics used:<br>Total ESXi Hosts CPU Utilization (GHz) = (vCenter) CPU\|Usage (MHz)<br>Total ESXi Hosts CPU Capacity (GHz) = (vCenter) CPU\|Total Capacity (MHz)<br>Total ESXi Hosts Memory Utilization (TB) = (vCenter) Memory\|Host Usage (KB)<br>Total ESXi Hosts Memory Capacity (TB) = (vCenter) Memory\|Total Capacity (KB) Organization/Physical<br>Total datastores capacity = (vCenter) Disk Space\|Total Capacity (GB)<br>Total datastores utilization = (vCenter) Disk Space\|Utilization (GB)<br>Total RDM = Sum((VM) Virtual Disk \ Configured Size) |
| Virtualization | Virtualization score is calculated based on the virtualization percentage of compute, storage, network, and desktop entered in the **Organization Details** page. The formula is a weighted average of virtualization percentages with 40%, 30%, 20%, and 10% corresponding weights.<br>**Note** At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter Server instances that are not associated with any physical data center. |
| **Green Supply** | |

## Table 9-2. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Power Source | The score is based on the carbon intensity of the power source used. For example, traditional coal based power sources will have a carbon intensity rate of 888 g/kWh while wind based power source will have a carbon intensity of 26 only (Reference : WNA Report - Comparison of Lifecycle Greenhouse Gas Emissions of Various Electricity Generation Sources.<br>**Note**: A 100% badge score indicates zero carbon emission.<br>Each value is calculated based on the share of the power sources entered in the **Organization Details** page, after which the average value is calculated. The value for each of the power resources is calculated as follows:<br>Green Factor Shares (100%) Adjusted Score (Green Factor * Shares)<br>Coal = 17% (Green Factor)<br>Oil = 31% (Green Factor)<br>Natural Gas = 53% (Green Factor)<br>Solar PV = 92% (Green Factor)<br>Biomass = 96% (Green Factor)<br>Nuclear = 97% (Green Factor)<br>Hydroelectric = 98% (Green Factor)<br>Wind = 98% (Green Factor)<br>**Note** Shares are user inputs in the **Organization Details** page.<br>**Note** At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter Server instances that are not associated with any physical data center. |
| Hardware Efficiency | The hardware efficiency badge is based on the fact that new generation hardware is more energy efficient than the old generation hardware - by running more workloads on the same sized hardware and also the ability of hardware to optimize energy consumption. The efficiency of hardware which includes server, storage, network, and desktop is based on the average age of the hardware entered in the **Organization Details** page. For example, if you had entered the age of your servers as 2 years, then the hardware efficiency for servers would be 80%, because the useful life duration is considered as 10 years.<br>**Note** At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter Server instances that are not associated with any physical data center. |

## Power Consumption Chart

The power consumption chart displays the cumulative power consumption for the last one month. It is calculated using the power usage metrics and is the sum of the last month's daily power consumption. When you click the Power Consumption option, you can also view the **Power Consumption** trend which can be daily, weekly, or monthly.

## Carbon Footprint Chart

The carbon footprint chart displays the carbon emission for the last one month. When you click the Carbon Footprint option, you can also view the **Carbon Footprint** trend which can be daily, weekly, or monthly.

## Environmental Impact

The environmental impact is depicted in terms of carbon emissions from the number of smart phones charged, as per the calculations based on Greenhouse Gas Equivalencies Calculator. For detailed info on this please refer to https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references#smartphones.

For information about Sustainability dashboards, see Sustainability Dashboards.

# Configuring Green Score Sustainability Data for Physical Data Centers

Virtualization plays a crucial role in data center consolidation and thereby helps in the reduction of hardware footprint in the data centers. Virtualization helps in the saving of power consumption and floor space in data centers which improves the overall efficiency and curbs carbon emissions that result from IT infrastructure growth. Use this page to configure green score data in VMware Aria Operations for physical data centers that have been added to a vCenter Server account, a VMware Cloud on Dell EMC account, or to a VMware Cloud Foundation account and configured for green score. Using green score, organizations can track where they are in their sustainability journey and how they are progressing over time.

**Note** Ensure that the following permissions are assigned:

- Home > Change Organization Settings
- Home > View Sustainability Page

## Configuring Green Score for the First Time for Physical Data Centers

1 From the **Home** page, click **Launchpad** > **Sustainability**.

2 From the **Sustainability** page, click **Configure** from the Green Score panel. Enter the details in the **Organization** tab.

**Note** You must configure sustainability data at the Organization level and only then will you be able to configure sustainability for physical data centers. For information about configuring sustainability at the organization level, see Configuring Green Score Sustainability Data at the Organization Level.

3 Click the **Save** button in the **Organization** tab to save the Organization level details and to activate the **Physical Data Centers** tab.

4 From the **Physical Data Centers** tab, click the **Add Physical Data Center** button to add a physical data center to a vCenter Server account, a VMware Cloud on Dell EMC account, or to a VMware Cloud Foundation account if you have not previously added a physical data center from the **Administration > Physical Data Center** page. For more information about adding physical data centers, see Adding Physical Data Centers in VMware Aria Operations.

5 When you create a physical data center, the green score configuration page for that physical data center is displayed and you can fill in all the details.

**Note** If you have already added a added a physical data center from the **Administration > Physical Data Center** page, a list of physical data centers appears on the **Physical Data Centers** tab. You can select each of the physical data centers from the vertical ellipses called **Add Greenscore Configuration**.

6 Select a physical data center from the Physical Data Centers map/list and then enter the details in the **Physical Data Centers** tab.

Table 9-3. Physical Data Center Details

| Option | Description |
|---|---|
| **Green Score Component Input - Virtualization Ratio(%)** | |
| Compute Virtualization | Enter the percentage of compute services that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Network Virtualization | Enter the percentage of network services that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Storage Virtualization | Enter the percentage of storage technologies that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Desktop Virtualization | Enter the percentage of desktops that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| **Power Sources (%)** | |

Table 9-3. Physical Data Center Details (continued)

| Option | Description |
| --- | --- |
| Add Power Source | Click this option to add the share of power sources that you use. The total should add up to 100%.<br>The power sources are coal, oil, natural gas, hydroelectric, solar PV, wind, biomass, and nuclear.<br>The default value for coal is 79.2% and the default value for hydroelectric is 20.8%. The default values are based on the global averages from the Greenhouse Gas (GHG) Emissions Intensity of Electricity Generation methods.<br>The actual global average value is 709g/kWh. |
| **Average Age of Hardware in years (as of today)** | |
| Servers (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |
| Storage (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |
| Network (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |
| Desktop (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |

After you click **Save**, the Sustainability data is calculated. Calculation of the sustainability data may take some time.

## Green Score Details

Green Score is calculated based on five components, which are: Workload Efficiency, Utilization of Physical resources, the extent of Virtualization, Power Source, and Hardware Efficiency. Each of these components has varying levels of impact on the carbon emissions within a data center and hence has a different weightage when contributing to a common Green Score.

- **Workload Efficiency**: Contributes to 22.5% weightage Efficiency = Wastage / (Wastage + Non Wastage).
- **Resource Utilization**: Contributes to 12.5% weightage based on the utilization of hardware from the absolute capacity without considering buffers and HA.
- **Virtualization**: Contributes to 15% weightage.
- **Power Source**: Contributes to 37.5% weightage. Scoring is based on the carbon intensity of the power source used.

- **Hardware Efficiency**: Contributes to 12.5% weightage, based on the fact that newer generation hardware will be more energy efficient.

Table 9-4. Green Score Details and Dependencies

| Name | Description |
| --- | --- |
| Physical Data Centers (map/list) | Displays the configured physical data centers in a map view or in a list. Select the physical data center for which you want to see green score details in a map view or from a list using the options at the top right corner.<br>To reset to the Organizational level, click the **X** button for the selected physical data center. |
| Physical Data Centers by Greenscore | Displays the number of all physical data centers. At the bottom of the panel, you can view the number physical data centers that are configured (green in color) and not configured (gray in color) for green score. For configured physical data centers, you can view the range of green score values applicable to the configured physical data center. |
| Greenscore Meter | Indicates how sustainable a company is. A higher score indicates a greater green score. Green score is calculated based on the different values calculated for workload efficiency, resource utilization, virtualization, power source, and hardware efficiency. |
| Greenscore Trend | Indicates the green score trend of the company. You can view daily, weekly, and monthly trends. |
| **Clean Demand** | |

## Table 9-4. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Workload Efficiency | Measures the wastage in the data center and calculates a score based on various waste sources. The workload efficiency value is calculated based on different metrics, such as powered off VMs, oversized VMs, idle VMs, snapshots, and orphaned disks. The metrics are as follows:<br>Total Consumption CPU = (vCenter) CPU\|vCPUs Allocated on all Powered On Consumers<br>Total Consumption Memory (GB) = (vCenter) Memory\| Host Usage (KB)<br>Total Consumption Disk (TB) = (vCenter) Disk Space\|Total Provisioned Disk Space (GB)<br>Oversized VM\|vCPU = (VM) Summary\|Oversized\|Virtual CPUs and Oversized VM\|Memory (GB) = (VM) Summary\| Oversized\|Memory (KB) - These metrics are used to calculate workload efficiency only for powered on VMs. This means that vCPUs and memory of powered off and oversized VMs are excluded from the calculation.<br>Idle VM\|Disk (TB) = (Cluster) Reclaimable\|Idle VMs\|Disk Space (GB)<br>Idle VM\|vCPU = (Cluster) Reclaimable\|Idle VMs\|CPU (vCPUs)<br>Idle VM\|Memory (GB) = (Cluster) Reclaimable\|Idle VMs\| Memory (KB)<br>Powered Off VM\|Disk (TB) = (Cluster) Reclaimable\| Powered Off VMs\|Disk Space (GB)<br>Snapshot\|Disk (TB) = (Cluster) Reclaimable\|VM Snapshots\| Disk Space (GB)<br>Orphaned VMDK\|Disk (TB) = (Datastore) Reclaimable\| Orphaned Disks\|Disk Space (GB) |
| **Lean Operations** | |

## Table 9-4. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Resource Utilization | Provides a value based on server and storage utilization. The score is based on the utilization of hardware from absolute capacity without considering buffers and HA. Here are the metrics used:<br>Total ESXi Hosts CPU Utilization (GHz) = (vCenter) CPU\|Usage (MHz)<br>Total ESXi Hosts CPU Capacity (GHz) = (vCenter) CPU\|Total Capacity (MHz)<br>Total ESXi Hosts Memory Utilization (TB) = (vCenter) Memory\|Host Usage (KB)<br>Total ESXi Hosts Memory Capacity (TB) = (vCenter) Memory\|Total Capacity (KB) Organization/Physical<br>Total datastores capacity = (vCenter) Disk Space\|Total Capacity (GB)<br>Total datastores utilization = (vCenter) Disk Space\|Utilization (GB)<br>Total RDM = Sum((VM) Virtual Disk \ Configured Size) |
| Virtualization | Virtualization score is calculated based on the virtualization percentage of compute, storage, network, and desktop entered in the corresponding physical data center's configuration. The formula is a weighted average of virtualization percentages with 40%, 30%, 20%, and 10% corresponding weights. |
| **Green Supply** | |

Table 9-4. Green Score Details and Dependencies (continued)

| Name | Description |
| --- | --- |
| Power Source | The score is based on the carbon intensity of the power source used. For example, traditional coal based power sources will have a carbon intensity rate of 888 g/kWh while wind based power source will have a carbon intensity of 26 only (Reference : WNA Report - Comparison of Lifecycle Greenhouse Gas Emissions of Various Electricity Generation Sources.<br>**Note**: A 100% badge score indicates zero carbon emission.<br>Each value is calculated based on the share of the power sources entered in the **Organization Details** page, after which the average value is calculated. The value for each of the power resources is calculated as follows:<br>Green Factor Shares (100%) Adjusted Score (Green Factor * Shares)<br>Coal = 17% (Green Factor)<br>Oil = 31% (Green Factor)<br>Natural Gas = 53% (Green Factor)<br>Solar PV = 92% (Green Factor)<br>Biomass = 96% (Green Factor)<br>Nuclear = 97% (Green Factor)<br>Hydroelectric = 98% (Green Factor)<br>Wind = 98% (Green Factor)<br>**Note** Shares are user inputs in the corresponding physical data center's page. |
| Hardware Efficiency | The hardware efficiency badge is based on the fact that new generation hardware is more energy efficient than the old generation hardware - by running more workloads on the same sized hardware and also the ability of hardware to optimize energy consumption. The efficiency of hardware which includes server, storage, network, and desktop is based on the average age of the hardware entered in the **Organization Details** page. However, you can change and input ages that are specific to the physical data center. For example, if you had entered the age of your servers as 2 years, then the hardware efficiency for servers would be 80%, because the useful life duration is considered as 10 years. |

## Power Consumption Chart

The power consumption chart displays the cumulative power consumption for the last one month. It is calculated using the power usage metrics and is the sum of the last month's daily power consumption. When you click the Power Consumption option, you can also view the **Power Consumption** trend which can be daily, weekly, or monthly.

## Carbon Footprint Chart

The carbon footprint chart displays the carbon emission for the last one month. When you click the Carbon Footprint option, you can also view the **Carbon Footprint** trend which can be daily, weekly, or monthly.

## Environmental Impact

The environmental impact is depicted in terms of carbon emissions from the number of smart phones charged, as per the calculations based on Greenhouse Gas Equivalencies Calculator. For detailed info on this please refer to https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references#smartphones.

For information about Sustainability dashboards, see Sustainability Dashboards.

# Configuring Automation Jobs

10

You can automate jobs to perform certain actions as per a schedule. You can create and manage automation jobs from the Automation Central page. Scheduling of jobs allows you to perform actions without manual supervision. For example, you can automate jobs to run during a maintenance window, which could be outside of working hours.

Read the following topics next:


- Automation Central
- Create Job from Automation Central
- Create Job from Reclaim or Rightsizing


## Automation Central

Automation Central is where you can create jobs to automate optimization actions which reclaim or rightsize VMs. You can also schedule additional actions such as powering on, powering off, and rebooting VMs and instances. Once you set up recurring jobs, you can track and obtain reports on them. You can customize jobs so that they only run based certain on parameters. For example, if you choose to delete a snapshot as an action, you can specify how old the snapshot must be before it is deleted. You can also schedule jobs from the Reclaim and Rightsizing pages, where you configure the job in the context of a recommendation provided by VMware Aria Operations.

### Actions that you can Automate

You can automate the following actions using Automation Central:

**Reclaim**

- Delete old snapshots
- Delete idle VMs
- Poweroff Idle VMs

- Delete powered off VMs

**Rightsize**

- Downsize oversized VMs
- Scale-up undersized VMs

**Other**

Schedule additional actions such as powering on, powering off, and rebooting VMs and instances. The additional actions are based on the public cloud (GCP, AWS, Microsoft Azure) or vCenter adapter.

## Where you find Automation Central

Click **Automation Central** in the left pane.

## How Automation Central Works

In the Automation Central page you see a list of upcoming jobs and a calendar under the **Schedule** tab. The calendar displays all the jobs that are scheduled for the current month. You can move between months to see more scheduled jobs.

### View the Summary of Scheduled Jobs

When you click on a date in the calendar, you see a summary of the job. The summary displays the frequency of the job, the type of job, and if the job is activated or deactivated. You can click **Preview** to see more details about the job, or click **Deactivate All Recurrences** to deactivate the job. To edit the job, click the **Edit** link.

### View a Report of Jobs

View reclamation and rightsizing reports. The reclamation report displays graphical and numerical data on the total cost saving, CPUs reclaimed, memory reclaimed and storage reclaimed for different time periods.

The rightsizing reports displays graphical and numerical data on the CPUs downsized, memory downsized, CPU oversized, and memory upsized for different time periods.

### View Job History

You can also view the history of configured jobs which have run. Click the **History** tab above the calendar to see the job name, and job details in a tabular format. Use the search to search for a Job or VM. Click the drop-down in the search box to perform an advanced search.

The job history page only displays the status of jobs. For detailed information about the task, go to **Administration > Recent Tasks**. Failed recurring jobs are triggered during the next run.

If VM Power Off is not allowed, the action fails and is indicate as such in the logs.

### View Configured Jobs

The **Jobs** tab is where you see a list of configured jobs. For each job, clicking the ellipses icon brings up a menu from where you can edit, delete, clone or deactivate the job. If a job that you created is not visible in the list, search for it by entering the name in the search box. Alternatively, you can check by the job status by clicking the drop-down in the search box for advanced search options.

## Prerequisites to Run Actions

Actions are run through VMware tools. vCenter privileges required to run each actions is documented in topic, Privileges Required for Configuring a vCenter Adapter Instance. Automation Central calls vCenter APIs with the credentials supplied to the vCenter adapter in VMware Aria OperationsVMware Aria Operations.

Automation Central is available for users who have the advanced license and above. You must have the necessary permission to schedule and manage jobs in the **Automation Central** page which you can access from the left menu. An administrator can manage these permissions in **Administration > Access Control > Roles**. Make sure that the **Manage Job Schedules** permission is selected in the Permissions section of the screen. This is available when you expand the **Automation Central** permission.

To schedule a job of a specific type, you must have permission to run the corresponding action defined under **Environment > Actions**. To view permissions assigned to the users and roles, navigate to **Administration > Access Control**.

**Operational Actions** must be enabled in the vCenter cloud adapter instance to run actions from Automation Central.

To use alternate user credentials for actions, see the topic, Configure a vCenter Server Cloud Account in VMware Aria Operations.

## Troubleshooting Automation Jobs

To see the logs for scheduled jobs, go to **Administration > Support Logs > Others** and look for the following logs:

- actionScheduler-.log
- actions-data-.log

Or, go to **Administration > Support Logs > Analytics** and search for action execution in analytics-.log.

# Create Job from Automation Central

Create a job to schedule an action to be performed automatically. You can select the type of action you want to perform, and then select the scope of the action. You can filter the scope based on attributes and metrics. Every action has a configuration option, which enables you to control the execution of the job based on conditions.

Procedure

1 In the Automation Central page, click **Add Job**.

The **Create New Job** page opens. This page displays the following cards:

- Reclaim
- Rightsize
- Billing . This card is available when you when you have VMware Cloud Director configured in VMware Aria Operations. For more information, see the topic, Scheduling Bill Generation Using Automation Central.
- Actions

2 Click on a card based on the type of job you want to create.

3 In the first step of the wizard, specify the following properties to create the action:

| Property | Description |
| --- | --- |
| **Name** | Specify a name for the job. This is displayed in the calendar. |
| **Description** | Provide a description for the job. |
| **Action Configuration** | Select this option for scheduling VM optimization actions. The choices are:<br>**Reclaim**:<br>a Delete old snapshots. When you select this option, you must specify at least one of the following:<br>■ The snapshot age in days.<br>■ The snapshot size in MB using the operators from the drop-down list.<br>■ The snapshot name, using one of the operators from the drop-down list.<br>**Note** The criteria for snapshot deletion is a combination of all the options you choose.<br>b Delete idle VMs. When you select this option, you must specify the idle time duration in days.<br>c Power off idle VMs. When you select this option, you must specify the idle time duration in days.<br>d Delete powered off VMs. When you select this option, you must specify the power off duration in days.<br>**Note** The time definition for powered off VMs, idle VMs, and snapshots is configured **Optimize > Reclaim >** page. Click the **Settings** link on the page to edit or view the reclamation settings. Ensure that you have reviewed the settings in the Reclaimation page before configuring the job here.<br>**Rightsize**:<br>a Downsize oversized VMs. When you select this option, you must select the rightsize resource type, and the downsizing limit amount.<br>b Scale-up undersized VMs. When you select this option, you must select the rightsize resource type.<br>**Note** VMware Aria Operations does not check if the VM hot add/remove setting is enabled. If the VM power off is not allowed, then the action will fail. You will find the error in Recent tasks and in the logs.<br>**Actions**<br>■ **Adapter Type**: Select a public cloud or vCenter adapter from the drop-down list. Options are, vCenter, GCP, AWS, Microsoft Azure.<br>■ **Object Type**: Select the object from the drop-down list on which you want the action to be performed. Options are, Virtual Machine, Host System, Cluster Compute Resource.<br>■ **Action**: Start typing the name of the action, or select an action from the drop-down list. The options are as follows: |

| Adapter Type | Available Action |
| --- | --- |
| vCenter | ■ Power Off VM<br>■ Power On VM<br>■ Reboot Guest OS For VM |

| Property | Description |
|---|---|
| | <table><tr><th>Adapter Type</th><th>Available Action</th></tr><tr><td></td><td>■ Suspend VM</td></tr><tr><td>GCP</td><td>■ Power Off CE Instance<br>■ Power On CE Instance<br>■ Reboot CE Instance</td></tr><tr><td>AWS</td><td>■ Power EC2 Instance<br>■ Power On EC2 Instance<br>■ Reboot EC2 Instance</td></tr><tr><td>Microsoft Azure</td><td>■ Power Off Azure Virtual Machine<br>■ Power On Azure Virtual Machine<br>■ Reboot Azure Virtual Machine</td></tr></table><br>**Note** If the number of resources for a job is ten or less, then the job runs at once. If the number of resources is more than ten, then the jobs run in groups of ten, in parallel. |

4 Click **Next**.

5 In the second step of the wizard, select scope of the object which contains the VMs on which VMware Aria Operations must run the automation jobs.

**Note**

- When you select scope for **Additional Actions**, the objects that you drag and drop must be either of the same object type, or the ancestor object (that will serve as a container) of the object type you have selected in the previous step. For example, if you selected **Host System** in the previous step, you need to select either a Host System object or an object that is higher in hierarchy (Cluster Compute Resources, Datacenter etc).
- VMware Aria Operations ignores object types which do not match the object type that you selected in the previous step. Invalid object types can be object of descendant object types, objects from a different management pack/solution, or from unrelated object hierarchy.

a Drag and drop the object type to the left pane.

b To delete members, click the X icon.

6 After you have dragged & dropped the objects for the action, click **Preview Scope** at the bottom right corner to see on which objects the jobs will potentially be executed. This list may be different after you apply some filter from the **Filter Criteria** page.

**Note** The final objects list on which the actions will be run can be seen by clicking the **Preview Scope** button for the relevant job in the **Schedule** tab.

7 Click **Next**.

8 In the **Filter Criteria** step of the wizard, you can create rules using which VMware Aria Operations filters objects on which the jobs are run. If you do not create a rule, all the VMs within the scope of the selected object type are affected.

   a In the **Set Filter Criteria** section, create rules using a combination of metrics, relationship, property, object name or tag values and logical expressions.

   b You can add more than one criteria by clicking either the ADD ANOTHER CRITERIA SET or ADD buttons. When you click ADD ANOTHER CRITERIA SET, VMware Aria Operations creates an OR logical expression between the criteria sets. When you click ADD button, VMware Aria Operations creates an AND logical expression between the criteria.

9 Click **Next**.

10 In the **Schedule** step of the wizard, configure a schedule for the action to run.

| Property | Description |
| --- | --- |
| **Start Date** | Set a start date which is on, or after the current date, for the job to start. Select a date from the calendar displayed on the page.<br>**Note** There will be a delay of up to five minutes for a job to run. The actual start time for an action also depends on the action itself and the number of target objects involved. |
| **Time zone** | For the start time and date to be calculated for the job, select a time zone. Choices are:<br>■ Browser - based on the current location reported by the browser.<br>■ Host - based on the current location of the host machine.<br>■ GMT based time zone - Based on a location calculated as per UTC+0. |

| Property | Description |
| --- | --- |
| **Start Time** | Select a start time for the job. The drop-down menu provides options in five minute intervals, starting closest to your current browser reported time. |
| **Recurrence** | Set a recurrence for the job. The choices are:<br>■ **One-Time**. No further configuration is possible.<br>■ **Daily**. Set the recurrence in days. Set how frequently it must run by providing a value in **Run Every** field.<br>■ **Weekly**. Select the days of the week when you want the job to run, by clicking the day of the week abbreviation.<br>■ **Monthly**. Select the months when you want the job to run, by clicking the month abbreviation. By default, all the months are selected. For the months that you select, you can configure the job to run:<br>  ■ On specific days of the month, by clicking the number. Or, the last day of the month, without specifying the exact date, by clicking **Last**.<br>  ■ On the specific number of the week (first, second, third, fourth or the last) in the month, combined with the specific day of the week.<br>  You can make multiple selections for each of the options in the drop down menu.<br>For the **Daily**, **Weekly**, **Monthly** options, you can set when the job must end based on a date, or the number of occurrences. |

a In the **Notifications** section, select the **Receive Updates on Job via Email** check box to receive notifications two hours before the job is set to run and after it has been executed. For the email to be sent, you must also select the email outbound plugin from the drop down menu, and enter the email address to which the email must be sent. If you have not created an email outbound plugin, see the topic, Outbound Settings in VMware Aria Operations.

**Note** You can send updates via email to only one email ID.

11 Click **Create** to complete the steps in the wizard and create the job.

# Create Job from Reclaim or Rightsizing

You can create an automation job based on the recommendation provided by VMware Aria Operations VMware Aria Operations in the Reclaim or Rightsizing pages. You cannot create an automation job outside of the context provided by VMware Aria Operations VMware Aria Operations here. Use Automation Central for that.

### Procedure

1 Do one of the following:

a In the menu, click **Home**, and then in the left pane, click **Optimize Capacity > Reclaim**.

b In the menu, click **Home**, and then in the left pane, click **Optimize Performance > Rightsizing**.

2 In the Reclaim or Rightsizing pages, do the following

   a Click the data center that you want to optimize.

   b In the table heading that displays, Select the types of VMs that you want to optimize.

   c Click the name of a listed cluster to show its VM list.

   d Select the checkbox next to the VM that you want to optimize.

   e Click **SCHEDULE ACTION**.

3 In the **Create Schedule Job** dialog box that opens, configure the following parameters:

| Property | Description |
| --- | --- |
| **Job Name** | Provide a name for the job. This information is displayed in the calendar in the Automation Central page. |
| **Job Description** | Provide a description for the job. |
| **Start Date** | From the date picker, select a date when the automation job should start. |
| **Time of Day** | ■ From the time picker, select the start time for the job.<br>■ From the drop down, select the time zone when the time that you selected is valid for. |
| **Receive updates on Job via Email** | Select this checkbox if you have an email server configured, and you want to receive email notification about the status of the job. Notifications are sent two hours before the execution of the job. |
| **Notification Method** | If you selected the previous option, select the email outbound plugin from the drop down menu, and enter the email address to which the email must be sent. |

4 Click **Create**.

Results

The automation job is created, and is available in the Automation Central page. From there, you can preview, edit, or delete the job.

**Note** The scope is not editable since the objects on which the job should be run have been already selected from the Reclaim/Rightsize page.

.

# Log Analysis with VMware Aria Operations for Logs

11

When VMware Aria Operations (SaaS) is integrated with VMware Aria Operations for Logs (SaaS), you can search for these logs, view the logs for an object, and launch the VMware Aria Operations for Logs service from VMware Aria Operations.

## Supported Objects

The following objects are supported for integration between VMware Aria Operations and VMware Aria Operations for Logs:

- vSphere
- VMware vSAN clusters, Disks and Disk Groups
- vCenter
- ESXi hosts, and virtual machines
- NSX Services:
  - Load Balancer Virtual Server
  - Load Balancer Pool
  - Load Balancer Service
  - Logical Switch
  - Logical Router
  - Firewall Section
  - Group
  - Transport Zone
  - Transport Node
  - Edge Cluster
  - Router Service
  - Management Appliances
  - Manager Cluster

- Load Balancers
- Manager Node
- NSX Manager Service

## Integrating the Services

The integration of the VMware Aria Operations and VMware Aria Operations for Logs services happens automatically if you can access both the services and the logs for the supported objects contain the following fields:

- A source field with a valid IP, and not an automated IP such as 169.x.x.x.
- A hostname field.

For information about signing up for and getting started with each service, see the following topics:

- Before You Begin with VMware Aria Operations
- Getting Started Checklist for VMware Aria Operations for Logs (SaaS)

If you are not a VMware Cloud on AWS user, you have to deploy a cloud proxy and then configure vCenter and an ESXi host to send logs for these objects to the cloud proxy for VMware Aria Operations for Logs.

- For more information on deploying a cloud proxy, see the topic, Deploy a Cloud Proxy for VMware Aria Operations for Logs (SaaS).
- For information about configuring vCenter and ESXi, see:
  - Forward vCenter Server Appliance Log Files to Remote Syslog Server
  - Configure Syslog on ESXi Hosts

If you are a VMware Cloud on AWS user, you do not have to deploy a cloud proxy. You also do not have to configure vCenter or an ESXi host, because audit and firewall logs for these objects are already available.

To publish the application or guest OS logs from a virtual machine, you must deploy a VMware Aria Operations for Logs Agent in the virtual machine. The agent helps forward the virtual machine guest OS logs to the cloud proxy for VMware Aria Operations for Logs. For information about VMware Aria Operations for Logs Agents and how to install and configure them, see the topic, Working with VMware Aria Operations for Logs Agents.

## Log Analysis

You can analyse logs in the following ways:

- From the left the menu, select **Troubleshooot > Log Analysis**.
- By selecting an inventory object from the left pane of the **Environment > Object Browser** page, and clicking the **Logs** tab.

You can also create a dashboard which displays the log analysis screen.

You can troubleshoot a problem in your environment by correlating the information in the logs with the metrics, which might help you determine the root cause of the problem. By default, the **Logs** tab displays different event types for the last five minutes. For vSphere objects, the logs are filtered to show the event types for the object that you select and logs from supported objects within that hierarchy.

The supported objects are vSphere, vCenter, ESXi hosts, and virtual machines. The supported VMware vSAN objects are cluster, Disk Groups and Disks. If an object that is not supported has objects that are supported within its hierarchy, the enclosing object shows logs from the objects within its hierarchy. For example, if a cluster has ESXi hosts and virtual machines within its hierarchy, the cluster shows logs from ESXi hosts and virtual machines.

The log analysis screen in the log analysis page, the logs tab of an object, and dashboards looks like the image below:

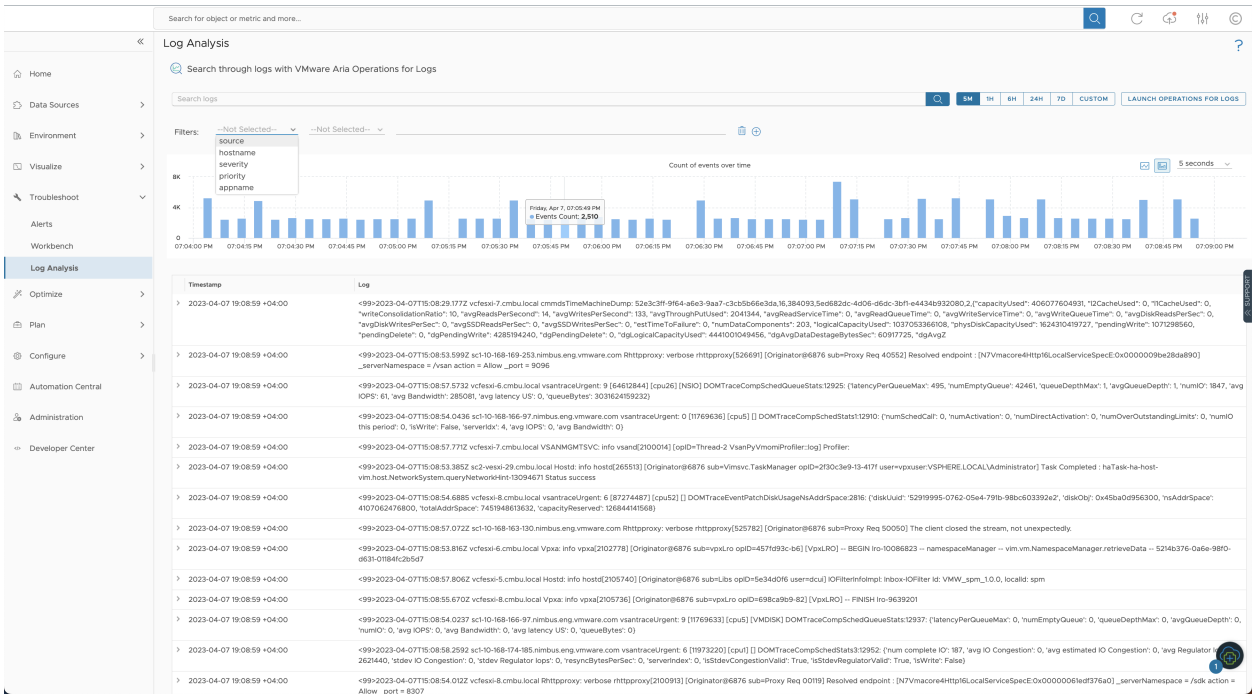


On this screen, you can quickly search for logs or filter down the logs displayed using the following options:

### Text Based Search

- Search bar: On the top of the screen is a search bar to search for logs.
- Time duration of logs: Change this value to display logs from a different time duration. The default is 5 minutes. Use **Custom** for a custom time duration.

### Advanced filtering

Use the advance filtering to narrow down the logs you are looking for. VMware Aria Operations parses the logs and gives you advanced filtering capabilities.

- Filters: Filters. Use source, hostname, severity, priority, and appname filters.
- Condition operators: Use filters with the contains, doesnotContain, startsWith, doesNotStartWith, matchesRegex, exist, and doesNotExist conditions.
- Query box: Input a query in the query box to filter the logs displayed.

The log analysis screen displays the event timeline in a bar chart by default. This tells you how the logs have spiked for the given time duration. You can correlate spikes in logs with metrics to narrow down the root cause of an issue. On the screen, you can change the event timeline frequency between 1 second, 5 seconds and 1 minute. You can switch from the bar graph to a line graph.

The logs are displayed in the bottom half of the page in a table which contains a timestamp and log column. In the table which displays the logs, you can expand a log to see more details.

## Viewing Logs in VMware Aria Operations for Logs

Click the **LAUNCH OPERATIONS FOR LOGS** button to view the VMware Aria Operations logs in VMware Aria Operations for Logs.

From the **Explore Logs** page in VMware Aria Operations for Logs, you can create queries to extract events based on timestamp, text, and fields in log events. VMware Aria Operations for Logs presents charts of the query results.

For information about VMware Aria Operations for Logs, see the VMware Aria Operations for Logs documentation.

# Configuring Alerts and Actions in VMware Aria Operations

12

Alerts are a way to be watchful of anything that is new or an issue that could be potentially dangerous to your environment. Whenever there is a problem in the environment, the alerts are generated. You can also create new alert definitions so that the generated alerts inform you about the problems in the monitored environment. In VMware Aria Operations, alerts and actions play key roles in monitoring the objects.

You may have several questions while working on alerts. Following are some of the key questions that will help you navigate through the alert documentation.

- Accessing Alerts VMware Aria Operations
- Understanding Alerts in VMware Aria Operations
- Viewing Alert Information
- Creating Alert Definitions
- Defining Symptoms for Alerts
- Defining Recommendations for Alerts
- Creating Notification Rules for Alerts
- Adding Outbound Notification Plug-Ins in VMware Aria Operations
- Creating Payload Templates for Outbound Plugins
- Exporting and Importing Outbound Settings
- Deactivating Alerts in VMware Aria Operations
- Grouping Alerts in VMware Aria Operations
- Intelligent Alerts

Actions allow you to make changes to the objects in your environment. When you grant a user access to actions in VMware Aria Operations, that user can take the granted action on any object that VMware Aria Operations manages. For details, see Actions in VMware Aria Operations.

Read the following topics next:

- Understanding Alerts in VMware Aria Operations
- Alert Definitions in VMware Aria Operations

- Symptom Definitions in VMware Aria Operations
- Recommendations in VMware Aria Operations
- Notifications in VMware Aria Operations
- Outbound Settings in VMware Aria Operations
- Payload Templates in VMware Aria Operations
- Managing Alert Groups in VMware Aria Operations
- Alert Definition Best Practices
- Create a Simple Alert Definition
- Create a New Alert Definition
- Create an Alert Definition for Department Objects
- Actions in VMware Aria Operations


# Understanding Alerts in VMware Aria Operations

This topic provides information on the different types of alerts in VMware Aria Operations, how to access them, and how to view more information about these alerts.

## Types of Alerts

Alerts in VMware Aria Operations are of three types. The alert type determines the severity of the problem.

### Health Alerts

The health alert list is all the generated alerts that are configured to affect the health of your environment and require immediate attention. You use the health alert list to evaluate, prioritize, and immediately begin resolving the problems.

### Risk Alerts

The risk alerts list is all the generated alerts that are configured to indicate risk in your environment. Address risk alerts in the near future, before the triggering symptoms that generated the alert negatively affect the health of your environment.

### Efficiency Alerts

The efficiency alerts list is all the generated alerts that are configured to indicate problems with the efficient use of your monitored objects in your environment. Address efficiency alerts to reclaim wasted space or to improve the performance of objects in your environment.

## Accessing Alerts VMware Aria Operations

The **All Alerts** or the **Administrative Alerts** page provides the list of all the alerts generated in VMware Aria Operations. Use the alert list to determine the state of your environment and to begin resolving problems.

### Where You Find the All Alerts Page

From the left menu, click **Troubleshoot > Alerts** .

### Where You Find the Administrative Alerts Page

As an admin, you can view the administrative alerts by clicking the warning icon next to the Alerts menu or from the left menu, click **Troubleshoot > Alerts** and then click the **Administrative Alerts** tab. You can view the **Administrative Alerts** page, only if you are a global admin user or if you have administrative privileges assigned to you.

### How the All Alerts and Administrative Alerts Pages Work

By default, only active alerts are initially listed, and the alerts are grouped by Time. Review and manage the alerts in the list using the toolbar options. Select multiple rows in the list using Shift+click, Control+click.

To see the alert details, click the alert name. The alert details appear on the right, including the symptoms triggered by the alert. The system offers recommendations for addressing the alert and link to run the recommendation. A Run Action button may appear in the details. Hover over the button to learn what recommendation is performed if you click the button. Alternatively, you can view the **Run** button and the **Suggested Fix** in the Alerts data grid. You can filter by alerts that have the Run option activated and perform the recommended task to address the alert from the Alerts data grid. Click the small box on the lower left of the alert list to include the **Suggested Fix** and **Run** columns in the data grid.

Click the name of the object on which the alert was generated to see the object details, and access additional information relating to metrics and events.

If you migrated alerts from a previous version of VMware Aria Operations, the alerts are listed with a cancelled status and alert details are not available.

### All Alerts and Administrative Alerts Options

The alert options include toolbar and data grid options. Use the toolbar options to sort the alert list and to cancel, suspend, or manage ownership. Use the data grid to view the alerts and alert details.

Select an alert from the list to activate the Actions menu:

### Table 12-1. Actions Menu

| Option | Description |
| --- | --- |
| Cancel Alert | Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.<br>Cancel alerts when you do not need to address them. Canceling an alert does not cancel the underlying condition that generated it. Canceling alerts is effective if the alert is triggered by fault and event symptoms, because these symptoms are triggered again only if subsequent faults or events occur on the monitored objects. If the alert was generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |
| Delete Canceled Alerts | Delete cancelled (inactive) alerts by doing a group selection or by individually selecting alerts. The option is deactivated for active alerts. |
| Suspend | Suspend an alert for a specified number of minutes.<br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.<br>The user who suspends the alert becomes the assigned owner. |
| Assign to | Assign the alert to a user. You can search for a specific username and click **Save** to assign the alert to the selected user. |
| Take Ownership | As the current user, you make yourself the owner of the alert.<br>You can only take ownership of an alert, you cannot assign ownership. |
| Release Ownership | Alert is released from all ownership. |
| Go to Alert Definition | Switches to the Alert Definitions page, with the definition for the previously selected alert displayed. |

### Table 12-1. Actions Menu (continued)

| Option | Description |
| --- | --- |
| Deactivate... | Provides two options to deactivate the alert:<br>**Note** To activate the Deactivate option, select **Definition** from the **Group By** drop-down list, and click on the name of the Alert Definition Group.<br>■ Deactivate the alert in all policies: This deactivates the alert for all objects for all the policies.<br>■ Deactivate alert in selected policies: This deactivates the alert for objects having the selected policy. |
| Open an external application | Actions you can run on the selected object.<br>For example, Open Virtual Machine in vSphere Client. |

### Table 12-2. Group By Options

| Option | Description |
| --- | --- |
| None | Alerts are not sorted into specific groupings. |
| Time | Group alerts by time triggered. This is the default option. You can also group by 1 hour, 4 hours, Today and Yesterday, days of current week, Last week and Older. |
| Criticality | Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the "All Alerts Data Grid Options" table, below. |
| Definition | Group alerts by definition, that is, group like alerts together. |
| Object Type | Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together. |
| Scope | Group alerts by scope. You can search for alerts within the selected scope. |

### Table 12-3. Quick Filters (Alert)

| Quick Filters | Descriptions |
| --- | --- |
| Filtering options | Limit the list of alerts to those matching the filters you choose.<br>For example, you might have chosen the Time option in the Group By menu. Now you can choose Status -> Active in the Quick Filters menu, and the All Alerts/Administrative Alerts page displays only the active alerts, ordered by the time they were triggered. |
| Options (see also the Group By and All Alerts Data Grid tables for more filter definitions) | |
| Alert id | ID given for an alert. |
| Alert | Name of the alert definition that generated the alert. |

Table 12-3. Quick Filters (Alert) (continued)

| Quick Filters | Descriptions |
| --- | --- |
| Owner | Name of operator who owns the alert. |
| Impact | Alert badge affected by the alert. The affected badge, health, risk, or efficiency, indicates the level of urgency for the identified problem. |
| Alert Subtype | Additional information about the type of alert that is triggered on a selected object. This helps you categorize the alerts in a detailed level other than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration. |
| Status | Current state of the alert.<br>Possible values include Active or Canceled. |
| Criticality | The level of importance of the alert in your environment.<br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.<br>The possible values include:<br>■ Critical<br>■ Immediate<br>■ Warning<br>■ Information |
| Triggered On | Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name.<br>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object. |
| Control State | State of user interaction with the alert. Possible values include:<br>■ Open. The alert is available for action and has not been assigned to a user.<br>■ Assigned. The alert is assigned to the user who is logged in when that user clicks **Take Ownership**.<br>■ Suspended. The alert was suspended for a specified amount of time. The alert is temporarily excluded from affecting the health, risk, and efficiency of the object. This state is useful when a system administrator is working on a problem and does not want the alert to affect the health status of the object. |
| Object Type | Type of object on which the alert was generated. |
| Created On | Date and time when the alert was generated. |

### Table 12-3. Quick Filters (Alert) (continued)

| Quick Filters | Descriptions |
| --- | --- |
| Updated On | Date and time when the alert was last modified.<br>An alert is updated whenever one of the following changes occurs:<br>■ Another symptom in the alert definition is triggered.<br>■ Triggering symptom that contributed to the alert is canceled. |
| Canceled On | Date and time when the alert canceled for one of the following reasons:<br>■ Symptoms that triggered the alert are no longer active. Alert is canceled by the system.<br>■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions are deactivated in the policy that is applied to the object.<br>■ Symptoms that triggered the alert are canceled because the corresponding symptom definitions were deleted.<br>■ Alert definition for this alert is deactivated in the policy that is applied to the object.<br>■ Alert definition is deleted.<br>■ User canceled the alert. |
| Action | Choose **Yes** to filter based on alerts that have the **Run** option activated. Choose **No** to filter based on alerts that have the **Run** option deactivated. |

The Alerts data grid provides the list of generated alerts used to resolve problems in your environment. An arrow in each column heading orders the list in ascending or descending order.

### Table 12-4. All Alerts and Administrative Alerts Data Grid

| Option | Description |
| --- | --- |
| Criticality | Criticality is the level of importance of the alert in your environment.<br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.<br>The possible values include:<br>■ Critical<br>■ Immediate<br>■ Warning<br>■ Information |
| Alert | Name of the alert definition that generated the alert.<br>Click the alert name to display the alert details to the right. |

Table 12-4. All Alerts and Administrative Alerts Data Grid (continued)

| Option | Description |
| --- | --- |
| Triggered On | Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name.<br>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object. |
| Created On | Date and time when the alert was generated. |
| Status | Current state of the alert.<br>Possible values include Active or Canceled. |
| Alert Type | Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/Hypervisor, Hardware, Storage, Network, Administrative, and Findings. |
| Alert Subtype | Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration. |
| Importance | Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm. |
| Suggested Fix | Displays the recommendation to address the alert. |
| Action | Click this button to perform the recommendation to address the alert. |

# Viewing Alert Information

When you click an alert from the all alerts list, the alert information appears on the right. View the alert information to see the symptoms which triggered the alert, recommendations to fix the underlying issue, and troubleshoot the cause of the alert.

## Different ways to view Alert information

- From the left menu, click **Troubleshoot > Alerts**, and then click an alert from the alert list.
- From the left menu, click **Environment > Object Browser**, then select a group, custom data center, application, or inventory object. Click the object and then the **Alerts** tab.
- In the menu, select Search and locate the object of interest. Click the object and then the **Alerts** tab.

The alert description is hidden when you open the alert information. Click **View Description** to see the description of the alert. View the time stamp of when the alert started, and when it was updated, below the alert title.

### Alert Details Tab

| Section | Description |
| --- | --- |
| Recommendations | View recommendations for the alert. Click **<** or **>** to cycle through the recommendations. To resolve the alert, click the **Run Action** button if it appears. |
| Other Recommendations | Collapse the section to view additional recommendations. See the links in the **Need More Information?** section to view additional metrics, events, or other details that appear as a link. |
| Alert Basis | |
| Active Only | This option is activated by default. When activated, all active symptoms/conditions that were met for the alert are displayed. When deactivated, all the symptoms/conditions of an alert are displayed. |
| Symptoms | View the symptoms that triggered the alert. Collapse each symptom to view additional information. |
| Conditions | View the conditions that triggered the alert. Collapse each condition to view additional information. |
| Troubleshoot with Logs | Opens the **Logs** tab in the Troubleshooting Workbench, with the current object in context. Displays logs from 15 minutes before when the alert was triggered. |
| Notes | Enter your notes about the alert and click **Submit** to save. |
| Close | Click the X icon to close the alert details tab. |

### Related Alerts Tab

The **Related Scope** displayed on the right, shows the objects that are one level above and one level below the object on which the alert was triggered. This topology is fixed. You cannot change the scope in the **Related Alerts** tab.

On the right, you can see the following:

- If the same alert was triggered on the object in the past 30 days. This helps you understand if this is a recurring problem or something new.
- If the same alert was triggered on other peers in the same environment, in the past 30 days. This helps you do a quick peer analysis to understand if others are impacted with the same problem.

- All the alerts triggered in the current topology. This helps you investigate if there are other alerts upstream or downstream in the environment which are impacting the health of the object.

**Potential Evidence Tab**

See the **Potential Evidence** tab for potential evidences around the problem, and to arrive at the root cause. This tab displays events, property changes, and anomalous metrics potentially relevant to the alert. The time range and the scope are fixed. To modify the scope or the time range and investigate further, click **Launch Workbench**. This runs the troubleshooting workbench.

The time range that is displayed in the potential evidence tab is two hours and thirty minutes before the alert was triggered. VMware Aria Operations looks for potential evidences in this time range.

## Intelligent Alerts

Every enterprise could have five or more monitoring tools that monitor various aspects of their data center operations around the clock. This could cause an alert flooding situation where multiple alerts are generated by a single monitoring tool or multiple tools for the same problem. As a result, IT administrators must sift through thousands of alerts to filter out the noise and focus on key issues, thus, increasing the sheer volume of the alerts and posing an alert storm or alert noise resulting in teams being unable to identify the most critical alerts. Alert flooding happens because monitoring tools lack the intelligence to understand that all alerts depict the same problem.

Machine Learning (ML) helps automate the management of complex systems that contain thousands of objects like VMs, Hosts, and Datastores, through monitoring millions of metrics, huge volumes of logs, and application traces, to capture a high-resolution image of the entire stack.

VMware Aria Operations through intelligent alert clustering helps eliminate business downtime that happens because of a lack of faster troubleshooting abilities and solving critical problems over multiple objects.

### Where You Find the Intelligent Alerts Tab

From the left menu, click **Troubleshoot > Alerts**, and then, click the **Intelligent Alerts** tab.

### How Intelligent Alert Clustering Works

Intelligent alerts, also known as alert clusters in VMware Aria Operations, groups related alerts together based on their creation time and their topology distance. This approach provides a more organized and efficient method for troubleshooting, compared to dealing with individual alerts arising from the same underlying issue. Alerts clustering is done based on DBScan algorithm. DBScan (Density-Based Spatial Clustering of Applications with Noise) is an unsupervised clustering Machine Learning algorithm that attempts to group data points closely packed into artificial clusters. In the context of VMware Aria Operations, DBScan has been

tailored into a streaming algorithm with specific parameters configured such as minimum points is set to five, time difference is set to fifteen minutes, and topology distance is set to one, for considering only direct children and parents. Two main views, Intelligent alert lifetime and Objects topology, are provided for alert cluster troubleshooting.

The **Intelligent Alert** tab displays the list of alert clusters in the left pane. Click any alert cluster to view the details in the right pane.

| Option | Description |
| --- | --- |
| Filters | You can filter the alert clusters by their status. Select **Active** or **Inactive** from the **Status** drop-down list and click **Apply**. |
| Alert Cluster | The alert cluster card displays the following:<br>■ **Status**: Displays if the status of the alert cluster is active or inactive.<br>■ **Object**: Displays the name of the object to which the cluster is assigned.<br>■ **Alert Chart**: Displays the number of alerts and objects along with the criticality of the alert. Hover over the graph to view the details.<br>■ **Start time**: Displays the time when the first cluster that satisfies the clustering condition is identified.<br>■ **End time**: Displays the time when the cluster no longer qualifies to be an alert cluster.<br>Click the alert cluster to view the details in the right pane. |
| Object | Name of the root object. |
| Start time/End time | The start time of the alert cluster is the time when the first cluster that satisfies the clustering condition is identified. The end time of the alert cluster is the time when the cluster no longer qualifies to be an alert cluster. |
| Alerts/Objects | Select **Alerts** to view the graphical representation of alerts in a specific period.<br>Select **Objects** to view the object-relationship chart for an alert cluster. Hover over the object and click **Details** to open the Summary page of the object. |
| How it Started | Click **How it Started** to view the lifetime of an alert cluster. Each bubble displays alerts and objects, hover over the bubble to view more details. |
| Troubleshoot | Click this to launch the Troubleshooting workbench for further troubleshooting. |

| Option | Description |
| --- | --- |
| Graph Chart | The graph chart displays the number of alerts by time, for the selected alert cluster.<br>Click the chart legend to filter alerts by:<br>■ Criticality<br>  ■ Critical<br>  ■ Immediate<br>  ■ Warning<br>  ■ Info<br>■ Objects<br>Click the Calendar icon to view past alerts by selecting the **Range** or by selecting a date in the **From** and **To** fields. |
| Group By | You can group alerts by:<br>■ Definition<br>■ Scope<br>■ Time<br>■ Criticality<br>■ Object Type |
| Filters | You can filter alerts by:<br>■ Alert ID<br>■ Alert name<br>■ Owner<br>■ Impact<br>■ Alert Type<br>■ Alert Subtype<br>■ Status |

# Alert Definitions in VMware Aria Operations

Alert definitions are a combination of symptoms and recommendations that you combine to identify problem areas in your environment and generate alerts on which you can act for those areas. You can then respond to the alerts with effective solutions that are provided in the recommendations.

# Where You Find Alert Definitions

To manage your alert definitions, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Alert Definitions**.

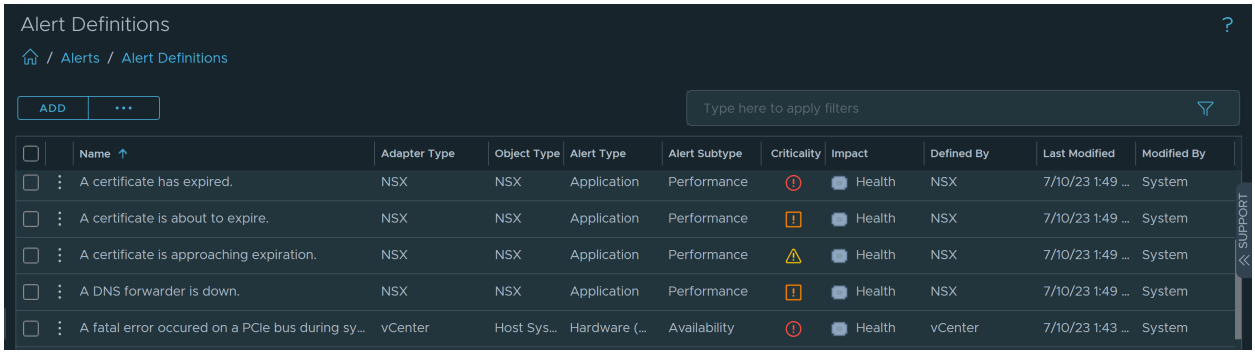


| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your alert definitions.<br>■ Add. Add an alert definition.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected definition.<br>■ Delete. Remove the selected definition.<br>■ Clone. Create a copy of the selected definition so that you can customize it for your needs.<br>■ Export. Downloads the alert definition.<br>■ Import. Allows you to import alert definitions. To import:<br>  ■ Click the **Import** option from the horizontal ellipsis.<br>  ■ Click **Browse** and select the file to import.<br>  ■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>  ■ Click **Import** to import the alert definition, and click **Done**. |
| Filters | Limits the list of alerts to those matching the filter you create.<br>You can also sort on the columns in the data grid. |
| Name | Name of the alert definition, which is also the name of the alert that appears when the symptoms are triggered. |
| Adapter Type | Adapter that manages the selected base object type. |
| Object Type | Base object type against which the alert is defined. |
| Alert Type | Metadata that is used to classify the alert when it is generated.<br>You define the value on the Alert Impact page of the workspace. |

| Option | Description |
| --- | --- |
| Alert Subtype | Subcategory of the alert type and is the metadata that is used to classify the alert when it is generated.<br>You define the value on the Alert Impact page of the workspace. |
| Criticality | Severity of the alert when it is generated. The criticality includes the following possible values:<br>■ Symptom. Alert is configured to display symptom based criticality.<br>■ Critical<br>■ Immediate<br>■ Warning<br>■ Info |
| Impact | Alert is configured to affect the Health, Risk, or Efficiency badge. |
| Defined by | Indicates who added the alert definition. The alert can be added by an adapter, a user, or the VMware Aria Operations system. |
| Last Modified | Displays the date on which the alert was last modified. |

Predefined alerts are provided in VMware Aria Operations as part of your configured adapters. Use Alert Definitions to manage your VMware Aria Operations alert library, and to add or modify the definitions.

### Modifying Alert Definitions

If you modify the alert impact type of an alert definition, any alerts that are already generated will have the previous impact level. Any new alerts will be at the new impact level. If you want to reset all the generated alerts to the new level, cancel the old alerts. If they are generated after cancellation, they will have the new impact level.

### Symptoms in Alert Definitions

Symptom definitions evaluate conditions in your environment that, if the conditions become true, trigger a symptom and can result in a generated alert. You can add symptom definitions that are based on metrics or super metrics, properties, message events, fault events, or metric events. You can create a symptom definition as you create an alert definition or as an individual item in the appropriate symptom definition list.

When you add a symptom definition to an alert definition, it becomes a part of a symptom set. A symptom set is the combination of the defined symptom with the argument that determines when the symptom condition becomes true.

An alert definition comprises one or more symptom sets. If an alert definition requires all of the symptom sets to be triggered before generating an alert, and only one symptom set is triggered, an alert is not generated. If the alert definition requires only one of several symptom sets to be triggered, then the alert is generated even though the other symptom sets were not triggered.

### Recommendations in Alert Definitions

Recommendations are the remediation options that you provide to your users to resolve the problems that the generated alert indicates.

When you add an alert definition that indicates a problem with objects in your monitored environment, add a relevant recommendation. Recommendations can be instructions to your users, links to other information or instruction sources, or VMware Aria Operations actions that run on the target systems.

## Creating Alert Definitions

The alert definition process includes adding symptoms that trigger an alert and recommendations that help you resolve the alert. The alert definitions you create with this process are saved to your VMware Aria Operations Alert Definition Overview list and actively evaluated in your environment based on your configured policies.

You can create or reuse existing symptoms and recommendations while defining an alert definition. If you create symptoms and recommendations, you add them to the definition, and they are added to the symptom and recommendations content libraries for future use. You also activate policies and select notifications for the alerts.

1 To create or edit your alert definitions, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Alert Definitions**.

2 Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition.

3 In the **Alert** tab, enter details of the alert.

| Option | Description |
| --- | --- |
| Name | Name of the alert as it appears when the alert is generated. |
| Description | Description of the alert as it appears when the alert is generated. Provide a useful description for your users. |
| Base Object Type | The object type against which the alert definition is evaluated and the alert is generated.<br>The drop-down menu includes all of the object types in your environment. You can define an alert definition based on one object type. |

| Option | Description |
| --- | --- |
| Impact | Under Advanced Settings, select the badge that is affected if the alert is generated.<br>You can select a badge based on the urgency of the alert.<br>■ Health. Alert requires immediate attention.<br>■ Risk. Alert should be addressed soon after it is triggered, either in days or weeks.<br>■ Efficiency. Alert should be addressed in the long term to optimize your environment. |
| Criticality | Severity of the alert that is communicated as part of the alert notification.<br>Select one of the following values.<br>■ Info. Informational purposes only. Does not affect badge color.<br>■ Warning. Lowest level. Displays yellow.<br>■ Immediate. Medium level. Displays orange.<br>■ Critical. Highest level. Displays red.<br>■ Symptom Based. In addition to alert criticality, each symptom includes a defined criticality. Criticality of the alert is determined by the most critical of all of the triggered symptoms. The color is dynamically determined accordingly. It you negate symptoms, the negative symptoms to not contribute to the criticality of a symptom-based alert. |
| Alert Type and Subtype | Select the type and subtype of alert.<br>This value is metadata that is used to classify the alert when it is generated, and the information is carried to the alert, including the alert notification.<br>You can use the type and subtype information to route the alert to the appropriate personnel and department in your organization. |

| Option | Description |
| --- | --- |
| Wait Cycle | The symptoms included in the alert definition remain triggered for this number of collection cycles before the alert is generated.<br>The value must be 1 or greater.<br>This setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition is added to the wait cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that after all of the symptoms are triggered at the desired symptom sensitivity level, the alert is immediately triggered. |
| Cancel Cycle | The symptoms are cancelled for this number of collection cycles after which the alert is cancelled.<br>The value must be 1 or greater.<br>This setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition is added to the cancel cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of the alert definition to 1. This configuration ensures that after all of the symptom conditions disappear after the desired symptom cancel cycle, the alert is immediately canceled. |

4 Click **Next** to add symptom definitions.

5 In the **Symptoms/Conditions**, drag the selected symptom/condition in to the left pane. Use the workspace on the left to specify whether all or any of the symptoms/conditions or symptom/condition sets must be true to generate an alert. As you add one or more symptoms, you create a symptom expression. If this expression is evaluated as true, then the alert is generated. You can similarly define one or more conditions for your alert, and when the conditions are met, the alert is generated. You can view the alert in the All Alerts page.

Table 12-5. Add Symptoms/Conditions Selection Options

| Option | Description |
| --- | --- |
| Defined On | Object that the symptom evaluates.<br>As you create alert definitions, you can select or define symptoms for the base object type and for related object types, based on the object relationship hierarchy. The following relationships are object types as they relate to the alert definition base object type.<br>■ Self. A base object type for the alert definition. For example, host system.<br>■ Descendant. An object type that is at any level below the base object type, either a direct or indirect child object. For example, a virtual machine is a descendant of a host system.<br>■ Ancestor. An object type that is one or more levels higher than the base object type, either a direct or indirect parent. For example, a data center and a vCenter Server are ancestors of a host system.<br>■ Parent. An object type that is in an immediately higher level in the hierarchy from the base object type. For example, a data center is a parent of a host system.<br>■ Child. An object type that is one level below the base object type. For example, a virtual machine is a child of a host system. |
| Symptoms tab | |

Table 12-5. Add Symptoms/Conditions Selection Options (continued)

| Option | Description |
| --- | --- |
| Select Symptom | Select the type of symptom definition that you are adding for the current Defined On object type.<br>■ Metric / Property. Add symptoms that use metric and property symptoms. These metrics are based on the operational or performance values, and configuration properties that VMware Aria Operations collects from target objects in your environment.<br>■ Message Event. Add symptoms that use message event symptoms. These symptoms are based on events received as messages from a component of VMware Aria Operations or from an external monitored system through the system's REST API.<br>■ Fault Event. Add symptoms that use fault symptoms. These symptoms are based on events that monitored systems publish. VMware Aria Operations correlates a subset of these events and delivers them as faults. Faults are intended to signify events in the monitored systems that affect the availability of objects in your environment.<br>■ Metric Event. Add symptoms that use metric event symptoms. These symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not VMware Aria Operations. These symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that VMware Aria Operations is actively monitoring.<br>■ Smart Early Warning. Add a symptom that uses a defined condition that is triggered when the number of anomalies on an object is over the trending threshold. This symptom represents the overall anomalous behavior of the object. Anomalies are based on VMware Aria Operations analysis of the number of applicable metrics that violate the dynamic threshold that determines the normal operating behavior of the object. This symptom is not configurable. You either use it or you do not use it. |
| Filter by Object Type | Available only when you select a Defined On value other than Self.<br>Limits the symptoms to those that are configured for the selected object type based on the selected Defined On relationship. |

### Table 12-5. Add Symptoms/Conditions Selection Options (continued)

| Option | Description |
| --- | --- |
| Create New Symptom | If symptoms that you need for your alert do not exist, you can create them.<br>Opens the symptoms definition dialog box.<br>Not available for Smart Early Warning symptoms, which are predefined in the system. |
| All Filters | Filter the list of symptom definitions. This selection is available when Defined On is set to **Self**, or when it is set to another relationship and you select an object from the Filter by Object Type drop-down menu.<br>■ Symptom. Type text to search on the name of the symptom definitions. For example, to display all symptom definitions that have efficiency in their name, type **`Efficiency`**.<br>■ Defined By. Type text to search for the name of the adapter that defines the symptom definitions. For example, to display all symptom definitions provided by the vCenter Adapter, type **`vCenter`**. To display only user-defined symptom definitions, type the search term **`User`**.<br>To clear a filter, click the double arrow icon that appears next to the filter name. |
| Quick filter (Name) | Search the list based on the symptom name. |
| Symptoms list | List of existing symptoms for the selected object type. To configure a symptom, drag it into the left workspace.<br>To combine symptoms that are based on multiple levels in the hierarchy, select the new Defined On level and Filter by Object Type before you select and drag the new symptom to the workspace. |
| Conditions tab | |
| Select Specific Object | Select a specific object based on its object type, adapter type, policy, collection state, and status. |
| Filter | Search the metrics based on object type. |
| Conditions list | List of metrics for the selected object type. To configure a condition, drag it into the left workspace. |

Use the workspace to configure the interaction of the symptoms, symptom sets, and conditions.

### Table 12-6. Symptom Sets in the Alert Definition Workspace

| Option | Description |
| --- | --- |
| Trigger alert when {operator} of the symptom sets are true | Select the operator for all of the added symptom/condition sets. Available only when you add more than one symptom/condition set.<br>■ All. All of the symptom/condition sets must be true before the alert is generated. Operates as a Boolean AND.<br>■ Any. One or more of the symptom/condition sets must be true before the alert is generated. Operates as a Boolean OR. |
| Symptoms | The symptom/condition sets comprise an expression that is evaluated to determine if an alert should be triggered.<br>To add one or more symptoms from the symptom list to an existing symptom set, drag the symptom from the list to the symptom set. To create a new symptom set for the alert definition, drag a symptom to the landing area outlined with a dotted line. |
| Symptom sets | Add one or more symptoms to the workspace, define the points at which the symptom sets are true, and specify whether all or any of the symptoms in the symptom set must be true to generate the alert.<br>A symptom set can include one or more symptoms/conditions, and an alert definition can include one or more symptom/condition sets.<br>If you create a symptom set where the Defined On object is Self, you can set the operator for multiple symptoms in the symptom set.<br>If you create a symptom set where the Defined On object is a relationship other than Self, you can set the operator and modify the triggering threshold. To configure the symptom set criteria, you set the options.<br>■ Value operator. Specifies how the value you provide in the value text box is compared to a number of related objects to evaluate the symptom/condition set as true.<br>■ Value text box. Number of objects of the specified relationship, based on the value type, that are required to evaluate the symptom/condition set as true.<br>■ Value type. Possible types include the following items:<br>  ■ Count. Exact number of related objects meet the symptom/condition set criteria.<br>  ■ Percent. Percentage of total related objects meet the symptom/condition set criteria.<br>  ■ Any. One or more of the related objects meet the symptom/condition set criteria.<br>  ■ All. All of the related objects meet the symptom/condition set criteria.<br>■ Symptom set operator. Operator applied between symptoms/conditions in the symptom set.<br>  ■ All. All of the symptoms/conditions must be true before the alert is generated. Operates as a Boolean AND.<br>  ■ Any. One or more of the symptoms/condition must be true before the alert is generated. Operates as a Boolean OR. |

### Table 12-6. Symptom Sets in the Alert Definition Workspace (continued)

| Option | Description |
| --- | --- |
| | When you include a symptom in a symptom set, the condition must become true to trigger the symptom set. However, you might want to configure a symptom set where the absence of a symptom condition triggers a symptom. To use the absence of the symptom condition, click the vertical ellipsis on the left of the symptom name and select **Invert Symptom**.<br>Although you can configure symptom criticality, if you invert a symptom, it does not have an associated criticality that affects the criticality of generated alerts. |

### Table 12-7. Conditions in the Alert Definition Workspace

| Option | Description |
| --- | --- |
| Alert is triggered when {operator} of the sets are true | Select the operator for all of the added condition sets. Available only when you add more than one condition set.<br>■ All. All of the condition sets must be true before the alert is generated. Operates as a Boolean AND.<br>■ Any. One or more of the condition sets must be true before the alert is generated. Operates as a Boolean OR. |
| Conditions | The condition sets comprise an expression that is evaluated to determine if an alert should be triggered.<br>■ Condition. Determines how the value you specify in the value text box is compared to the current value of the metric or property when the condition is evaluated.<br>■ Value. Value that specifies the threshold.<br>■ Criticality level. Severity of the symptom/condition when it is triggered.<br>■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom/condition is triggered. The default value is 1, which means that the symptom/condition is triggered in the same collection cycle when the condition became true.<br>**Note** You cannot edit the wait cycle while defining conditions for Properties and Population.<br>■ Cancel Cycle. The symptom/condition is canceled after the trigger condition is false for this number of collection cycles after which the symptom/condition is cancelled. The default value is 1, which means that the symptom/condition is canceled in the same cycle when the condition becomes false.<br>**Note** You cannot edit the cancel cycle while defining conditions for Properties and Population.<br>To add one or more conditions from the condition list to an existing symptom/condition set, drag the condition from the list to the symptom/condition set. |

6 Click **Next** to add recommendations.

7 In the **Recommendations** tab, drag the selected recommendation in to the left pane. Use the workspace on the left to to change the priority order.

Table 12-8. Add Recommendations Options in the Alert Definition Workspace

| | |
|---|---|
| Create New Recommendation | If recommendations that you need to resolve the symptoms in the problem do not exist, you can create them. |
| All Filters | Filter the list of recommendations.<br>■ Description. Type text to search on the name of the recommendation. For example, to display all recommendations that have memory in their name, type **`Memory`**.<br>■ Defined By. Type text to search for the name of the adapter that defines the recommendation. For example, to display all recommendations provided by the vCenter Adapter, type **`vCenter`**.<br>To clear a filter, click the double arrow icon that appears next to the filter name. |
| Quick filter (Name) | Limits the list based on the text you enter. |
| List of available recommendations. | List of existing recommendations that you can drag to the workspace.<br>Recommendations are instructions and, where possible, actions that assist you with resolving alerts when they are triggered. |
| Recommendation workspace | Add one or more recommendations to the workspace.<br>If you add more than one recommendation, you can drag the recommendations to change the priority order. |

8 Click **Next** to activate policies.

9 In the **Policies** tab, you can view the policy tree in the left pane and you can either select the default policy or any other policy from the tree.

You can automate the recommended action that has the highest priority by changing the **Status** to **Activated** in the right pane. Whenever the alert is executed on an object within the policy, the recommended action will be executed on the object.

**Note** To deactivate alerts associated with a specific policy, deselect the policy in the left pane, and click **Update**.

You can also customize thresholds for a policy by clicking the policy and editing the trigger value in the right pane. Editing the threshold of conditions will affect its alert definition in the selected policy.

**Note** If you create an alert without enabling any policies, then the alert remains inactive.

10 Click **Create** to create the alert. The new alert appears in the list of alert definitions.

# Symptom Definitions in VMware Aria Operations

Symptoms are conditions that indicate problems in your environment. You can define symptoms in VMware Aria Operations and add them to alert definitions so that you know when a problem occurs with your monitored objects.

As data is collected from your monitored objects, the data is compared to the defined symptom condition. If the condition is true, then the symptom is triggered.

To define symptoms, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**.

You can define symptoms based on:

- Metric/Property Symptoms
- Message Event Symptoms
- Fault Symptoms
- Metric Event Symptoms

The symptoms defined in your environment are managed in the Symptom Definitions. When the symptoms that are added to an alert definition are triggered, they contribute to a generated alert.

### Define Symptoms to Cover All Possible Severities and Conditions

Use a series of symptoms to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of Warning while `Volume reached capacity limit` might have a severity level of Critical. The first symptom is not an immediate threat. The second symptom is an immediate threat.

## Defining Symptoms for Alerts

The symptoms defined in your environment are managed in the **Symptom Definitions** page.

You can create a symptom definition while creating an alert definition or as an individual item from the **Symptom Definitions** page.

You can define symptoms based on:

- Defining Metric/Property Symptoms
- Defining Message Events
- Defining Fault Symptoms
- Defining Metric Event Symptoms

## Metric/Property Symptoms

Metric/Property symptoms are based on the operational/performance values or the configuration properties that VMware Aria Operations collects from target objects in your environment.

### Metric Symptom Definitions

You can configure the symptoms to evaluate static thresholds or dynamic thresholds. You define symptoms based on metrics so that you can create alert definitions that let you know when the performance of an object in your environment is adversely affected.

#### Static Thresholds

Metric symptoms that are based on a static threshold compare the currently collected metric value against the fixed value you configure in the symptom definition.

For example, you can configure a static metric symptom where, when the virtual machine CPU workload is greater than 90, a critical symptom is triggered.

#### Dynamic Thresholds

Metric symptoms that are based on dynamic thresholds compare the currently collected metric value against the trend identified by VMware Aria Operations, evaluating whether the current value is above, below, or generally outside the trend.

For example, you can configure a dynamic metric symptom where, when the virtual machine CPU workload is above the trended normal value, a critical symptom is triggered.

The Metric Symptom Definitions is a list of the metric-based symptoms defined in your VMware Aria Operations environment. You use the information in the list to evaluate the defined metric threshold triggering states and determine if you want to add, edit, or clone symptoms.

### Property Symptom Definitions

You define symptoms based on properties so that you can create alert definitions that let you know when changes to properties on your monitored objects can affect the behavior of the objects in your environment.

The Property Symptom Definitions is a list of the property-based symptoms in your VMware Aria Operations environment. You use the information in the list to evaluate the defined property triggering states and determine whether to add, edit, or clone symptoms.

### Where You Find Metric/Property Symptoms

To manage symptoms based on metrics/properties, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions > Metric / Property**.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

### Table 12-9. Metric/Property Symptoms Options

| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.<br>■ Add. Add a symptom definition.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge.<br>■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge.<br>■ Clone. Create a copy of the selected symptom definition.<br>■ Export. Downloads the symptom definition.<br>■ Import. Allows you to import symptom definitions. To import:<br>  ■ Click the **Import** option from the horizontal ellipsis.<br>  ■ Click **Browse** and select the file to import.<br>  ■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>  ■ Click **Import** to import the symptom definition, and click **Done**. |
| Filter options | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Name | Descriptive name of the symptom. |
| Criticality | Severity of the symptom when it is triggered. |
| Object Type | Base object type against which the symptom is defined. |
| Metric Name | Text string that is used as a reference key for the metric. You can use the metric key to locate additional information about how the system statistics are derived from the metric. |

Table 12-9. Metric/Property Symptoms Options (continued)

| Option | Description |
| --- | --- |
| Operator | Operator used to compare the current value to the threshold value, and trigger the symptom. |
| Value | Text string that is the compared value for the property. |
| Defined By | Indicates whether the symptom was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the symptom was last modified. |
| Modified By | Displays the name of the user who last modified the symptom. |

## Defining Metric/Property Symptoms

A metric symptom is triggered when a metric is compared to the configured static or dynamic thresholds, and the symptom condition is evaluated as true in VMware Aria Operations.

### Defining Metric Symptoms

If the symptom is based on a static threshold, the metric is compared based on the configured operator and the provided numeric value. If the symptom is based on a dynamic threshold, the metric is compared based on whether the current value is above, below, or abnormal compared to the calculated trend value.

1 To define symptoms based on metrics, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions > Metric / Property**. Click **Add** to define a metric-based symptom in the workspace.

2 Enter the following details.

| Option | Description |
| --- | --- |
| Base Object Type | Object against which the symptom is evaluated.<br>Based on the select object type, the list of available metrics displays only the metrics applicable to the object type. |
| Symptom Type | Select **Metrics** from the Symptom Type drop-down list. |
| Select Specific Object | If a metric or supermetric is not listed in the common metric or supermetric list, based on the selected based object type, use Select Object to inspect the metrics or supermetrics of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a metric or supermetric for a specific object, the symptom definition is applicable to all objects with that metric or supermetric in your environment. |
| Search | Use a word search to limit the number of items that appear in the list. |

| Option | Description |
| --- | --- |
| Metric list | List of metrics for the selected base object type.<br>Click and drag the metric to the left pane. |
| Symptom Definition workspace<br>You can define symptoms based on static or dynamic thresholds. | |
| Threshold | Determines if the symptom is static or dynamic.<br>■ Static thresholds are fixed values that trigger symptoms as true. You can configure one threshold for each symptom. You can also create multiple symptoms for multiple thresholds.<br>For example, configure one symptom where the CPU use is greater than 90 percent and another where the CPU usage is less than 40 percent. Each is a separate symptom and can be added individually to an alert definition.<br>■ Dynamic thresholds are based on VMware Aria Operations trended data where the triggering value is determined through the analytics. If the current value of the metric does not fall in the trended range, the symptom is triggered. |

| Option | Description |
| --- | --- |
| Static Threshold configuration options | If you select Static Threshold, configure the options for this threshold type.<br>■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Condition. Determines how the value you specify in the value text box is compared to the current value of the metric when the symptom is evaluated.<br>■ Value. Value that the condition evaluates.<br>■ Criticality. Severity of the symptom when it is triggered.<br>■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.<br>■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.<br>■ Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered.<br>■ Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking **Select Specific Object** next to the search box.<br>■ To enable the metric threshold evaluation based on real-time data collected every 20 seconds, click the **Near Real-Time Monitoring** check box.<br>**Note** If alerts have symptom definitions with near real-time monitoring enabled, the alerts will be triggered based on the 20 seconds cycle instead of five minutes. The near real-time data collection should be enabled through all cloud accounts. This setting can be controlled through symptom overrides in policies. |
| Dynamic Threshold configuration options | If you select Dynamic Threshold, configure the options for this threshold type. |

| Option | Description |
| --- | --- |
| | ■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Condition. Relationship of the current value to trended range based on the following options:<br>  ■ Above Threshold. If current value is above trended range, the symptom is triggered.<br>  ■ Below Threshold. If the current value is below the trended range, the symptom is triggered.<br>  ■ Abnormal. If the current value is either above or below the trended range, the symptom is triggered.<br>■ Criticality. Severity of the symptom when it is triggered.<br>■ Evaluate on instanced metrics. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for CPU usage, when the check box is not selected, the symptom is triggered based on the object's CPU usage. However, if you select the check box, the system also evaluates CPU usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered.<br>■ Exclude the following instances of the metric. To exclude specific instanced metrics from the symptom, drag the metric instances from the left pane. If you cannot locate the metric instance you want to exclude, you can search for it in another object that uses the metric by clicking **Select Object** next to the **Metrics** field. |

3 Click **Save**.

### Defining Property Symptom

A property symptom is triggered when the defined threshold is compared with the current property value and the comparison is evaluated as true.

1 To define symptoms based on properties, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions > Metric / Property**. Click **Add** to define a property-based symptom in the workspace.

2 Enter the following details.

| Option | Description |
| --- | --- |
| Base Object Type | Object against which the symptom is evaluated.<br>Based on the selected object type, the list of available properties displays only the properties applicable to the object type. |
| Symptom Type | Select **Properties** from the Symptom Type drop-down list. |
| Select Specific Object | If a property is not listed in the common properties list, based on the selected based object type, use Select Object to inspect the properties of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a property for a specific object, the symptom definition is applicable to all objects with that property in your environment. |
| Search | Use a word search to limit the number of items that appear in the list. |
| Property list | List of properties for the selected base object type.<br>Click and drag the property to the left pane. |

| Option | Description |
| --- | --- |
| Symptom Definition workspace<br>The properties are configured values that are compared to the value you specify. You can configure a single property symptom or add multiple symptoms.<br>For example, if you need an alert when a particular property, such as Memory Hot Add, is no longer at the value required, you can configure a symptom and add it to an alert definition. | |
| Property | Configure the options:<br>■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Condition. Determines how the value you specify in the value text box is compared to the current value of the property for an object when the symptom definition is evaluated.<br>■ Value. Value that the condition evaluates.<br>■ Criticality. Severity of the symptom when it is triggered.<br>■ Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom is triggered. The default value is 1, which means that the symptom is triggered in the same collection cycle when the condition became true.<br>■ Cancel Cycle. The symptom is canceled after the trigger condition is false for this number of collection cycles after which the symptom is cancelled. The default value is 1, which means that the symptom is canceled in the same cycle when the condition becomes false.<br>■ Evaluate on instanced properties. Select this check box so that the system evaluates the object level symptom as well as the instance level symptom. For example, for memory usage, when the check box is not selected, the symptom is triggered based on the object's memory usage. However, if you select the check box, the system also evaluates memory usage of each of the cores. If any of the cores is found to be crossing the threshold, the symptom is triggered.<br>■ Drop instances to exclude. To exclude specific instanced properties from the symptom, drag the property instances from the right pane. If you cannot locate the property instance you want to exclude, you can search for it in another object that uses the property by clicking **Select Specific Object** next to the search box. |

3 Click **Save**.

## Message Event Symptoms

Message event symptoms are based on events received as messages from a component of VMware Aria Operations or from an external monitored system through the system's REST API. You define symptoms based on message events to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

The adapters for the external monitored systems and the REST API are inbound channels for collecting events from external sources. Adapters and the REST server both run in the VMware Aria Operations system. The external system sends the messages, and VMware Aria Operations collects them.

You can create message event symptoms for the supported event types. The following list is of supported event types with example events.

- System Performance Degradation. This message event type corresponds to the EVENT_CLASS_SYSTEM and EVENT_SUBCLASS_PERFORM_DEGRADATION type and subtype in the VMware Aria Operations API SDK.
- Change. The VMware adapter sends a change event when the CPU limit for a virtual machine is changed from unlimited to 2 GHz. You can create a symptom to detect CPU contention issues as a result of this configuration change. This message event type corresponds to the EVENT_CLASS_CHANGE and EVENT_SUBCLASS_CHANGE type and subtype in the VMware Aria Operations API SDK.
- Environment Down. The VMware Aria Operations adapter sends an environment down event when the collector component is not communicating with the other components. You can create a symptom that is used for internal health monitoring. This message event type corresponds to the EVENT_CLASS_ENVIRONMENT and EVENT_SUBCLASS_DOWN type and subtype in the VMware Aria Operations API SDK.
- Notification. This message event type corresponds to the EVENT_CLASS_NOTIFICATION and EVENT_SUBCLASS_EXTEVENT type and subtype in the VMware Aria Operations API SDK.

### Where You Find Message Event Symptoms

To manage symptoms based on message events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**. Select the **Message Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

The Message Event Symptom Definitions is a list of the message event-based symptoms defined in your VMware Aria Operations environment. You use the information in the list to evaluate the defined message events and to determine if you want to add, edit, or clone symptoms.

| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.<br>■ Add. Add a symptom definition.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge.<br>■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge.<br>■ Clone. Create a copy of the selected symptom definition.<br>■ Export and Import. Export the file as xml from one VMware Aria Operations so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file. |
| Filter options | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Name | Descriptive name of the symptom. |
| Adapter Type | Adapter type for which the symptom is configured. |
| Object Type | Base object type against which the symptom is defined. |
| Event Type | Defined event classification type. |
| Operator | Operator used to compare the message from the incoming event against the event message specified in the symptom. |
| Event Message | Text string that is compared to the message in the incoming event using the specified operator. |
| Criticality | Severity of the symptom when it is triggered. |
| Defined By | Indicates whether the symptom was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the symptom was last modified. |
| Modified By | Displays the name of the user who last modified the symptom. |

### Defining Message Events

You can define message event systems in VMware Aria Operations so that you can create one or more of the symptoms that you can add to an alert definition. A message event symptom is triggered when a message in an incoming event matches the text string in the symptom, based on the specified operator.

1 To define symptoms based on message events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**.

2 Select the **Message Event** tab, and then click **Add**.

**Note** **Near Real-Time Monitoring** is triggered automatically if the adapter instance has real-time monitoring enabled.

3 Enter the following details.

Table 12-10. Symptoms Workspace Options for Message Events

| Option | Description |
| --- | --- |
| Based Object Type | Object against which the symptom is evaluated. |
| Symptom Type | Displays the symptom type as **Message Event**. |
| Select the Type of Event | Select the type of incoming event against which you are matching the events as they arrive. The incoming event must contain the following type and subtype combinations.<br>■ System Degradation<br>■ Change<br>■ Environment<br>■ Notification<br>■ Data Availability<br>■ Collector Down<br>■ Object Error<br>Click and drag the message event to the left pane. |

Table 12-10. Symptoms Workspace Options for Message Events (continued)

| Option | Description |
| --- | --- |
| Symptom Definition workspace<br>The Message Event text string is compared to the message in the incoming event by using the specified operator. You can configure a single message event symptom or add multiple symptoms.<br>For example, the VMware adapter sends a change event when the CPU limit for a virtual machine was changed from unlimited to 2 GHz. You can create a symptom to detect CPU contention issues as a result of this configuration change. | |
| Message Event | Configure the options:<br>■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Condition. Determines how the value you specify in the value text box is compared to the current value of the property for an object when the symptom definition is evaluated.<br>■ Value. Value that the condition evaluates.<br>■ Criticality. Severity of the symptom when it is triggered. |

4 Click **Save**.

## Fault Symptoms

Fault symptoms are based on events published by monitored systems. VMware Aria Operations correlates a subset of these events and delivers them as faults. Faults are intended to signify events in the monitored systems that affect the availability of objects in your environment. You define symptoms based on faults to include in alert definitions that use these symptoms. When the configured symptom condition is true, the symptom is triggered.

You can create fault symptoms for the supported published faults. Some object types have multiple fault definitions from which to choose, while others have no fault definitions.

If the adapter published fault definitions for an object type, you can select one or more fault events for a given fault while you define the symptom. The symptom is triggered if the fault is active because of any of the chosen events. If you do not select a fault event, the symptom is triggered if the fault is active because of a fault event.

### Where You Find Fault Symptoms

To manage symptoms based on fault message events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**. Select the **Fault** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

The Fault Symptom Definitions is a list of the fault-based symptoms defined in your VMware Aria Operations environment. You use the information in the list to evaluate the defined fault message events and to determine whether to add, edit, or clone symptoms.

| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.<br>■ Add. Add a symptom definition.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge.<br>■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge.<br>■ Clone. Create a copy of the selected symptom definition.<br>■ Export and Import. Export the file as xml from one VMware Aria Operations so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file. |
| Filter options | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Name | Descriptive name of the symptom. |
| Adapter Type | Adapter type for which the symptom is configured. |
| Object Type | Base object type against which the symptom is defined. |
| Fault | Selected fault based on object type. |
| Defined By | Indicates whether the symptom was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the symptom was last modified. |
| Modified By | Displays the name of the user who last modified the symptom. |

### Defining Fault Symptoms

You can define fault symptoms in VMware Aria Operations that are based on events published by the monitored systems, so that you can add one or more symptoms to an alert definition. A fault symptom is triggered when a fault is active on the base object because of the occurrence of any of the fault events selected in the symptom definition.

1 To define symptoms based on fault message events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**.

2 Select the **Fault** tab and then, click **Add**.

3 Enter the following details.

| Option | Description |
| --- | --- |
| Based Object Type | Object against which the symptom is evaluated. |
| Symptom Type | Displays the symptom type as **Fault**. |
| Fault Definitions | Select the fault definition for the selected base object type. Some object types do not have fault definitions, and other types have multiple definitions.<br>Click and drag the fault definition to the left pane. |
| Symptom Definition workspace<br>The fault events are published events from monitored systems. You can configure a single fault event symptom or add multiple symptoms.<br>For example, if your base object is host and you drag the Hardware sensor fault for unknown type fault definition, you then select one of two text strings indicating a fault. | |
| Fault Symptom | Configure the options:<br>■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Value. Select one or more fault events that activate the fault. If you do not select a string, then any of the provided strings are evaluated.<br>■ Criticality. Severity of the symptom when it is triggered. |

4 Click **Save**.

## Metric Event Symptoms

Metric event symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not VMware Aria Operations.

Metric event symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that VMware Aria Operations is actively monitoring.

The metric event thresholds, which determine whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system, represent the type and subtype combination that is specified in the incoming metric event.

- Above Threshold. Corresponds to type and subtype constants EVENT_CLASS_HT and EVENT_SUBCLASS_ABOVE defined in the VMware Aria Operations API SDK.
- Below Threshold. Corresponds to type and subtype constants EVENT_CLASS_HT and EVENT_SUBCLASS_BELOW defined in the VMware Aria Operations API SDK.

- Equal Threshold. Corresponds to type and subtype constants EVENT_CLASS_HT and EVENT_SUBCLASS_EQUAL defined in the VMware Aria Operations API SDK.
- Not Equal Threshold. Corresponds to type and subtype constants EVENT_CLASS_HT and EVENT_SUBCLASS_NOT_EQUAL defined in the VMware Aria Operations API SDK.

### Where You Find Metric Event Symptoms

To manage symptoms based on metric events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**. Click the **Metric Event** tab.

You can also define symptoms as you are defining alerts in the Alert Definition Workspace.

The Metric Event Symptom Definitions is a list of the metric event-based symptoms defined in your VMware Aria Operations environment. You use the information in the list to evaluate the defined threshold triggering states for the metric events and to determine if you want to add, edit, or clone symptoms.

| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your symptoms. You can select multiple symptoms using Ctrl+click or Shift+click.<br>■ Add. Add a symptom definition.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected symptom definition. Any changes you make affect the alert definitions that include this symptom. You cannot edit a symptom that manages a badge.<br>■ Delete. Remove the selected symptom definition. You cannot delete an alert that is used in an alert definition. To delete a symptom, you must first remove it from the alert definitions in which it is used. You cannot delete a symptom that manages a badge.<br>■ Clone. Create a copy of the selected symptom definition.<br>■ Export and Import. Export the file as xml from one VMware Aria Operations so that you can import the file on another instance. When you import the file, if you encounter a conflict, you can override the existing file or not import the new file. |
| Filter options | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Name | Descriptive name of the symptom. |
| Adapter Type | Adapter type for which the symptom is configured. |
| Object Type | Base object type against which the symptom is defined. |
| Event Metric | Selected event metric based on object type. |

| Option | Description |
| --- | --- |
| Event Type | Specifies whether the metric was above, below, equal to, or not equal to the threshold set by the monitoring system. |
| Defined By | Indicates whether the symptom was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the symptom was last modified. |
| Modified By | Displays the name of the user who last modified the symptom. |

### Defining Metric Event Symptoms

You can define metric event symptoms that are based on reported violations of metric thresholds from monitored systems, so that you can create one or more of the symptoms that you can add to an alert definition in VMware Aria Operations.

1 To define symptoms based on metric events, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Symptom Definitions**.

2 Select the **Metric Event** tab and then, click **Add**.

3 Enter the following details.

| Option | Description |
| --- | --- |
| Base Object Type | Select the base object type.<br>Object against which the symptom is evaluated.<br>Based on the select object type, the list of available metrics displays only the metrics applicable to the object type. |
| Symptom Type | Displays the symptom type as **Metric Event**. |
| Search | Use a word search to limit the number of items that appear in the list. |
| Select Specific Object | If a metric or supermetric is not listed in the common metric or supermetric list, based on the selected based object type, use Select Object to inspect the metrics or supermetrics of a selected object so that you can locate the property that you must use to create the symptom. Even though you select a metric or supermetric for a specific object, the symptom definition is applicable to all objects with that metric or supermetric in your environment. |
| Metric Event list | Select the metric event for the selected base object type.<br>Click and drag the metric to the left pane. |

| Option | Description |
| --- | --- |
| Symptom Definition Workspace<br>You can configure a single threshold or add multiple thresholds.<br>For example, configure a symptom where, when the virtual machine CPU usage is above the threshold defined in the monitored system, the metric event is above the threshold on the system. | |
| Metric Event | Configure the options:<br>■ Symptom name. Name of the symptom as it appears in the symptom list when configuring an alert definition, as it appears when the alert is generated, and when viewing triggered symptoms.<br>■ Condition. Select whether the metric is above, below, equal to, or not equal to the threshold set on the monitored system.<br>■ Criticality. Severity of the symptom when it is triggered.<br>Click **Save**. |

4 Click **Save**.

## Understanding Negative Symptoms for VMware Aria Operations Alerts

Alert symptoms are conditions that indicate problems in your environment. When you define an alert, you include symptoms that generate the alert when they become true in your environment. Negative symptoms are based on the absence of the symptom condition. If the symptom is not true, the symptom is triggered.

To use the absence of the symptom condition in an alert definition, you negate the symptom in the symptom set.

All defined symptoms have a configured criticality. However, if you negate a symptom in an alert definition, it does not have an associated criticality when the alert is generated.

All symptom definitions have a configured criticality. If the symptom is triggered because the condition is true, the symptom criticality will be the same as the configured criticality. However, if you negate a symptom in an alert definition and the negation is true, it does not have an associated criticality.

When negative symptoms are triggered and an alert is generated, the effect on the criticality of the alert depends on how the alert definition is configured.

The following table provides examples of the effect negative symptoms have on generated alerts.

Table 12-11. Negative Symptoms Effect on Generated Alert Criticality

| Alert Definition Criticality | Negative Symptom Configured Criticality | Standard Symptom Configured Criticality | Alert Criticality When Triggered |
|---|---|---|---|
| Warning | One Critical Symptom | One Immediate Symptom | Warning. The alert criticality is based on the defined alert criticality. |
| Symptom Based | One Critical Symptom | One Warning Symptom | Warning. The negative symptom has no associated criticality and the criticality of the standard symptom determines the criticality of the generated alert. |
| Symptom Based | One Critical Symptom | No standard symptom included | Info. Because an alert must have a criticality and the negative alert does not have an associated criticality, the generated alert has a criticality of Info, which is the lowest possible criticality level. |

# Recommendations in VMware Aria Operations

Recommendations are probable solutions for an alert generated in VMware Aria Operations. You can create a library of recommendations that include instructions to your environment administrators or actions that they can run to resolve an alert.

Recommendations provide your network engineers or virtual infrastructure administrators with information to resolve alerts.

Depending on the knowledge level of your users, you can provide more or less information, including the following options, in any combination.

- One line of instruction.
- Steps to resolve the alert on the target object.
- Hyperlink to a Web site, runbook, wiki, or other source.
- Action that makes a change on the target object.

When you define an alert, provide as many relevant action recommendations as possible. If more than one recommendation is available, arrange them in priority order so that the solution with the lowest effect and highest effectiveness is listed first. If no action recommendation is available, add text recommendations. Be as precise as possible when describing what the administrator should do to fix the alert.

# Where You Find Recommendations

To define recommendations, click **Configure** > **Alerts**, and then in the right pane, click **Recommendations**.

You can also define recommendations when you create an alert definition.

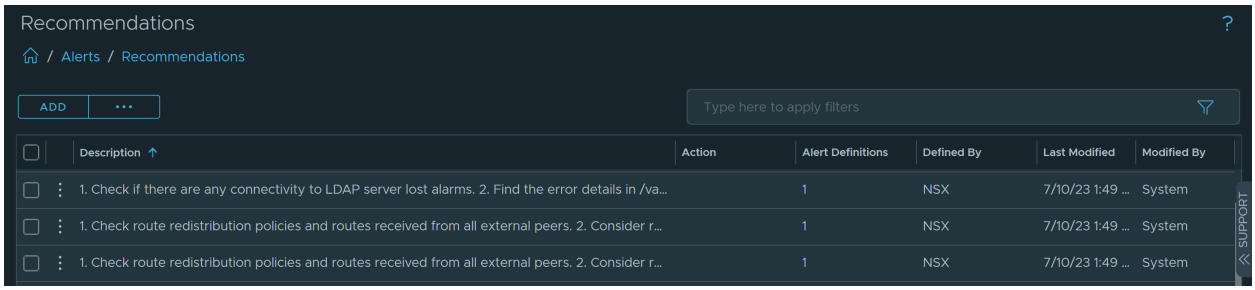


| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your recommendations.<br>■ Add. Add a recommendation.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Edit. Modify the selected recommendation.<br>■ Delete. Remove the selected recommendation.<br>■ Clone. Create a copy of the selected recommendation so that you can create a new recommendation that uses the current one.<br>■ Export. Downloads the recommendations.<br>■ Import. Allows you to import recommendations. To import:<br>  ■ Click the **Import** option from the horizontal ellipsis.<br>  ■ Click **Browse** and select the file to import.<br>  ■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>  ■ Click **Import** to import the recommendation, and click **Done**. |
| Quick Filter | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Description | Displays the recommendation text that is provided when the alert is generated.<br>Click this link to view the Details Page. On this page, you can view the alert definitions assigned for a particular recommendation. To remove the selected recommendation from all alert definitions, click **Edit** and then, in the **Edit Recommendation** page, click **Remove from all**. |
| Action | If the recommendation includes running an action, the name of the action is displayed. |

| Option | Description |
| --- | --- |
| Alert Definitions | Displays the number of alert definitions assigned for a particular recommendation. |
| Defined By | Indicates whether the recommendation was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the recommendation was last modified. |
| Modified By | Displays the name of the user who last modified the recommendation. |

## Defining Recommendations for Alerts

You can create recommendations that are solutions to alerts generated in VMware Aria Operations. The recommendations are intended to ensure that your network operations engineers and virtual infrastructure administrators can respond to alerts as quickly and accurately as possible. A recommendation can contain links to useful Web sites or local runbooks, instructions as text, or actions that you can initiate from VMware Aria Operations.

1 To define recommendations, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Recommendations**.

2 Click **Add** and enter the following details.

| Option | Description |
| --- | --- |
| Description | Enter the description of what must be done to resolve the triggered alert.<br>The description can include steps a user must take to resolve the alert or it might be instructions to notify a virtual infrastructure administrator.<br>This is a text field. |
| Create a hyperlink | To create a hyperlink, enter a description, select the text, and click the Hyperlink icon to make the text a hyperlink to a Website or local wiki page. |
| Action (Optional) | |
| Adapter Type | Select an adapter type from the drop-down list to narrow down the list of actions displayed in the Actions field. |
| Action | You can add an action as a method to resolve a triggered symptom or a generated alert. Actions must already be configured in VMware Aria Operations.<br>You must provide text in the text box to describe the action before you can save the recommendation. |

3 Click **Save**.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`
- `Set CPU Count for VM Power Off Allowed`
- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

## Notifications in VMware Aria Operations

Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside VMware Aria Operations. You can configure notification rules for the supported outbound alerts so that you can filter the alerts that are sent to the selected external system.

You can use the notifications list to manage your rules and then use the notification rules to limit the alerts that are sent to the external system. To use notifications, the supported outbound alert plug-ins must be added and running.

With notification rules, you can limit the data that is sent to the following external systems.

- Standard Email. You can create multiple notification rules for various email recipients based on one or more of the filter selections. If you add recipients but do not add filter selections, all the generated alerts are sent to the recipients.
- REST. You can create a rule to limit alerts that are sent to the target REST system so that you do not need to implement filtering on that target system.
- SNMP Trap. You can configure VMware Aria Operations to log alerts on an existing SNMP Trap server in your environment.
- Log File. You can configure VMware Aria Operations to log alerts to a file on each of your VMware Aria Operations nodes.

You configure notification options to specify which alerts are sent out for the Standard Email, REST, SNMP, and Log File outbound alert plug-ins. For the other plug-in types, all the alerts are sent when the target outbound alert plug-in is activated.

The most common outbound alert plug-in is the Standard Email plug-in. You configure the Standard Email plug-in to send notifications to one or more users when an alert is generated that meets the criteria you specify in the notification settings.

# Where You Find Notifications

To manage your notifications, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Notifications**.

**Note** To use notifications, the supported outbound alert plug-ins must be added and running.

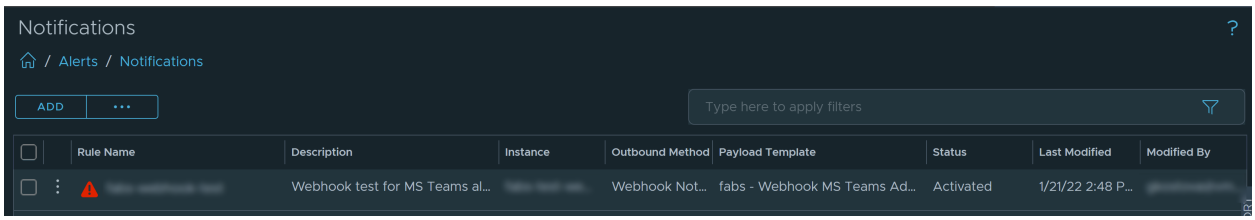


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your notification rules.<br>■ Add. Opens the Add Rule dialog box where you configure the filtering options for the notification rule.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Delete. Removes the selected rule.<br>■ Deactivate or Activate. Deactivates or activates the selected rule(s).<br>■ Export or Import. Export the selected notifications to a ".xml" file so that you can import it on another VMware Aria Operations instance. |
| Quick Filter (Action Name) | Limits the list to actions matching the filter. You can filter by:<br>■ Rule Name<br>■ Instance<br>■ Status<br>■ Modified By |
| Rule Name | Name you assigned when you created the notification rule. Click the vertical ellipsis to perform the following actions.<br>■ Edit. Allows you to edit the selected rule.<br>■ Clone. Allows you to clone an existing notification rule and edit the attributes of the cloned notification rule. You can create multiple alert notification rules so that you can send the same alert notifications to different outbound settings.<br>**Note** You can clone only one alert notification rule at a time.<br>■ Delete. Removes the selected rule.<br>■ Deactivate or Activate. Deactivates or activates the selected rule.<br>■ Export. Downloads the notification setting. |
| Description | Description of the notification rule. |
| Instance | Name of the configured outbound alert instance for the notification rule.<br>Instances are configured as part of the outbound alerts and can indicate different email servers or sender addresses for alert notifications. |
| Outbound Method | Displays the type of the outbound method that is configured. |
| Payload Template | Displays the payload template that is used. |
| Status | Displays if the rule is activated or not. |

| Option | Description |
| --- | --- |
| Email Address | If the rule is for standard email notifications, the alert recipient email addresses are listed. |
| Object Name | If the rule specifies a notification for a particular object, the object name is listed. |
| Children | If the rule specifies a notification for a particular object and selected child objects, the child object types are listed. |
| Last Modified | Displays the date on which the rule was last modified. |
| Modified By | Displays the name of the user who last modified the rule. |

## Creating Notification Rules for Alerts

You add, manage, and edit your notification rules in VMware Aria Operations. To send notifications to a supported system, you must configure and activate the settings for outbound alerts.

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances. For details on configuring outbound plug-ins, see Adding Outbound Notification Plug-Ins in VMware Aria Operations.

You use the Notifications page to manage your alert notification rules. The rules determine which VMware Aria Operations alerts are sent to the supported target systems.

Notification rules are filters that limit the data sent to external systems by using outbound alert plug-ins that are supported, configured, and running. Rather than sending all alerts to all your email recipients, you can use notification rules to send specific alerts. For example, you can send health alerts for virtual machines to one or more of your network operations engineers. You can send critical alerts for selected hosts and clusters to the virtual infrastructure administrator for those objects. Before you can create and manage notification rules, you must configure the outbound alert plug-in instances.

1 To manage your notifications, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Notifications**. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.

2 Enter the following notification details.

| Option | Description |
| --- | --- |
| Name | Name of the rule that you use to manage the rule instance. |
| Description | Description of the rule. |
| Notification Status | Either activate or deactivate a notification setting. Deactivating a notification will stop the alert notification for that setting and activating it will activate the alert notification. |

3 Click **Next**.

4 Define criteria for the notification rule.

| Option | Description |
| --- | --- |
| Object Scope | |
| Criteria | Object Type, Object, Tags, Applications, and Tiers for which you are filtering the alert notifications.<br>After you select the type, you select the specific instance. For example, if you select **Object**, you then select the specific object by name and determine whether to include any child objects. |
| Alert Scope | |
| Category | Alert Types/Subtypes, Alert Impact, or Alert Definition that triggers the alert.<br>After you select the criteria, you can configure the specific selections associated with the criteria. For example, if you select **Alert Definition**, you then select the alert definition that limits the data to alerts with this definition. You can select multiple alert definitions as conditions for a notification to trigger. |
| Criticality | Defined criticality of the alert that results in the data being sent to an external system. For example, if you select **Critical**, then the data that is sent to the external system must also be labeled as critical. |
| Control State | State of the alert, either opened, assigned, or suspended. |
| Notify On | |
| Status | Current state of the alert, either canceled, updated, or new. |
| Notification Heartbeat | |
| Heartbeat | Set this to **Active** if you want to send repeat notifications for the active alerts. The frequency of the notification depends on the collection interval set for the adapter whose object is being evaluated. By default, this checkbox is not selected.<br>**Note** Setting this option to **Active** can cause a potential surge in the notifications that are generated. |
| Advanced Filters: By Collector | |
| Collector/Group | Select a collector or group if you want to receive notifications for the objects that receive data from the selected collector/group. |
| **Note** If you do not define any alert filters in the Define Criteria tab, then the notification will be sent for all the alerts without applying any conditions for the object scope, alert scope, or alert state. | |

5 Click **Next**.

6 Select the outbound method that you want to use to send your notification.

| Option | Description |
| --- | --- |
| Outbound Method | ■ Select Plug-In Type: Type of plugin. Select one of the outbound alert plug-in types: Standard Email Plugin, Rest Notification Plugin, Webhook Notification Plugin, Service-Now Notification Plugin, and Slack Plugin.<br>■ Select Instance: Select the configured instance for the type of plug-in.<br>■ Create New Instance: You can also create a new outbound instance for the plug-in type you select.<br>For details, see Adding Outbound Notification Plug-Ins in VMware Aria Operations. |

7 Click **Next**.

8 Select the payload template.

| Option | Description |
| --- | --- |
| Payload Template | Select the payload template that you want to include in the notification. Each plug-in has its default template and you can select the default template if no customization is required. The template includes additional information about the alert or the object that is displayed in the notification. You can also customize your payload for a Webhook Notification Plugin. For details on creating payload templates, see Creating Payload Templates for Outbound Plugins. |
| The values in this tab differ based on the outbound plug-in you have selected in the previous step. | |

| Option | Description |
| --- | --- |
| Outbound Method -Standard Email Plugin | If you are configuring notifications for standard email, you can add recipients and associated information.<br>■ Recipient(s). Enter the email addresses of the individuals to whom you are sending email messages that contain alert notifications. If you are sending to more than one recipient, use a semicolon (;) between addresses.<br>■ Cc Recipients. Enter the email addresses of the individuals that have to be cc'd for the email.<br>■ Bcc Recipients. Enter the email addresses of the individuals that have to be bcc'd for the email.<br>■ Notify again. Number of minutes between notifications messages for active alerts. Leave the text box empty to send only one message per alert.<br>■ Max Notifications. Number of times to send the notification for the active alert. Leave the text box empty to send only one message per alert.<br>■ Delay to notify. Number of minutes to delay before sending a notification when a new alert is generated. For example, if the delay is 10 minutes and a new alert is generated, the notification is not sent for 10 minutes. If the alert is canceled in those 10 minutes, the notification is not sent. The notification delay reduces the number of notifications for alerts that are canceled during that time.<br>■ Description. Enter the text to include in the email message. For example, the Attention Host Management team. |

| Option | Description |
| --- | --- |
| Outbound Method - Service-Now Notification Plugin | If you are configuring notifications for a Service-Now notification plug-in, you can add instances and associated information.<br>■ Caller. Enter the name of the person who reported the incident or who is affected by the incident.<br>■ Category. Specify the category to which the incident belongs.<br>■ Sub Category. Specify the sub-category to which the incident belongs.<br>■ Business Service. Specify the business service of the incident.<br>■ Contact Type. Enter the contact type.<br>■ State. Enter the incident state in digits.<br>■ Resolution Code. Enter the resolution code for the incident.<br>■ Resolution notes. Enter the resolution notes for the incident.<br>■ On hold reason. Enter the reason as to why the incident is on hold.<br>■ Impact. Set the incident impact in digits. Impact measures the business criticality of the affected service.<br>■ Urgency. Set urgency for the incident in digits. Urgency defines the number of days taken to resolve an incident.<br>■ Priority. Enter the priority for the incident. Priority defines the sequence in which the incident must be resolved.<br>■ Assignment Group. Enter the assignment group for the incident.<br>■ Assigned To. Enter the details of the person to whom the incident is assigned.<br>■ Severity. Set the severity for the incident in digits.<br>■ Upon Approval. Specify the next steps to be taken upon incident approval.<br>■ Problem. Enter the details of the related problem if it exists.<br>■ Cause by change. Enter the change request which triggered the incident.<br>■ Change Request. Enter the details for the related change list if it exists. |
| Outbound Method - Slack Plugin | If you are configuring notifications for a Slack plugin, add the Webhook URL of Slack. For example, the Webhook URL is in the format: **`https://hooks.slack.com/services/T00000000/B00000000/XXXXXXXXXXXXXXXXXXXXXXXX`**.<br>Create and authorize an app within Slack to obtain the Webhook URL. For details on creating and authorizing an app within Slack, refer to the Slack Documentation. |

| Option | Description |
| --- | --- |
| | Once you have created the notification rule, the alerts are displayed within that particular Slack channel with a link to the alert. Click the link to view the details of the alert on the Object Summary page. |

9 Click **Next** to test the notification

10 In the **Test Notification** tab, click **Initiate Process** to initiate the notification configuration validation process.

11 If you activate the **Filter the alert definitions and objects based on criteria outlined in the Define Criteria section** option, then the alert definitions and objects displayed below are based on the criteria outlined in the **Define Criteria** section.

12 Select an alert definition and an object for validation.

13 Click **Validate Configuration**.

- Notification Validation Steps: View the steps involved in validating the notification configuration. The steps also indicate an error, if any. The validation steps differ based on the selected Outbound method.

| Outbound Method | Validation Steps |
| --- | --- |
| Standard Email Plugin | ■ Validate Criteria<br>■ Establish Endpoint Connection<br>■ Certificates<br>■ Authentication<br>■ Send Notification<br>**Note** The Standard Email Plugin does not provide a response that can be validated in VMware Aria Operations. |
| SNMP Trap Plugin | ■ Validate Criteria<br>■ Establish Endpoint Connection<br>■ Send Notification<br>**Note** The SNMP Trap Plugin does not provide a response that can be validated in VMware Aria Operations. |
| Webhook Notification Plugin | ■ Validate Criteria<br>■ Establish Endpoint Connection<br>■ Certificates<br>■ Authentication<br>■ Send Notification<br>■ Endpoint Receives Notification |

| Outbound Method | Validation Steps |
| --- | --- |
| Slack Plugin | ■ Validate Criteria<br>■ Establish Endpoint Connection<br>■ Authentication<br>■ Send Notification<br>■ Endpoint Receives Notification |
| ServiceNow Notification Plugin | ■ Validate Criteria<br>■ Establish Endpoint Connection<br>■ Authentication<br>■ Send Notification<br>■ Endpoint Receives Notification |

- Response: The Response tab in the right pane displays if the test passed successfully or if there were any errors.
- Body: Displays the content of the notification.

14 Click **Create** to create the notification rule. You can view the rule you created under **Alerts > Notifications**.

## User Scenario: Create a VMware Aria Operations Email Alert Notification

As a virtual infrastructure administrator, you need VMware Aria Operations to send email notifications to your advanced network engineers when critical alerts are generated for mmbhost object, the host for many virtual machines that run transactional applications, where no one has yet taken ownership of the alert.

### Prerequisites

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Ensure that at least one instance of the Standard Email Plug-In is configured and running. See Add a Standard Email Plug-In for VMware Aria Operations Outbound Alerts.

### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Notifications**.

2 Click **Add** to add a notification rule.

3 In the **Name** text box, enter a name similar to **`Unclaimed Critical Alerts for mmbhost`**.

4 Set the **Notification Status**, you can either activate or deactivate a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.

5 In the **Define Criteria** tab, select the objects and alerts for which you want to receive notifications.

   a From the **Criteria** drop-down menu, select **Object**.

   b Locate and select the object from the list.

6 Configure the Alert Scope.

   a From the **Category** drop-down menu, select **Alert Impact**, and from the adjacent drop-down menu, select **Health**.

   b From the **Criticality** drop-down menu, select **Critical**.

7 In the Notify On section, select **Open** from the **Status** drop-down menu.

   The Open state indicates that no engineer or administrator has taken ownership of the alert.

8 In the **Set Outbound Method** tab, select **Standard Email Plug-In** from the **Outbound method** drop-down menu, and then select the configured instance of the email plug-in.

9 In the **Select Payload Template** tab, configure the email options.

   a In the **Recipients** text box, enter the email addresses of the members of your advance engineering team, separating the addresses with a semi-colon (;).

   b To send a second notification if the alert is still active after a specified amount of time, enter the number of minutes in the **Notify again** text box.

   c Type number of notifications that are sent to users in the **Max Notifications** text box.

10 Click **Create**.

#### Results

You created a notification rule that sends an email message to the members of your advance network engineering team when any critical alerts are generated for the mmbhost object and the alert is not claimed by an engineer. This email reminds them to look at the alert, take ownership of it, and work to resolve the triggering symptoms.

#### What to do next

Respond to alert email notifications.

## Notifications - User Scenario: Create a Webhook Alert Notification

As a virtual infrastructure administrator, you need VMware Aria Operations to send alerts in JSON or XML to a Webhook with any endpoint REST API that accepts these messages. You want only alerts where the virtualization alerts that affect availability alert types to go to an external application. You can then use the provided information to initiate a remediation process in that application to address the problem indicated by the alert. The notification configuration limits the alerts sent to the outbound alert instance to those matching the notification criteria.

### Prerequisites

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Ensure that at least one instance of the Webhook Notification Plugin is configured and running. See Add a Webhook Notification Plugin for Outbound Instance.

### Procedure

1. From the left menu, click **Configure > Alerts**, and then in the right pane, click **Notifications**.
2. Click **Add** to add a notification rule.
3. In the **Name** text box, enter a name similar to **`Virtualization Alerts for Availability`**.
4. Set the **Notification Status**, you can either activate or deactivate a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.
5. In the **Define Criteria** tab, select the objects and alerts for which you want to receive notifications.
   a. From the **Criteria** drop-down menu, select **Object**.
   b. Locate and select the object from the list.
6. Configure the Alert Scope.
   a. From the **Category** drop-down menu, select **Alert Type**, and from the **Alert Types/ Subtypes** menu, select **Availability** under **Virtualization/Hypervisor Alerts**.
   b. From the **Criticality** drop-down menu, select **Warning**.
7. In the Notify On section, select **New** from the **Status** drop-down menu.

   The New status indicates that the alert is new to the system and not updated.
8. In the **Set Outbound Method** tab, select **Webhook Notification Plugin** from the **Outbound method** drop-down menu, and then select the configured instance of the Webhook plugin.
9. In the **Select Payload Template** tab, select the **Default Webhook Template**.
10. Click **Create**.

### Results

You created a notification rule that sends the alert text to the target REST-activated system. Only the alerts where the configured alert impact is Virtualization/Hypervisor Availability and where the alert is configured as a warning are sent to the target instance using the Webhook plugin.

### What to do next

You created a notification rule that sends the alert text to the target Webhook system.

# Outbound Settings in VMware Aria Operations

You use the Outbound Settings to manage your communication settings so that you can send information to users or applications outside of VMware Aria Operations.

You manage your outbound options from this page, including adding or editing outbound plug-ins, and turning the configured plug-ins on or off. When activated, the plug-in sends a message to users as email notifications, or sends a message to other applications.

Outbound plug-in settings determine how the supported external notification systems connect to their target systems. You configure one or more instances of one or more plug-in types so that you can send data about generated notifications outside of VMware Aria Operations.

You configure each plug-in with the required information, including destination locations, hosts, ports, user names, passwords, instance name, or other information that is required to send notifications to those target systems. The target systems can include email recipients, log files, or other management products.

Some plug-ins are included with VMware Aria Operations, and others might be added when you add a management pack as a solution.

The configuration options vary depending on which plug-in you select from the **Plug-In Type** drop-down menu.

To add outbound notification plug-in, see Adding Outbound Notification Plug-Ins in VMware Aria Operations.

## Where You Find Outbound Settings

To manage your outbound settings, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

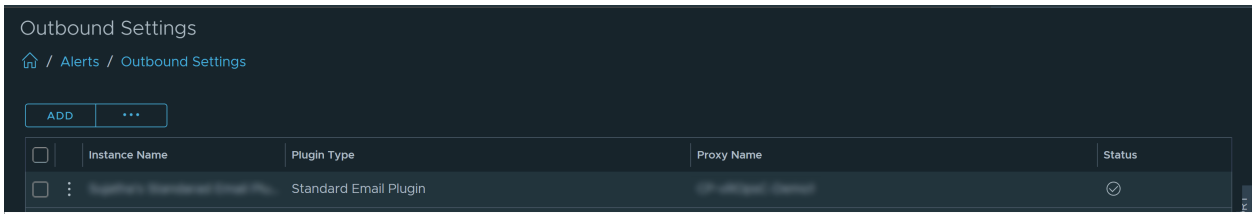

| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your Outbound Plug-Ins.<br>■ Add. Opens the Outbound Plug-In dialog box where you configure the connection options for the instance.<br>Select an existing plugin and click the vertical ellipsis to perform the following actions.<br>■ Edit. Modify the Outbound Plug-In instance details.<br>■ Delete. Removes the selected plug-in instance.<br>■ Activate or Deactivate. Starts or stops the plug-in instance. Deactivating an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.<br>■ Export. Downloads the outbound settings.<br>■ Import. Allows you to import outbound settings. To import:<br>  ■ Click the **Import** option from the horizontal ellipsis.<br>  ■ Click **Browse** and select the file to import.<br>  ■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>  ■ Click **Import** to import outbound settings, and click **Done**. |
| Instance Name | Name that you assigned when you created the plug-in instance.<br>Click the vertical ellipsis to perform the following actions.<br>■ Edit. Allows you to edit the selected payload template.<br>■ Delete. Removes the selected payload template.<br>■ Activate or Deactivate. Starts or stops the plug-in instance. Deactivating an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.<br>■ Export. Downloads the outbound settings. |
| Plug-In Type | Type of configured plug-in for the plug-in instance. The types of plug-ins vary depending on the solutions you have added to your environment. |
| Status | Specifies whether the plug-in is currently running. |

## List of Outbound Plugins

VMware Aria Operations provides outbound plug-ins. This list includes the name of the plug-in and whether you can filter the outbound data based on your notification settings.

If the plug-in supports configuring notification rules, then you can filter the messages before they are sent to the target system. If the plug-in does not support notifications, all messages are sent to the target system, and you can process them in that application.

If you installed other solutions that include other plug-in options, they appear as a plug-in option with the other plug-ins.

Messages and alerts are sent only when the plug-in is activated.

Table 12-12. Notification Support for Outbound plug-ins

| Outbound plug-in | Configure Notification Rules |
| --- | --- |
| Automated Action plug-in | No<br>The Automated Action plug-in is activated by default. If automated actions stop working, select the Automated Action plug-in and activate it if necessary. If you edit the Automated Action plug-in, you only have to provide the instance name. |
| REST Notification plug-in | Yes |
| Standard Email plug-in | Yes |
| SNMP Trap plug-in | Yes |
| Webhook Notification Plugin | Yes |
| Slack plug-in | Yes |
| Service-Now Notification plug-in | Yes |
| Vmware Hosted Email plug-in | No<br>The Vmware Hosted Email plug-in is activated by default. |

## Adding Outbound Notification Plug-Ins in VMware Aria Operations

You add outbound plug-in instances so that you can notify users about alerts or capture alert data outside of VMware Aria Operations.

You can configure one or more instances of the same plug-in type if you need to direct alert information to multiple target systems.

The Automated Action plug-in is activated by default. If automated actions stop working, check the Automated Action plug-in and activate it if necessary. If you edit the Automated Action plug-in, you only need to provide the instance name.

- Add a Standard Email Plug-In for VMware Aria Operations Outbound Alerts
- Add an SNMP Trap Plug-In for VMware Aria Operations Outbound Alerts

  You add an SNMP Trap plug-in when you want to configure VMware Aria Operations to log alerts on an existing SNMP Trap server in your environment.
- Add a Service-Now Notification Plug-In for Outbound Alerts

  You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with VMware Aria Operations. Service Now creates an incident whenever an alert is triggered in VMware Aria Operations.

- Notifications - Add a Slack Plugin for Outbound Notifications

  You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run VMware Aria Operations using alert links to troubleshoot further.

- Add a Webhook Notification Plugin for Outbound Instance

  You can integrate Webhook with any endpoint REST API and configure outbound payload.

- Sample Email Alert

  Here is a sample email for a newly created alert.

- VMware Hosted Email plug-in for Outbound Settings

  Use the VMware hosted email plug-in for alert notifications and reports delivery in VMware Aria Operations.

## Add a Standard Email Plug-In for VMware Aria Operations Outbound Alerts

You add a Standard Email Plug-In so that you can use Simple Mail Transfer Protocol (SMTP) to email VMware Aria Operations alert notifications to your virtual infrastructure administrators, network operations engineers, and other interested individuals. Alternately, you can also use your cloud proxy to communicate with your email server.

### Prerequisites

Ensure that you have an email user account that you can use as the connection account for the alert notifications. If you choose to require authentication, you must also know the password for this account.

### Procedure

1. From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.
2. Click **Add**, and from the **Plug-In Type** drop-down menu, select **Standard Email Plugin**.

   The dialog box expands to include your SMTP settings.
3. Optionally, click the **Use Proxy** check box and then select the cloud proxy from the **Collector/Group** drop-down menu to establish a connection.

   **Note** You must use a cloud proxy if your mail or rest server is located outside of AWS.
4. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.

5 Configure the SMTP options appropriate for your environment.

| Option | Description |
| --- | --- |
| **Use Secure Connection** | Activates secure communication encryption using SSL/TLS. If you select this option, you must select a method in the **Secure Connection Type** drop-down menu.<br>**Note** In VMware Aria Operations, the Standard Email Plugin always uses the 'starttls' command when you select the **Use Secure Connection** checkbox and select **TLS** as the **Secure Connection Type**. |
| **Requires Authentication** | Activates authentication on the email user account that you use to configure this SMTP instance. If you select this option, you must provide a password for the user account. |
| **SMTP Host** | URL or IP address of your email host server. |
| **SMTP Port** | Default port SMTP uses to connect with the server. |
| **Secure Connection Type** | Select either SSL/TLS as the communication encryption method used in your environment from the drop-down menu. You must select a connection type if you select Use Secure Connection.<br>**Note** In VMware Aria Operations, the Standard Email Plugin always uses the 'starttls' command when you select the **Use Secure Connection** checkbox and select **TLS** as the **Secure Connection Type**. |
| **Sender Email Address** | Email address that appears on the notification message. |
| **Sender Name** | Displayed name for the sender email address. |
| **Credential Type** | Select the **Credential Type** from the list. If your Endpoint URL does not need any authentication, then select **No Credential** from the **Credential Type** list. |
| **Credential** | Add or edit the **Credential** details. Click the plus icon to enter the details of the new credentials in the **Create New Credential** pane, and click **Save**.<br>This field appears only when you select **Basic Authentication** as the **Credential Type**.<br>**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The VMware Aria Operations administrator must assign the credentials from the **Orphan and Unassigned** page. For more information, see the topic, Managing Orphaned and Unassigned Content. |
| **Receiver Email Address** | Receiver's email address. |

6 Click **Save**.

7 To start the outbound alert service for this plug-in, select the instance in the list and click **Activate** on the toolbar.

### Results

This instance of the Standard Email Plug-In for outbound SMTP alerts is configured and running.

### What to do next

Create notification rules that use the Standard Email Plug-In to send a message to your users about alerts requiring their attention. See User Scenario: Create a VMware Aria Operations Email Alert Notification .

## Add an SNMP Trap Plug-In for VMware Aria Operations Outbound Alerts

You add an SNMP Trap plug-in when you want to configure VMware Aria Operations to log alerts on an existing SNMP Trap server in your environment.

You can provide filtering when you define a Notification using an SNMP Trap destination.

### Prerequisites

Ensure that you have an SNMP Trap server configured in your environment, and that you know the IP address or host name, port number, and community that it uses.

### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

2 Click **Add**, and from the **Plug-In Type** drop-down menu, select **SNMP Trap Plugin**.

   The dialog box expands to include your SNMP trap settings.

3 Enter an **Instance Name**.

4 (Optional) Click the **Use Proxy** check box and then select the cloud proxy from the **Collector/Group** drop-down menu to establish a connection.

   **Note** You must use a cloud proxy if your SNMP destination host is located outside of AWS.

5 Configure the SNMP trap settings appropriate to your environment.

| Option | Description |
|---|---|
| **Destination Host** | IP address or fully qualified domain name of the SNMP management system to which you are sending alerts. |
| **Port** | Port used to connect to the SNMP management system. Default port is 162. |
| **Community** | Text string that allows access to the statistics. SNMP Community strings are used only by devices that support SNMPv1 and SNMPv2c protocol. |
| **Username** | User name to configure SNMP trap settings in your environment. If the user name is specified, SNMPv3 is considered as the protocol by the plugin.<br>If left blank, SNMPv2c is considered as the protocol by the plugin.<br>**Note** SNMP uses User Datagram Protocol (UDP) as its transport protocol. |
| **Authentication Protocol** | Authentication algorithms available are SHA-224, SHA-256, SHA-384, SHA-512. |
| **Authentication Password** | Authentication password. |
| **Privacy Protocol** | Privacy algorithms available are AES192, AES256. |

| Option | Description |
| --- | --- |
| **Privacy Password** | Privacy password. |
| **Engine ID** | Engine ID serves as an identifier for the agent. It is used with a hashing function to generate localized keys for authentication and encryption of SNMP v3 messages.<br>It is mandatory to specify the Engine ID when configuring the SNMP Trap plugin. If you do not add the Engine ID and save the SNMP Trap plugin instance, the field is auto-generated the next time you edit the settings. |

6 Click **Test** to validate the connection.

**Note** The Community and Username options are mutually exclusive. Define either one of them to avoid an error. If you add a user name, you can optionally define the Authentication Protocol and Authentication Password followed by the Privacy Protocol and Privacy Password. The privacy protocol and its password cannot be defined independent of the authentication protocol and its password.

#### Results

This instance of the SNMP Trap plug-in is configured and running.

#### What to do next

When the plug-in is added, Creating Notification Rules for Alerts for receiving the SNMP traps.

## Add a Service-Now Notification Plug-In for Outbound Alerts

You add a Service-Now Notification plug-in when you want to integrate Service Now ticketing system with VMware Aria Operations. Service Now creates an incident whenever an alert is triggered in VMware Aria Operations.

Using Service-Now Notification Plug-In you can send alert notifications to the Service Now ticketing system to create incidents. The incident includes information like the Caller, Category, Subcategory, Business Service, and other attributes related to alerts.

#### Prerequisites

Ensure that you have log in credentials for Service-Now.

Ensure that you are assigned with IT Infrastructure Library (ITIL) role in Service Now.

#### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

2 Click **Add** and from the **Plug-In Type** drop-down menu, select **Service-Now Notification Plug-in**.

The dialog box expands to include your plug-in instance settings.

3 Enter an **Instance Name**.

4 Enter the Service Now URL.

```
https://dev22418.service-now.com/
```

5 Enter a value for the Connection Count.

The connection count represents the maximum number of open connections allowed per node in VMware Aria Operations.

6 Select the **Credential Type** from the list.

**Note** If your Endpoint URL does not need any authentication, then select **No Credential** from the **Credential Type** list.

7 Add or edit the **Credential** details. Click the plus icon to enter the details of the new credentials in the **Create New Credential** pane, and click **Save**. This field appears only when you select **Basic Authentication** as the **Credential Type**.

**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The VMware Aria Operations administrator must assign the credentials from the **Orphan and Unassigned** page. For more information, see the topic, Managing Orphaned and Unassigned Content.

8 To verify the specified paths, credentials, and permissions, click **Test**.

9 Click **Save**.

#### Results

This instance of the Service-Now Notifications plug-in is configured and running.

#### What to do next

When the plug-in is added, Creating Notification Rules for Alerts for creating incidents in Service-Now ticketing system.

## Notifications - Add a Slack Plugin for Outbound Notifications

You can add a Slack plug-in to forward alerts and configure multiple notification rules with different slack channels. The Slack plug-in allows you to receive pre-formatted alert details with alert fields and helps you run VMware Aria Operations using alert links to troubleshoot further.

#### Prerequisites

#### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

2 Click **Add** and from the **Plugin Type** drop-down menu, select **Slack Plugin**.

The dialog box expands to include your plug-in instance settings.

3 Enter an **Instance Name**.

4 Enter a value for the **Connection Count**.

The connection count represents the maximum number of open connections allowed per node in VMware Aria Operations.

5 To verify the specified paths, credentials, and permissions, click **Test**.

6 Click **Save**.

#### Results

This instance of the Slack plugin is configured and running.

#### What to do next

When the plugin is added, Creating Notification Rules for Alerts for different slack channels.

### Add a Webhook Notification Plugin for Outbound Instance

You can integrate Webhook with any endpoint REST API and configure outbound payload.

#### Prerequisites

Ensure that you have login credentials for Webhook.

#### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

2 Click **Add** and from the **Plugin Type** drop-down menu, select **Webhook Notification Plugin**.

The dialog box expands to include your plugin instance settings.

3 Enter an **Instance Name**.

4 Select the **Use Proxy** checkbox and then select the **Collector/Group** from the list.

**Note** When the credential type is **OAuth Authentication**, you can only select a single cloud proxy and not a collector group.

5 Enter the Webhook URL.

**Note** For notifications that reference webhook outbound instances, the URL that you enter here serves as a base URL that is combined with the endpoint URL suffix defined in related webhook payload templates.

6 Enter a value for the **Connection Count**.

The connection count represents the maximum number of open connections allowed per node in VMware Aria Operations.

7 Select the **Credential Type** from the list.

**Note** If your Endpoint URL does not need any authentication, then select **No Credential** from the **Credential Type** list.

8 Add or edit the **Credential** details. Click the plus icon and enter the details of the new credentials in the right pane, and click **Save**.

**Note** When you upgrade to the latest version of VMware Aria Operations, all credentials get unassigned. The VMware Aria Operations administrator must assign the credentials from the **Orphan and Unassigned** page. For more information, see the topic, Managing Orphaned and Unassigned Content.

The fields in the **Create New Credential** pane appear based on the **Credential Type** you select.

| Credential Type | Fields |
| --- | --- |
| Basic Authentication | Enter the Name, User Name, and Password. |
| Bearer Token | Enter the Name of the credential and the Token. |

| Credential Type | Fields |
| --- | --- |
| OAuth Authentication | Enter the following details:<br>■ Name: Enter a name for the authentication.<br>■ Grant Type: Select either **Client Credentials** or **Password Credentials**.<br>■ Authentication URL: Enter the URL from where the access token can be retrieved.<br>■ Client ID: Enter the client ID for the authentication URL.<br>■ Client Secret: Enter the client secret for the authentication URL.<br>■ User Name: Enter the user name for the authentication URL.<br>**Note** This field appears only when the grant type is **Password Credentials**.<br>■ Password: Enter the password for the authentication URL.<br>**Note** This field appears only when the grant type is **Password Credentials**.<br>■ Scope: Enter the labels to specify the access token. The labels specify the permissions that the access tokens will have.<br>■ Send Credentials: Select either **In auth header** or **In body**.<br>  ■ In auth header: Sends the Client ID and Client Secret in the header.<br>  ■ In body: Sends the Client ID and Client Secret in the payload body.<br>**Note** This field appears only when the grant type is **Client Credentials**.<br>■ Access Token Path: Enter your access token path.<br>■ Validity Token Path: To keep track of when the token is going to expire, enter the validity token path and select the format from the drop-down list. You can choose one of the following formats:<br>  ■ Second<br>  ■ Milisecond<br>  ■ Absolute Time<br>■ Header Name: Enter a header name. By default, the header name is 'Authorization'.<br>■ Prefix: Enter a prefix. By default, the prefix is 'Bearer'.<br>■ Use Proxy: Select this check box to activate the **Collector/Group** drop-down list.<br>■ Collector/Group: Select the cloud proxy from the list.<br>**Note** When the credential type is **OAuth Authentication**, you can only select a single cloud proxy and not a collector group. |
| Certificate | Enter the following details:<br>■ Name: Enter the name of the certificate.<br>■ Certificate: Enter the certificate in the X.509 format.<br>■ Certificate Key: Enter the private key. The formats supported are Open SSL, PKCS1, and PKCS8 |
| API Key | Enter the Name, API Key, and the API Value. |

9 To verify the specified paths, credentials, and permissions, click **Test**.

**Note**

- The **Test** feature does not currently support the Custom Templates. If the template contains a custom headers, the test might fail.

10 Click **Save**.

Results

This instance of the Webhook Notification plugin is configured and running.

## Sample Email Alert

Here is a sample email for a newly created alert.

```
Alert Definition Name: Node is experiencing swapping due to memory pressure
Alert Definition Description: Node is experiencing swapping due to memory pressure
Object Name : VMware Aria Operations Node-VMwareAria Cluster Node
Object Type : vC-Ops-Node
Alert Impact: risk
Alert State : warning
Alert Type : Application
Alert Sub-Type : Performance
Object Health State: info
Object Risk State: warning
Object Efficiency State: info
Control State: Open
Symptoms:
SYMPTOM SET - self
```

| Symptom Name | Object Name | Object ID | Metric | Message Info |
|---|---|---|---|---|
| Node swap usage at Warning level | VMware Aria Operations Node-VMware Aria Cluster Node | 50ec874a-2d7d-4e78-98b1-afb26fd67e58 | Swap\|Workload | 59.183 > 30.0 |

```
Recommendations:
Notification Rule Name: rule1
Notification Rule Description:
Alert ID : badc2266-935d-4fb9-8594-e2e71e4866fc
VCOps Server - VMwareAriaClusterNode

Alert details(link)
```

## VMware Hosted Email plug-in for Outbound Settings

Use the VMware hosted email plug-in for alert notifications and reports delivery in VMware Aria Operations.

Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Outbound Settings**.

   The Outbound Settings page appears.

2 The VMware hosted email plug-in is enabled by default. Click the vertical ellipsis and click **Edit**.

3 Enter the **Receiver's Email Address** and click **Send** to test the connection. The **Send** button appears after you enter the email address.

   **Note** You can receive a maximum of 24 thousand emails in a day.

4 Click **Save**.

## Exporting and Importing Outbound Settings

As a VMware Aria Operations admin, you can backup the content before upgrading, export all the outbound plugin configurations, and import it into a different VMware Aria Operations instance. You can also export the content from VMware Aria Operations on-prem to VMware Aria Operations.

**Note** Any user with "Manage" Outbound Settings permission can export and import outbound plugin configurations.

Procedure

1 Export an outbound setting.

   a From the left menu, click **Configure > Alerts**, and in the right pane, click **Outbound Settings**.

   b Select the outbound settings that you want to export and click the horizontal ellipses and select **Export**.

   c Setup a new password to export data. The password should be at least 14 characters long.

   d Click **Export**.

      The outbound setting data is exported in the .json format. A password is used to encrypt the data in the file using the AES algorithm with 128 bit key. Use the same password while importing this file.

2 Import an outbound setting.

**Note** Before importing the outbound setting, ensure that you have exported the outbound plugin configurations.

a From the left menu, click **Configure > Alerts**, and in the right pane, click **Outbound Settings**.

b Click the horizontal ellipses and select **Import**.

c Click **Browse** to select the `.json` file and enter the password that you had set while exporting the content.

d If there is a conflict while importing the content, you can either overwrite the existing outbound settings or skip the import, which is the default.

e Click **Import** to import outbound settings to the destination setup.

**Note** While importing outbound settings on VMware Aria Operations, HTTP Proxy configurations will be excluded.

# Payload Templates in VMware Aria Operations

Payload is an essential information in the data block that you send or receive from the server. Use the **Payload Templates** page to view the list of payload templates available for each plug-in in VMware Aria Operations. You can add, manage, and edit your payload templates from this page. Default payload templates are provided for each plug-in type.

## Where You Find Payload Templates

To manage your payload templates, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Payload Templates**.

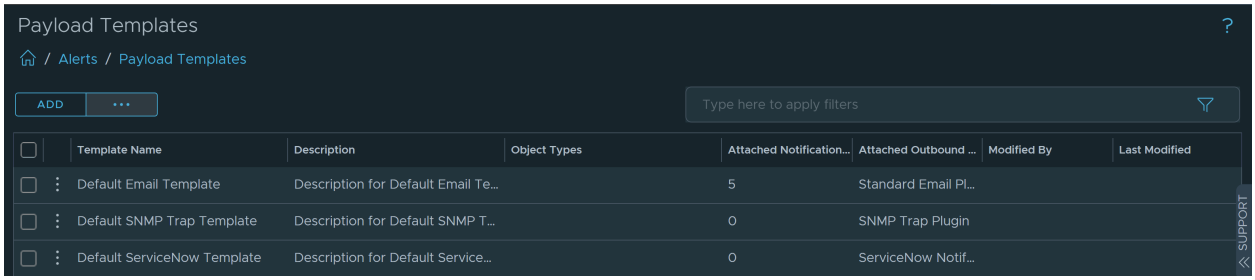

| Option | Description |
| --- | --- |
| Toolbar options | Use the toolbar options to manage your notification rules.<br>■ Add. Use the Create Payload Template dialog box to create new payload templates.<br>Click the horizontal ellipsis to perform the following actions.<br>■ Delete. Removes the selected payload template.<br>■ Export. Downloads the payload template.<br>**Note** Export and delete actions are not supported for the default payload templates available for each plug-in.<br>■ Import. Allows you to import payload templates. To import:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select the file to import.<br>■ Select if you want to Overwrite or Skip the import in case of a conflict.<br>■ Click **Import** to import the payload template, and then click **Done**. A message with the number of imported and skipped files will appear. |
| Quick Filter | Limits the list based on the text you type. It considers only text from templates Name column.<br>You can also sort columns in the data grid. You can view a blue arrow next to the column according to which sorting was performed, pointing up or down based on the sorting order (ascending or descending). |
| Template Name | Name of the payload template.<br>Click the vertical ellipsis to perform the following actions.<br>■ Edit. Allows you to edit the selected payload template.<br>■ Clone. Clones the selected payload template.<br>■ Delete. Removes the selected payload template.<br>■ Export. Downloads the payload template.<br>■<br>**Note** Edit, Delete, and Export actions are not supported for the default payload templates available for each plug-in. |
| Description | Description of the payload template. |
| Object Types | Base object type against which the payload template is defined, if any. |
| Attached Notification Rules | Notification rule attached to the payload template. |
| Attached Outbound Methods | Outbound plugin type attached to the payload template. |
| Modified By | Name of the last person to modify the payload template. |
| Last Modified | Date on which the payload template was last modified. |

## Creating Payload Templates for Outbound Plugins

You can create a payload template for an outbound plugin of your choice in VMware Aria Operations.

Use payload templates to configure the payload of an email, and customize the subject line, and email body. You can enable your own input properties and different payloads for updated and canceled alerts.

You can customize payloads only for Standard Email Plug-in and Webhook Notification Plug-in.

### Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Payload Templates**. On the toolbar, click **Add** to create a new payload template.

2 In the **Details** tab, enter the basic details of the payload template.

| Option | Description |
| --- | --- |
| Name | Provide a name for the payload template. |
| Description | Enter a description for the payload template. |
| Outbound Method | Outbound plugin for which you want to create a new payload template.<br>Select Plug-In Type: Type of plugin. Select one of the outbound alert plug-in types: Standard Email Plugin, Webhook Notification Plugin, Service-Now Notification Plugin, and Slack Plugin. |

3 Click **Next**.

4 In the **Object Content** tab, define the object details that you want to include in the notifications.

**Note** To create a payload template, it is mandatory to add the object type for all the outbound plugin types except for the Standard Email Plugin and Webhook Notification Plugin.

| Option | Description |
| --- | --- |
| Add Object Type | Select an object type from the list. Once you select the object type, define the metrics, properties, parents, and ancestors associated with the object type that you want to include in the notification.<br>For example, if you select Datacenter as an object type, click **Edit** to define the metrics, properties, parents, and ancestors associated with it.<br>From the right pane, double-click or drag the metrics and properties into the **Metrics and Properties** box. You can select up to 30 metrics and properties.<br>Define the parent host or cluster information that you want to include in the notification. From the right pane, double-click or drag the parent and ancestors information into the **Parents and Ancestors** box.<br>The information that you define here will be included in the alert notification for all the plug-ins. However, for a Webhook Notification Plugin and Standard Email Plugin, the information will be included only when you define the values in the **Payload Details** tab. |

5 Click **Create** to create the new payload template or click **Next** if you are creating a payload template for a Standard Email Plugin or a Webhook Notification Plugin.

6 In the **Payload Details** tab, enter the payload details that you want to include in the notification.

**Note** This tab is available only when you are creating a payload template for a Standard Email Plugin or a Webhook Notification Plugin.

| Option | Description |
| --- | --- |
| Do you want to add template input properties? | Select **Yes** to add input properties and enter the **Key**, **Type**, **Display Name**, and **Description** of the input property. Otherwise, select **No**.<br>**Note** The input properties are specific to your endpoint. Once you define the input properties in the template, you must provide the appropriate values in each rule where this template will be used. |
| Do you want different payload details for new, updated, and canceled alerts? | Select **Yes** to define different payload details for new, updated, and canceled alerts. Otherwise, select **No**. |
| The following fields appear while creating a payload template for the Standard Email Plugin. | |
| Subject | Enter a subject for the email notification. |

| Option | Description |
| --- | --- |
| Body | Enter the content for the email notification. You can also search for parameters in the right pane. Click the copy icon next to the parameter to copy the parameter and you can paste the parameter in the email body.<br>You can use the options in the toolbar to edit, format, and highlight the email content.<br>**Note** You can set up different email content for new, updated, and canceled alerts. |
| The following fields appear while creating a payload template for the Webhook Notification Plugin. | |
| Endpoint URL | Enter the endpoint URL. The endpoint URL will be appended to the base URL provided in the related webhook outbound instance.<br>**Note** The entire URL is encoded. However, there is an exception to use the character '/' in the URL. |
| Content Type | Select the content type for the payload. |
| Custom Headers | Enter the **HTTP Custom Header Name** and **Value**. Click the plus icon to add multiple custom headers.<br>**Note** For webhook payloads using token-based authentication, add an Authorization Header in the format required by the endpoint. |
| HTTP Method | Select the HTTP method of request. |
| Payload of the request | Payload for the selected plug-in type. It displays information based on the selected metrics, properties, ancestors, and object types.<br>You can search for parameters in the right pane. Click the copy icon next to the parameter to copy the parameter and you can paste the parameter in the **Payload of the request** box. |

7 Click **Create**.

Once the payload template is created, you can view it in the **Payload Templates** page. After selecting a payload template in the notification rule, you can view the payload template details in the **Notifications** page.

# Managing Alert Groups in VMware Aria Operations

For easy and better management of alerts, you can arrange them as a group as per your requirement.

It is complicated to identify a problem in large environments as you receive different kind of alerts. To manage alerts easily, group them by their definitions.

For example, there are 1000 alerts in your system. To identify different types of alerts, group them based on their alert definitions. It is also easy to detect the alert having the highest severity in the group.

When you group alerts, you can view the number of times the alerts with the same alert definition are triggered. By grouping alerts, you can perform the following tasks easily and quickly:

- Find the noisiest alert: The alert that has triggered maximum number of times is known as the noisiest alert. Once you find it, you can deactivate it to avoid further noise.
- Filter alerts: You can filter alerts based on a substring in alert definitions. The result shows the group of alerts that contain the substring.

**Note**

- If you cancel or deactivate an alert group, the alerts are not canceled instantly. It might take some time if the group is large.
- Only one group can be expanded at a time.
- The number next to the group denotes the number of alerts in that particular group.
- The criticality sign indicates the highest level of severity of an alert in a group.

## Grouping Alerts in VMware Aria Operations

You can group alerts by time, criticality, definition, and object type.

To group alerts:

### Procedure

1 From the left menu, click **Troubleshoot** > **Alerts**.

2 Select from the various options available from the **Group By** drop-down menu.

## Deactivating Alerts in VMware Aria Operations

In an alerts group, you can deactivate an alert by a single click.

To deactivate an alert:

1 From the left menu, click **Troubleshoot** > **Alerts**.

2 From the **Group By** drop-down, select **Definition**, and click on the name of the Alert Definition Group.

3 From the data grid, click **Actions** > **Deactivate**.

   You can deactivate the alerts by two methods:

   - Deactivate Alert in All Policies: Deactivates the alert for all the objects for all the policies.

- Deactivate Alert in Selected Policies: Deactivates the alert for the objects having the selected policy.

# Alert Definition Best Practices

As you create alert definitions for your environment, apply consistent best practices so that you optimize alert behavior for your monitored objects.

## Alert Definitions Naming and Description

The alert definition name is the short name that appears in the following places:

- In data grids when alerts are generated
- In outbound alert notifications, including the email notifications that are sent when outbound alerts and notifications are configured in your environment

Ensure that you provide an informative name that clearly states the reported problem. Your users can evaluate alerts based on the alert definition name.

The alert definition description is the text that appears in the alert definition details and the outbound alerts. Ensure that you provide a useful description that helps your users understand the problem that generated the alert.

## Wait and Cancel Cycle

The wait cycle setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition goes into effect after the wait cycle for the symptom definition results in a triggered symptom. In most alert definitions you configure the sensitivity at the symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that the alert is immediately generated after all of the symptoms are triggered at the desired symptom sensitivity level.

The cancel cycle setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition goes into affect after the cancel cycle for the symptom definition results in a cancelled symptom. In most definitions you configure the sensitivity at the symptom level and configure the cancel cycle of alert definition to 1. This configuration ensures that the alert is immediately cancelled after all of the symptoms conditions disappear after the desired symptom cancel cycle.

## Create Alert Definitions to Generate the Fewest Alerts

You can control the size of your alert list and make it easier to manage. When an alert is about a general problem that can be triggered on a large number of objects, configure its definition so that the alert is generated on a higher level object in the hierarchy rather than on individual objects.

As you add symptoms to your alert definition, do not overcrowd a single alert definition with secondary symptoms. Keep the combination of symptoms as simple and straightforward as possible.

You can also use a series of symptom definitions to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of Warning while `Volume reached capacity limit` might have a severity level of Critical. The first symptom is not an immediate threat, but the second one is an immediate threat. You can then include the Warning and Critical symptom definitions in a single alert definition with an Any condition and set the alert criticality to be Symptom Based. These settings cause the alert to be generated with the right criticality if either of the symptoms is triggered.

## Avoid Overlapping and Gaps Between Alerts

Overlaps result in two or more alerts being generated for the same underlying condition. Gaps occur when an unresolved alert with lower severity is canceled, but a related alert with a higher severity cannot be triggered.

A gap occurs in a situation where the value is <=50% in one alert definition and >=75% in a second alert definition. The gap occurs because when the percentage of volumes with high use falls between 50 percent and 75 percent, the first problem cancels but the second does not generate an alert. This situation is problematic because no alert definitions are active to cover the gap.

## Actionable Recommendations

If you provide text instructions to your users that help them resolve a problem identified by an alert definition, precisely describe how the engineer or administrator should fix the problem to resolve the alert.

To support the instructions, add a link to a wiki, runbook, or other sources of information, and add actions that you run from VMware Aria Operations on the target systems.

# Create a Simple Alert Definition

While troubleshooting, you can now quickly create an alert for a particular object type or a metric in a quick and efficient way.

You can create a simple alert definition from the following locations.

- From the left menu, click **Troubleshoot > Workbench** and select the metric for which you want to create an alert. You can create an alert from the **Potential Evidence** or the **Metrics** tab.
- From the left menu, click **Troubleshoot > Alerts**. Select an alert and click the **Potential Evidence** tab.
- From the Home page, you can search for the specific object type or metric in the search bar.

Procedure

1 Click the drop-down menu available in the right side of the widget and select the **Create an Alert Definition** option.

**Note** You cannot create alert definitions for Badge, Time, and Capacity Remaining metrics.

2 In the Create Alert Definition page, enter the **Name** and **Description** of the alert.

3 Set thresholds, criticality, and the number of wait cycles. Click **Show Advanced Settings** to set Wait Cycle and Cancel Cycle.

**Note** The Object Type or Metric/Property are pre-selected and cannot be edited.

4 Click **Create**.

The new alert is created and the policy the object belongs to and its children policies are activated for the alert.

# Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for VMware Aria Operations to alert you. When the alert is triggered on your host system, VMware Aria Operations alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have VMware Aria Operations notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Alert Definitions**.

2 Enter **`capacity`** in the search text box.

Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.

3 Click **Add** to create a new capacity alert definition for your host systems.

a In the alert definition workspace, for the Name and Description, enter **`Hosts - Alert on Capacity Exceeded`**.

b For the Base Object Type, select **vCenter Adapter > Host System**

c Under **Advanced Settings**, select the following options.

| Option | Selection |
|---|---|
| Impact | Select **Risk**. |
| Criticality | Select **Immediate**. |
| Alert Type and Subtype | Select **Application : Capacity**. |
| Wait Cycle | Select **1**. |
| Cancel Cycle | Select **1**. |

d In the **Symptoms/Conditions** workspace, select the following options.

| Option | Selection |
|---|---|
| Defined On | Select **Self**. |
| Symptom Definition Type | Select **Metric / Property**. |
| Quick filter (Name) | Enter `capacity`. |

e From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the left pane.

In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.

f For Add Recommendations, enter `virtual machine` in the quick filter text box.

g Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the left pane.

This recommendation is set to Priority 1.

4 Click **Save** to save the alert definition.

Your new alert appears in the list of alert definitions.

### Results

You have added an alert definition to have VMware Aria Operations alert you when the capacity of your host systems begins to run out.

## Create an Alert Definition for Department Objects

As a virtual infrastructure administrator, you are responsible for the virtual machines and hosts that the accounting department uses. You can create alerts to manage the accounting department objects.

You received several complaints from your users about delays when they are using their accounting applications. Using VMware Aria Operations, you identified the problem as related to CPU allocations and workloads. To better manage the problem, you create an alert definition with tighter symptom parameters so that you can track the alerts and identify problems before your users encounter further problems.

Using this scenario, you create a monitoring system that monitors your accounting objects and provides timely notifications when problems occur.

## Add Description and Base Object to Alert Definition

To create an alert to monitor the CPUs for the accounting department virtual machines and monitor host memory for the hosts on which they operate, you begin by describing the alert.

When you name the alert definition and define alert impact information, you specify how the information about the alert appears in VMware Aria Operations. The base object is the object around which the alert definition is created. The symptoms can be for the base object and for related objects.

Procedure

1 From the left menu, click **Configure > Alerts**, and then in the right pane, click **Alert Definitions**.

2 Click **Add** to add a definition.

3 Type a name and description.

   In this scenario, type **`Acct VM CPU early warning`** as the alert name, which is a quick overview of the problem. The description, which is a detailed overview. should provide information that is as useful as possible. When the alert is generated, this name and description appears in the alert list and in the notification.

4 From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Host System**.

   This alert is based on host systems because you want an alert that acts as an early warning to possible CPU stress on the virtual machines used in the accounting department. By using host systems as the based object type, you can respond to the alert symptom for the virtual machines with bulk actions rather than responding to an alert for each virtual machine.

5 Click **Advanced Settings** and configure the metadata for this alert definition.

   a From the **Impact** drop-down menu, select **Risk**.

      This alert indicates a potential problem and requires attention in the near future.

   b From the **Criticality** drop-down menu, select **Immediate**.

      As a Risk alert, which is indicative of a future problem, you still want to give it a high criticality so that it is ranked for correct processing. Because it is designed as an early warning, this configuration provides a built-in buffer that makes it an immediate risk rather than a critical risk.

c From the **Alert Type and Subtype** drop-down menu, select **Performance** under **Virtualization/Hypervisor**.

d To ensure that the alert is generated during the first collection cycle after the symptoms become true, set the **Wait Cycle** to **1**.

e To ensure that the an alert is removed as soon as the symptoms are no longer triggered, set the **Cancel Cycle** to **1**.

The alert is canceled in the next collection cycle if the symptoms are no long true.

These alert impact options help you identify and prioritize alerts as they are generated.

#### Results

You started an alert definition where you provided the name and description, selected host system as the base object type, and defined the data that appears when the alert generated.

#### What to do next

Continue in the workspace, adding symptoms to your alert definition. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

## Add a Virtual Machine CPU Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add symptoms to your VMware Aria Operations alert definition after you provide the basic descriptive information for the alert. The first symptom you add is related to CPU usage on virtual machines. You later use a policy and group to apply alert to the accounting virtual machines.

This scenario has two symptoms, one for the accounting virtual machines and one to monitor the hosts on which the virtual machines operate.

#### Prerequisites

Begin configuring the alert definition. See Add Description and Base Object to Alert Definition.

#### Procedure

1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next** and configure the symptoms.

2 Begin configuring the symptom set related to virtual machines CPU usage.

a From the **Select Symptom** drop-down menu, select **Metric / Property**.

b From the **Defined On** drop-down menu, select **Child**.

c From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.

d Click **Create New** to open the **Add Symptom Definition** workspace window.

3 Configure the virtual machine CPU usage symptom in the **Add Symptom Definition** workspace window.

   a From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Virtual Machine**.

      The collected metrics for virtual machines appears in the list.

   b In the metrics list **Search** text box, which searches the metric names, type **`usage`**.

   c In the list, expand **CPU** and drag **Usage (%)** to the workspace on the left.

   d From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use VMware Aria Operations analytics to identify the trend metric values for objects.

   e In the **Symptom Definition Name** text box, type a name similar to **`VM CPU Usage above trend`**.

   f From the criticality drop-down menu, select **Warning**.

   g From the threshold drop-down menu, select **Above Threshold**.

   h Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for 3 collection cycles before the symptom is triggered. This wait avoids triggering the symptom when there is a short spike in CPU usage.

   i Click **Save**.

   The dynamic symptom, which identifies when the usage is above the tracked trend, is added to the symptom list.

4 In the **Alert Definition Workspace** window, drag **VM CPU Usage above trend** from the symptom definition list to the symptom workspace on the left.

   The Child-Virtual Machine symptom set is added to the symptom workspace.

5 In the symptoms set, configure the triggering condition so that when the symptom is true on half of the virtual machines in the group to which this alert definition is applied, the symptom set is true.

   a From the value operator drop-down menu, select **>**.

   b In the value text box, enter **`50`**.

   c From the value type drop-down menu, select **Percent**.

### Results

You defined the first symptom set for the alert definition.

### What to do next

Add the host memory usage symptom to the alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

## Add a Host Memory Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add a second symptom to your VMware Aria Operations alert definition after you add the first symptom. The second symptom is related to host memory usage for the hosts on which the accounting virtual machines operate.

### Prerequisites

Add the virtual machine CPU usage symptom. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

### Procedure

1. In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next**.
2. Configure the symptom related to host systems for the virtual machines.
   a. From the **Select Symptom** drop-down menu, select **Metric / Property**.
   b. From the **Defined On** drop-down menu, select **Self**.
   c. Click **Create New** to add new symptom.
3. Configure the host system symptom in the **Add Symptom Definition** workspace window.
   a. From the **Base Object Type** drop-down menu, expand **vCenter Adapters** and select **Host System**.
   b. In the metrics list, expand **Memory** and drag **Usage (%)** to the workspace on the left.
   c. From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use VMware Aria Operations analytics to identify the trend metric values for objects.
   d. In the **Symptom Definition Name** text box, enter a name similar to **`Host memory usage above trend`**.
   e. From the criticality drop-down menu, select **Warning**.
   f. From the threshold drop-down menu, select **Above Threshold**.
   g. Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for three collection cycles before the symptom is triggered. This wait avoids triggering the symptom when a short spike occurs in host memory usage.
   h. Click **Save**.

The dynamic symptom identifies when the hosts on which the accounting virtual machines run are operating above the tracked trend for memory usage.

The dynamic symptom is added to the symptom list.

4 In the **Alert Definition Workspace** window, drag **Host memory usage above trend** from the symptoms list to the symptom workspace on the left.

The Self-Host System symptom set is added to the symptom workspace.

5 On the Self-Host System symptom set, from the value type drop-down menu for **This Symptom set is true when**, select **Any**.

With this configuration, when any of the hosts running accounting virtual machines exhibit memory usage that is above the analyzed trend, the symptom condition is true.

6 At the top of the symptom set list, from the **Match {operator} of the following symptoms** drop-down menu, select **Any**.

With this configuration, if either of the two symptom sets, virtual machine CPU usage or the host memory, are triggered, an alert is generated for the host.

### Results

You defined the second symptom set for the alert definition and configured how the two symptom sets are evaluated to determine when the alert is generated.

### What to do next

Add recommendations to your alert definition so that you and your engineers know how to resolve the alert when it is generated. See Add Recommendations to the Alert Definition.

## Add Recommendations to the Alert Definition

To resolve a generated alert for the accounting department's virtual machines, you provide recommendations so that you or other engineers have the information you need to resolve the alert before your users encounter performance problems.

As part of the alert definition, you add recommendations that include actions that you run from VMware Aria Operations and instructions for making changes in vCenter Server that resolve the generated alert.

### Prerequisites

Add symptoms to your alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

### Procedure

1 In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, **Alert Impact**, and **Add Symptom Definitions**, click **Next** and add the recommended actions and instructions.

2 Click **Create New Recommendation** and select an action recommendation to resolve the virtual machine alerts.

   a In the **Description** text box, enter a description of the action similar to **`Add CPUs to virtual machines`**.

   b From the **Actions** drop-down menu, select **Set CPU Count for VM**.

   c Click **Create**.

3 Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory problems similar to this example.

   **`If this host is part of a DRS cluster, check the DRS settings to verify that the load balancing setting are configured correctly. If necessary, manually vMotion the virtual machines.`**

4 Click **Create**.

5 Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory alerts.

   a Enter a description of the recommendation similar to this example.

   **`If this is a standalone host, add more memory to the host.`**

   b To make the URL a hyperlink in the instructions, copy the URL, for example, https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html, to your clipboard.

   c Highlight the text in the text box and click the hyperlink icon.

   d Paste the URL in the **Create a hyperlink** text box and click **OK**.

   e Click **Create**.

6 In the **Alert Recommendation Workspace**, drag **Add CPUs to virtual machines**, **If this host is part of a DRS cluster**, and the **If this is a standalone host** recommendations from the list to the recommendation workspace in the order presented.

7 Click **Next** to select policies and view notifications.

8 Click **Create**.

### Results

You provided the recommended actions and instructions to resolve the alert when it is generated. One of the recommendations resolves the virtual machine CPU usage problem and the other resolves the host memory problem.

### What to do next

Create a group of objects to use to manage your accounting objects. See Create a Custom Accounting Department Group.

## Create a Custom Accounting Department Group

To manage, monitor, and apply policies to the accounting objects as a group, you create a custom object group.

### Prerequisites

Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.

### Procedure

1. From the left menu, click **Environment** > **Custom Groups**.
2. Click **Add** to create a new custom group.
3. Type a name similar to **`Accounting VMs and Hosts`**.
4. From the **Group Type** drop-down menu, select **Department**.
5. From the **Policy** drop-down menu, select **Default Policy**.

   When you create a policy, you apply the new policy to the accounting group.
6. In the Define membership criteria area, from the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Host System**, and configure the dynamic group criteria.
   a. From the criteria drop-down menu, select **Relationship**.
   b. From the relationships options drop-down menu, select **Parent of**.
   c. From the operator drop-down menu, select **contains**.
   d. In the **Object name** text box, enter **`acct`**.
   e. From the navigation tree drop-down list, select **vSphere Hosts and Clusters**.

   You created a dynamic group where host objects that are the host for virtual machines with acct in the virtual machine name are included in the group. If a virtual machine with acct in the object name is added or moved to a host, the host object is added to the group.
7. Click **Preview** in the lower-left corner of the workspace, and verify that the hosts on which your virtual machines that include acct in the object name appear in the **Preview Group** window.
8. Click **Close**.
9. Click **Add another criteria set**.

   A new criteria set is added with the OR operator between the two criteria sets.

10 From the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Virtual Machine**, and configure the dynamic group criteria.

   a From the criteria drop-down menu, select **Properties**.

   b From the **Pick a property** drop-down menu, expand **Configuration** and double-click **Name**.

   c From the operator drop-down menu, select **contains**.

   d In the **Property value** text box, enter **`acct`**.

   You created a dynamic group where virtual machine objects with acct in the object name are included in the group that depends on the presence of those virtual machines. If a virtual machine with acct in the name is added to your environment, it is added to the group.

11 Click **Preview** in the lower-left corner of the workspace, and verify that the virtual machines with acct in the object name are added to the list that also includes the host systems.

12 Click **Close**.

13 Click **OK**.

   The Accounting VMs and Hosts group is added to the Groups list.

#### Results

You created a dynamic object group that changes as virtual machines with acct in their names are added, removed, and moved in your environment.

#### What to do next

Create a policy that determines how VMware Aria Operations uses the alert definition to monitor your environment. See Create a Policy for the Accounting Alert.

## Create a Policy for the Accounting Alert

To configure how VMware Aria Operations evaluates the accounting alert definition in your environment, you configure a policy that determines behavior so that you can apply the policy to an object group. The policy limits the application of the alert definition to only the members of the selected object group.

When an alert definition is created, it is added to the default policy and activated, ensuring that any alert definitions that you create are active in your environment. This alert definition is intended to meet the needs of the accounting department, so you deactivate it in the default policy and create a new policy to govern how the alert definition is evaluated in your environment, including which accounting virtual machines and related hosts to monitor.

#### Prerequisites

- Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.

- Verify that you created a group of objects that you use to manage you accounting objects. See Create a Custom Accounting Department Group.

Procedure

1 From the left menu, click **Configure** > **Policies**.

2 Click the **Policy Library** tab and then, click **Add** .

3 Type a name similar to **`Accounting Objects Alerts Policy`** and provide a useful description similar to the following example.

```
This policy is configured to generate alerts when
Accounting VMs and Hosts group objects are above trended
 CPU or memory usage.
```

4 Select **Default Policy** from the **Start with** drop-down menu.

5 On the left, click **Customize Alert / Symptom Definitions** and deactivate all the alert definitions except the new Acct VM CPU early warning alert.

   a In the Alert Definitions area, click **Actions** and select **Select All**.

      The alerts on the current page are selected.

   b Click **Actions** and select **Deactivate**.

      The alerts indicate Deactivated in the State column.

   c Repeat the process on each page of the alerts list.

   d Select **Acct VM CPU early warning** in the list, click **Actions** and select **Activate**.

      The Acct VM CPU early warning alert is now activated.

6 On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.

7 Click **Save**.

Results

You created a policy where the accounting alert definition exists in a custom policy that is applied only to the virtual machines and hosts for the accounting department.

What to do next

Create an email notification so that you learn about alerts even you when you are not actively monitoring VMware Aria Operations. See Configure Notifications for the Department Alert.

## Configure Notifications for the Department Alert

To receive an email notification when the accounting alert is generated, rather than relying on your ability to generally monitor the accounting department objects in VMware Aria Operations, you create notification rules.

Creating an email notification when accounting alerts are triggered is an optional process, but it provides you with the alert even when you are not currently working in VMware Aria Operations.

### Prerequisites

- Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.
- Verify that standard email outbound alerts are configured in your system. See Add a Standard Email Plug-In for VMware Aria Operations Outbound Alerts.

### Procedure

1. From the left menu, click **Configure > Alerts**, and then in the right pane, click **Notifications**.
2. Click **Add** to add a notification rule.
3. Configure the communication options.
   a. In the **Name** text box, type a name similar to **`Acct Dept VMs or Hosts Alerts`**.
   b. From the **Select Plug-In Type** drop-down menu, select **StandardEmailPlugin**.
   c. From the **Select Instance** drop-down menu, select the standard email instance that is configured to send messages.
   d. In the **Recipient(s)** text box, type your email address and the addresses of other recipients responsible for the accounting department alerts. Use a semicolon between recipients.
   e. Leave the **Notify again** text box blank.

      If you do not provide a value, the email notice is sent only once. This alert is a Risk alert and is intended as an early warning rather than requiring an immediate response.

   You configured the name of the notification when it is sent to you and the method that is used to send the message.
4. In the Filtering Criteria area, configure the accounting alert notification trigger.
   a. From the **Notification Trigger** drop-down menu, select **Alert Definition**.
   b. Click **Select Alert Definitions**.
   c. Select **Acct VM CPU early warning** and click **Select**.
5. Click **Save**.

### Results

You created a notification rule that sends you and your designated engineers an email message when this alert is generated for your accounting department alert definition.

### What to do next

Create a dashboard with alert-related widgets so that you can monitor alerts for the accounting object group. See Create a Dashboard to Monitor Department Objects.

## Create a Dashboard to Monitor Department Objects

To monitor all the alerts related to the accounting department object group, you create a dashboard that includes the alert list and other widgets. The dashboard provides the alert data in a single location for all related objects.

Creating a dashboard to monitor the accounting virtual machines and related hosts is an optional process, but it provides you with a focused view of the accounting object group alerts and objects.

### Prerequisites

Create an object group for the accounting department virtual machines and related objects. See Create a Custom Accounting Department Group.

### Procedure

1. From the left menu, click **Visualize** > **Dashboards**, and then click **Create**.
2. In the Dashboard Configuration definition area, type a tab name similar to **`Accounting VMs and Hosts`** and configure the layout options.
3. Click **Widget List** and drag the following widgets to the workspace.
   - **Alert List**
   - **Efficiency**
   - **Health**
   - **Risk**
   - **Top Alerts**
   - **Alert Volume**

   The blank widgets are added to the workspace. To change the order in which they appear, you can drag them to a different location in the workspace.
4. On the Alert List widget title bar, click **Edit Widget** and configure the settings.
   a. In the **Title** text box, change the title to **`Acct Dept Alert List`**.
   b. For the **Refresh Content** option, select **On**.
   c. Type **`Accounting`** in the **Search** text box and click **Search**.

      The Accounting value corresponds to the name of the object group for the accounting department virtual machines and related hosts.
   d. In the filtered resource list, select the **Accounting VMs and Hosts** group.

      The Accounting VMs and Hosts group is identified in the Selected Resource text box.
   e. Click **OK**.

The Acct Dept Alert List is now configured to display alerts for the Accounting VMs and Hosts group objects.

5 Click **Widget Interactions** and configure the following interactions.

   a For Acct Dept Alert List, leave the selected resources blank.

   b For Top Alerts, Health, Risk, Efficiency, and Alert Volume select **Acct Dept Alert List** from the **Selected Resources** drop-down menu.

   c Click **Apply Interactions**.

   With the widget interaction configured in this way, the select alert in the Acct Dept Alert List is the source for the data in the other widgets. When you select an alert in the alert list, the Health, Risk, and Efficiency widgets display alerts for that object, Top Alerts displays the topic issues affecting the health of the object, and Alert Volume displays an alert trend chart.

6 Click **Save**.

Results

You created a dashboard that displays the alerts related to the accounting virtual machines and hosts group, including the Risk alert you created.

# Actions in VMware Aria Operations

Actions are the ability to update objects or read data about objects in monitored systems, and are commonly provided in VMware Aria Operations as part of a solution. The actions added by solutions are available from the object Actions menu, list and view menus, including some dashboard widgets, and can be added to alert definition recommendations.

The possible actions include read actions and update actions.

The read actions retrieve data from the target objects.

The update actions modifies the target objects. For example, you can configure an alert definition to notify you when a virtual machine is experiencing memory issues. Add an action in the recommendations that runs the Set Memory for Virtual Machine action. This action increases the memory and resolves the likely cause of the alert.

To see or use the actions for your vCenter Server objects, you must activate actions in the vCenter Adapter for each monitored vCenter Server instance. Actions can only be viewed and accessed if you have the required permissions.

# Actions, Modified Objects, and Object Levels in VMware Aria Operations Actions

The list of actions includes the name of the action, the objects that each one modifies, and the object levels at which you can run the action. You use this information to ensure that you correctly apply the actions as alert recommendations and when the actions are available in the **Actions** menu.

## Actions and Modified Objects

VMware Aria Operations actions make changes to objects in your managed vCenter Server instances.

When you grant a user access to actions in VMware Aria Operations, that user can take the granted action on any object that VMware Aria Operations manages.

## Action Object Levels

The actions are available when you work with different object levels, but they modify only the specified object. If you are working at the cluster level and select **Power On VM**, all the virtual machines in the cluster for which you have access permission are available for you to run the action. If you are working at the virtual machine level, only the selected virtual machine is available.

| Action | Modified Object | Object Levels |
|---|---|---|
| Rebalance Container | Virtual Machines | ■ Data Center<br>■ Custom Data Center |
| Delete Idle VM | Virtual Machines | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Set DRS Automation | Cluster | ■ Clusters |
| Move VM | Virtual Machine | ■ Virtual Machines |
| Power Off VM | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Shut Down Guest OS for VM | Virtual Machine<br>VMware Tools must be installed and running on the target virtual machines to run this action. | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Power On VM | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Delete Powered Off VM | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |

| Action | Modified Object | Object Levels |
| --- | --- | --- |
| Set Memory for VM<br>and<br>Set Memory for VM Power Off Allowed | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Set Memory Resources for VM | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Set CPU Count for VM<br>and<br>Set CPU Count for VM Power Off Allowed | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Set CPU Resources for VM | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Set CPU Count and Memory for VM<br>and<br>Set CPU Count and Memory for VM Power Off Allowed | Virtual Machine | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Delete Unused Snapshots for VM | Snapshot | ■ Clusters<br>■ Host Systems<br>■ Virtual Machines |
| Delete Unused Snapshots for Datastore | Snapshot | ■ Clusters<br>■ Datastores<br>■ Host Systems |
| Execute Script | Virtual Machine | ■ Virtual Machine |
| Get Top Processes | Virtual Machine | ■ Virtual Machine |
| Apply Guest User Mapping | vCenter Server | ■ vCenter Server<br>**Note** This action is deprecated and will be removed in the next release. |
| Clear Guest User Mapping | vCenter Server | ■ vCenter Server<br>**Note** This action is deprecated and will be removed in the next release. |

| Action | Modified Object | Object Levels |
|---|---|---|
| Export Guest User Mapping | vCenter Server | ■ vCenter Server<br>**Note** This action is deprecated and will be removed in the next release. |
| Configure Included Services | Service Discovery Adapter Instance | ■ Service Discovery Adapter Instance<br>**Note** This action is deprecated and will be removed in the next release. |

# Viewing Actions List in VMware Aria Operations

Actions are the method you use to configuration changes on managed objects that you initiate from VMware Aria Operations. These actions are available to add to alert recommendations.

Actions are defined to run on the target object from different object levels, allowing you to add actions as recommendations for alert definitions that are configured for different base objects. The Actions page is a list of actions available in your environment.

## Where You Find the Actions List

To view the available actions, from the left menu, click **Configure > Alerts**, and then in the right pane, click **Actions**.

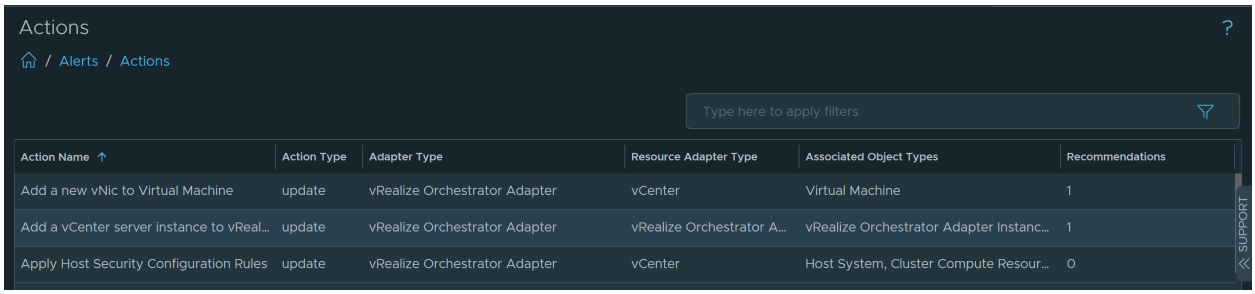


| Option | Description |
|---|---|
| Filter options | Limits the list based on the text you type.<br>You can also sort on the columns in the data grid. |
| Action Name | Name of the action. Duplicate names indicate that the action name is provided by more than one adapter or has more than one associated object.<br>Click this link to view Details page. On the Details Page, you can click on Recommendations to view the associated recommendations. |
| Action Type | Type of action that the action performs, either read or update.<br>■ Update actions make changes to the target objects.<br>■ Read actions retrieve data from the target objects. |
| Adapter Type | Name of the configured adapter that provides the action. |

| Option | Description |
| --- | --- |
| Resource Adapter Type | Adapter that provides the action. |
| Associated Object Types | Indicates the object level at which the action instance runs. |
| Recommendations | Indicates whether the action is used in at least one recommendation. |

The actions `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, these actions can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`
- `Set CPU Count for VM Power Off Allowed`
- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

## Actions Supported for Automation

Recommendations can identify ways to remediate problems indicated by an alert. Some of these remediations can be associated with actions defined in your VMware Aria Operations instance. You can automate several of these remediation actions for an alert when that recommendation is the first priority for that alert.

You activate actionable alerts in your policies. By default, automation is deactivated in policies. To configure automation for your policy, in the menu, click **Configure > Policies > Policy Library**. Then, to edit a policy, access the **Alert / Symptom Definitions** workspace, and select **Local** for the **Automate** setting in the Alert / Symptom Definitions pane.

When an action is automated, you can use the **Automated** and **Alert** columns in **Administration > Recent Tasks** to identify the automated action and view the results of the action.

- VMware Aria Operations uses the **automationAdmin** user account to trigger automated actions. For these automated actions that are triggered by alerts, the Submitted By column displays the **automationAdmin** user.
- The Alert column displays the alert that triggered the action. When an alert is triggered that is associated to the recommendation, it triggers the action without any user intervention.

The following actions are supported for automation:

- Delete Powered Off VM
- Delete Idle VM

- Move VM
- Power Off VM
- Power On VM
- Set CPU Count And Memory for VM
- Set CPU Count And Memory for VM Power Off Allowed
- Set CPU Count for VM
- Set CPU Count for VM Power Off Allowed
- Set CPU Resources for VM
- Set Memory for VM
- Set Memory for VM Power Off Allowed
- Set Memory Resources for VM
- Shut Down Guest OS for VM

## Roles Needed to Automate Actions

To automate actions, your role must have the following permissions:

- Create, edit, and import policies in **Configure > Policies > Policy Library**.
- Create, clone, edit, and import alert definitions in **Configure > Alerts > Alert Definitions**.
- Create, edit, and import recommendation definitions in **Configure > Alerts > Recommendations**.

**Important** You set the permissions used to run the actions separately from the alert and recommendation definition. Anyone who can modify alerts, recommendations, and policies can also automate the action, even if they do not have permission to run the action.

For example, if you do not have access to the Power Off VM action, but you can create and modify alerts and recommendations, you can see the Power Off VM action and assign it to an alert recommendation. Then, if you automate the action in your policy, VMware Aria Operations uses the `automationAdmin` user to run the action.

## Example Action Supported for Automation

For the Alert Definition named `Virtual machine has chronic high CPU workload leading to CPU stress`, you can automate the action named `Set CPU Count for VM`.

When CPU stress on your virtual machines exceeds a critical, immediate, or warning level, the alert triggers the recommended action without user intervention.

## Integration of Actions with VMware Aria Automation

VMware Aria Operations restricts actions on Datacentres and Custom Datacentres that contains VMware Aria Automation managed child objects such as, cluster compute resources, hosts, and VMs.

You can turn on or turn off the actions on VMware Aria Automation managed objects by modifying the **Operational Actions** from the respective vCentre in Cloud Accounts or by creating a new role with limited action ability on VMware Aria Automation managed objects.

### Actions Determine Whether Objects Are Managed

Actions check the objects in the VMware Aria Automation managed resource container to determine which objects are being managed by VMware Aria Automation.

Actions such as Rebalance Container check the child objects of the data center container or custom data center container to determine whether the objects are managed by VMware Aria Automation. If the objects are being managed, the action does not appear on those objects.

## Working with Actions That Use Power Off Allowed

Some of the actions provided with VMware Aria Operations require the virtual machines to shut down or power off, depending on the configuration of the target machines, to run the actions. You should understand the impact of the Power Off Allowed option before running the actions so that you select the best options for your target virtual machines.

### Power Off and Shut Down

The actions that you can run on your vCenter Server instances include actions that shut down virtual machines and actions that power off virtual machines. It also includes actions where the virtual machine must be in a powered off state to complete the action. Whether the VM is shut down or powered off depends on how it is configured and what options you select when you run the action.

The shut-down action shuts down the guest operating system and then powers off the virtual machine. To shut down a virtual machine from VMware Aria Operations, the VMware Tools must be installed and running on the target objects.

The power off action turns off the VM without regard for the state of the guest operating system. In this case, if the VM is running applications, your user might lose data. After the action is finished, for example, modifying the CPU count, the virtual machine is returned to the power state it was in when the action began.

### Power Off Allowed and VMware Tools

For the actions where you are increasing the CPU count or the amount of memory on a VM, some operating systems support the actions if the Hot Plug is configured on the VM. For other operating systems, the virtual machine must be in a powered off state to change the configuration. To accommodate this need where the VMware Tools is not running, the Set CPU Count, Set Memory, and Set CPU Count and Memory actions include the Power Off Allowed option.

If you select Power Off Allowed, and the machine is running, the action verifies whether VMware Tools is installed and running.

- If VMware Tools is installed and running, the virtual machine is shut down before completing the action.
- If VMware Tools is not running or not installed, the virtual machine is powered off without regard for the state of the operating system.

If you do not select Power Off Allowed and you are decreasing the CPU count or memory, or the hot plug is not activated for increasing the CPU count or memory, the action does not run and the failure is reported in Recent Tasks.

## Power Off Allowed When Changing CPU Count or Memory

When you run the actions that change the CPU count and the amount of memory, you must consider several factors to determine if you want to use the Power Off Allowed option. These factors include whether you are increasing or decreasing the CPU or memory and whether the target virtual machines are powered on. If you increase the CPU or memory values, whether hot plug is activated also affects how you apply the option when you run the action.

How you use Power Off Allowed when you are decreasing the CPU count or the amount of memory depends on the power state of the target virtual machines.

### Table 12-13. Decreasing CPU Count and Memory Behavior Based On Options

| Virtual Machine Power State | Power Off Allowed Selected | Results |
|---|---|---|
| On | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, decreases the CPU or memory, and powers the machine back on.<br>If VMware Tools is not installed, the action powers off the virtual machine, decreases the CPU or memory, and powers the machine back on. |
| On | No | The action does not run on the virtual machine. |
| Off | Not applicable. The virtual machine is powered off. | The action decreases the value and leaves the virtual machine in a powered off state. |

How you use Power Off Allowed when you are increasing the CPU count or the amount of memory depends on several factors, including the state of the target virtual machine and whether hot plug is activated. Use the following information to determine which scenario applies to your target objects.

If you are increasing the CPU count, you must consider the power state of the virtual machine and whether CPU Hot Plug is activated when determining whether to apply Power Off Allowed.

Table 12-14. Increasing CPU Count Behavior.

| Virtual Machine Power State | CPU Hot Plug Activated | Power Off Allowed Selected | Results |
|---|---|---|---|
| On | Yes | No | The action increases the CPU count to the specified amount. |
| On | No | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the CPU count, and powers the machine back on.<br>If VMware Tools is not installed, the action powers off the virtual machine, increases the CPU count, and powers the machine back on. |
| Off | Not applicable. The virtual machine is powered off. | Not required. | The action increases the CPU count to the specified amount. |

If you are increasing the memory, you must consider the power state of the virtual machine, whether Memory Hot Plug is activated, and whether there is a Hot Memory Limit when determining how to apply Power Off Allowed.

### Table 12-15. Increasing Memory Amount Behavior

| Virtual Machine Power State | Memory Hot Plug Activated | Hot Memory Limit | Power Off Allowed Selected | Results |
|---|---|---|---|---|
| On | Yes | New memory value ≤ hot memory limit | No | The action increases the memory the specified amount. |
| On | Yes | New memory value > hot memory limit | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.<br>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on. |
| On | No | Not applicable. The hot plug is not activated. | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.<br>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on. |
| Off | Not applicable. The virtual machine is powered off. | Not applicable. | Not required | The action increases the memory the specified amount. |

# Configuring Objects

13

Using the power of object management - including metrics and alerts - you can monitor objects, applications, and systems that must stay up and running. Some metrics and alerts are prepackaged into dashboards and policies; others you combine into custom tools

VMware Aria Operations discovers objects in your environment and makes them available to you. With the information that VMware Aria Operations provides, you can quickly access and configure any object. For example, you can determine if a datastore is connected or providing data, or you can power on a virtual machine.

Read the following topics next:


- Object Discovery


## Object Discovery

Its ability to monitor and collect data on objects in your systems environment makes VMware Aria Operations a critical tool in maintaining system uptime and ensuring ongoing good health for all system resources from virtual machines to applications to storage - across physical, virtual, and cloud infrastructures.

Following are examples of objects that can be monitored.

- vCenter Server
- Virtual machines
- Servers/hosts
- Compute resources
- Resource pools
- Data centers
- Storage components
- Switches
- Port groups
- Datastores

## Adapters – Key to Object Discovery

VMware Aria Operations collects data and metrics from objects using adapters, that are the central components of management packs. You can customize adapter instances for your virtual environment using cloud accounts and other accounts. VMware Aria Operations uses cloud accounts to manage the communication and integration with other products, applications, and functions.

- Cloud Accounts - You can configure cloud adapter instances and collect data from cloud solutions that are already installed in your cloud environment from the cloud accounts page.
- Other Accounts - You can view and configure native management packs and other solutions that are already installed and configure adapter instances from the other accounts page.
- Repository - You can activate or deactivate native management packs and add or upgrade other management packs from the Repository page.

The screenshot displays the list of available solutions in VMware Aria Operations . You must first Activate the solution before adding and configuring the accounts.

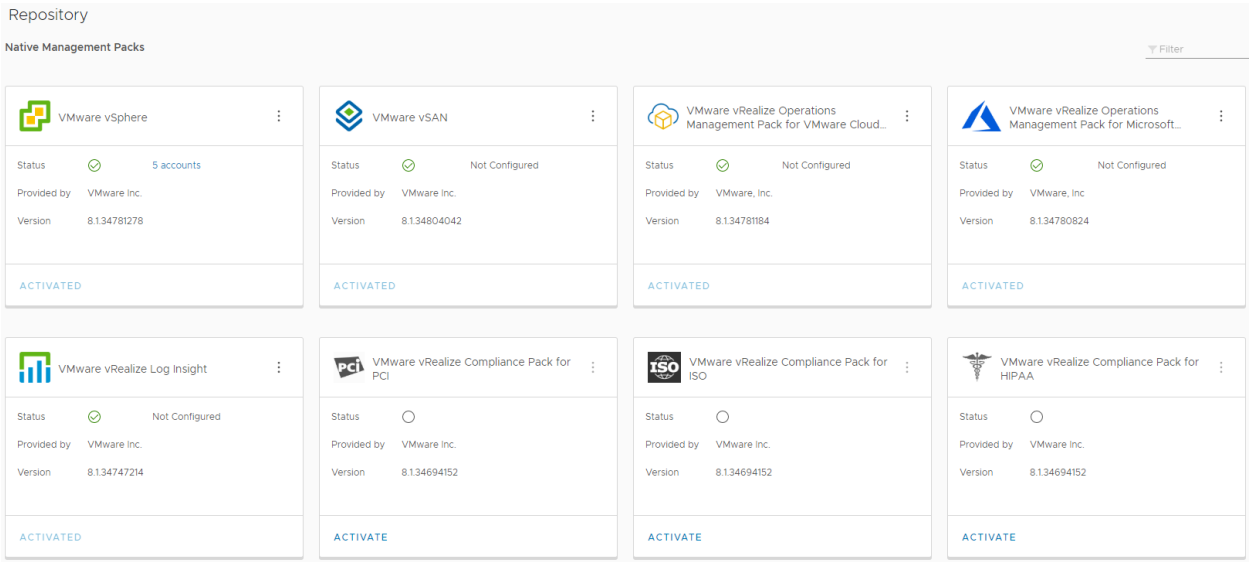


For complete information on configuring management packs and adapters, see Chapter 4 Integrating Data Sources with VMware Aria Operations

When you create a new adapter instance, it begins discovering and collecting data from the objects designated by the adapter, and notes the relationships between them. Now you can begin to manage your objects.

## Workload Management Inventory Objects

VMware Aria Operations discovers the following workload management objects and their child objects using the vCenter adapter:

- Tanzu Kubernetes cluster
- vSphere Pods

- Namespace

A cluster with Kubernetes activated, running on vSphere, is called a Supervisor Cluster. In the VMware Aria Operations inventory, the summary tab of the Supervisor Cluster indicates that it has workload management activated. The Supervisor Cluster contains specific objects that activate the capability to run Kubernetes workloads within ESXi. VMware Aria Operations collects metrics and data for the Supervisor Cluster. Supervisor Clusters contain Namespaces, which are resource pools that have dedicated memory, CPU, and storage.

Namespaces contain virtual machines with k8s activated. They are called k8s control VMs. These VMs are managed by vSphere. Therefore, you cannot take action on these VMs from within VMware Aria Operations.

DevOps engineers can run workloads on containers running inside vSphere Pods. They can create Tanzu k8s cluster inside a Namespace. A vSphere Pod is a VM with a small footprint that runs one or more Linux containers. It is the equivalent of a k8s pod. A Tanzu Kubernetes cluster is a full distribution of the open-source Kubernetes container orchestration software that is packaged, signed, and supported by VMware.

To understand the vSphere Tanzu Kubernetes architecture, see *Configuring and Managing vSphere with Kubernetes* in the vSphere documentation.

Workload management objects are excluded from the following workflows:

- Compliance
- Reclaim
- Rightsizing
- Workload optimization

## About Objects

Objects are the structural components of your mission-critical IT applications: virtual machines, datastores, virtual switches and port groups are examples of objects.

Because downtime equals cost - in unused resources and lost business opportunities - it's crucial that you successfully identify, monitor and track objects in your environment. The goal is to proactively isolate, troubleshoot and correct problems even before users are aware that anything is wrong.

When a user actually reports an issue, the solution should be quick and comprehensive.

For a complete list of objects that can be defined in VMware Aria Operations refer to Object Discovery.

VMware Aria Operations gives you visibility into objects including applications, storage and networks across physical, virtual and cloud infrastructures through a single interface that relates performance information to positive or negative events in the environment.

## Managing Objects

When you monitor a large infrastructure, the number of objects and corresponding metrics in VMware Aria Operations grows rapidly, especially as you add solutions that extend dynamic monitoring and alerts to more parts of your infrastructure. VMware Aria Operations gives you ample tools to stay abreast of events and issues.

## Adding Objects and Configuring Object Relationships

VMware Aria Operations automatically discovers objects and their relationships once you create an adapter instance. You have the added ability to manually add any objects that you want monitored and to configure object relationships using abstract concepts rather than the connections recorded by VMware Aria Operations. Where VMware Aria Operations might discover the classic parent-child relationships between objects, you can create relationships between objects that might not normally be related. For example, you could configure all the datastores supporting a company department to be related.

When objects are related, a problem with one object appears as an anomaly on related objects. So object relationships can help you to identify problems in your environment quickly. The object relationships that you create are called custom groups.

## Custom Groups

To create an automated management system you need some way to organize objects so that you can quickly gain insights. You can achieve a high level of automation using custom groups. You have multiple options for tailoring group attributes to support your monitoring strategy.

For example, you can designate a group either to be static or to be updated automatically with membership criteria that you designate. Consider a non-static group of all virtual machines that are powered on and have OS type Linux. When you power on a new Linux VM, it is automatically added to the group and the policy is applied.

For additional flexibility, you can also specify individual objects to be always included or excluded from a given custom group. Or you can have a different set of alerts and capacity calculations for your production environment versus your testing environments.

## Managing Applications

VMware Aria Operations allows you to create containers or objects that can contain a group of virtual machines or other objects in different structural tiers. This new application can then be managed as a single object, and have health badges and alarms aggregated from the child objects of the group.

For example, the system administrator of an online training system might request that you monitor components in the Web, application and database tiers of the training environment. You build an application that groups related training objects together in each tier. If a problem occurs with one of the objects, it is highlighted in the application display and you can investigate the source of the problem

### The Power of Object Management

Using the power of object management, including metrics and alerts - some prepackaged into dashboards and policies, others that you combine into custom monitoring tools - you'll keep a close watch on the objects, applications and systems that must stay up and running.

## Managing Objects in Your Environment

An object is the individual managed item in your environment for which VMware Aria Operations collects data, such as a router, switch, database, virtual machine, host, and vCenter Server instances.

The system requires specific information about each object. When you configure an adapter instance, VMware Aria Operations performs object discovery to start collecting data from the objects with which the adapter communicates.

An object can be a single entity, such as a database, or a container that holds other objects. For example, if you have multiple Web servers, you can define a single object for each Web server and define a separate container object to hold all of the Web server objects. Groups and applications are types of containers.

Categorize your objects using tags, so that you can easily find, group, or filter them later. A tag type can have multiple tag values. You or VMware Aria Operations assigns objects to tag values. When you select a tag value, VMware Aria Operations displays the objects associated with that tag. For example, if a tag type is Lifecycle and tag values are Development, Test, Pre-production, and Production, you might assign virtual machine objects VM1, VM2, or VM3 in your environment to one or more of these tag values, depending on the virtual machine function.

### Adding an Object to Your Environment

You might want to add an object by providing its information to VMware Aria Operations. For example, some solutions cannot discover all the objects that might be monitored. For these solutions, you must either use manual discovery or manually add the object.

When you add an individual object, you provide specific information about it, including the kind of adapter to use to make the connection and the connection method. For example, a vSAN adapter does not know the location of the vSAN devices that you want to monitor.

#### Prerequisites

Verify that an adapter is present for the object you plan to add. See Chapter 4 Integrating Data Sources with VMware Aria Operations.

Verify that an adapter is present for the object you plan to add. See the Getting Started with *VMware Aria Operations* .

**Note** Objects added to VMware Aria Operations via API will require an OSI license per object.

#### Procedure

1 From the left menu, click **Environment**, and then click **Inventory**.

2 On the toolbar, click the plus sign.

3 Use the topic menus to reveal all fields and provide the required information.

| Option | Description |
| --- | --- |
| **Display name** | Enter a name for the object. For example, enter **`vSAN-Host1`**. |
| **Description** | Enter any description. For example, enter **`vSAN-Host monitored with vSAN adapter`** |
| **Adapter type** | Select an adapter type. For example, select **vSAN Adapter**. |
| **Adapter instance** | Select an adapter instance. |
| **Object type** | Select an object type. For a vSAN adapter, you might select vSAN-Host. When you select the object type, the dialog box selections change to include information you provide so that VMware Aria Operations can find and connect with the selected object type. |
| **Host IP address** | Enter the host IP. For example, enter the IP address of vSAN-Host1. |
| **Port number** | Accept the default port number or enter a new value. |
| **Credential** | Select the Credential, or click the plus sign to add new login credentials for the object. |
| **Collection interval** | Enter the collection interval, in minutes. For example, if you expect the host to generate performance data every 5 minutes, set the collection interval to 5 minutes. |
| **Dynamic Thresholding.** | Accept the default, Yes. |

4 Click **OK** to add the object.

### Results

vSAN-Host1 appears in the Inventory as a host object type for the vSAN adapter type.

### What to do next

When you add an individual object, VMware Aria Operations does not begin collecting metrics for the object until you turn on data collection. See Inventory : List of Objects.

For each new object, VMware Aria Operations assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags. See Creating and Assigning Tags.

For each new object, VMware Aria Operations assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags.

## Configuring Object Relationships

VMware Aria Operations shows the relationship between objects in your environment. Most relationships are automatically formed when the objects are discovered by an installed adapter. In addition, you can use VMware Aria Operations to create relationships between objects that might not normally be related.

Objects are related physically, logically, or structurally.

- Physical relationships represent how objects connect in the physical world. For example, virtual machines running on a host are physically related.
- Logical relationships represent business silos. For example, all the storage objects in an environment are related to one another.
- Structural relationships represent a business value. For example, all the virtual machines that support a database are structurally related.

Solutions use adapters to monitor the objects in your environment so that physical relationship changes are reflected in VMware Aria Operations. To maintain logical or structural relationships, you can use VMware Aria Operations to define the object relationships. When objects are related, a problem with one object appears as an influence on related objects. So object relationships can help you to identify problems in your environment quickly.

Apart from the parent-child relationship, you can also define new relationships in VMware Aria Operations. The relationship between objects in your environment can be one-to-many, many-to-one, or one-one, the relationship can be defined in horizontal , vertical, or diagonal levels.

## Creating and Assigning Tags

A large enterprise can have thousands of objects defined in VMware Aria Operations. Creating object tags and tag values makes it easier to find objects and metrics. With object tags, you select the tag value assigned to an object and view the list of objects that are associated with that tag value.

A tag is a type of information, for example, Adapter Types. Adapter Types is a predefined tag. Tag values are individual instances of that type of information. For example, when the system discovers objects using the vCenter Adapter, it assigns all the objects to the vCenter Adapter tag value under the Adapter Types tag.

You can assign any number of objects to each tag value, and you can assign a single object to tag values under any number of tags. You typically look for an object by looking under its adapter type, its object type, and possibly other tags.

If an object tag is locked, you cannot add objects to it. VMware Aria Operations maintains locked object tags.

### What to read next

- Predefined Object Tags

  VMware Aria Operations includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

- Add an Object Tag and Assign Objects to the Tag

  An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

- Use a Tag to Find an Object

  The quickest way to find an object in VMware Aria Operations is to use tags. Using tags is more efficient than searching through the entire object list.

### Predefined Object Tags

VMware Aria Operations includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

For example, when you add an object, the system assigns it to the tag value for the collector it uses and the kind of object that it is. VMware Aria Operations creates tag values if they do not already exist.

If a predefined tag has no values, there is no object of that tag type. For example, if no applications are defined, the applications tag has no tag values.

Each tag value appears with the number of objects that have that tag. Tag values that have no objects appear with the value zero. You cannot delete the predefined tags or tag values.

**Table 13-1. Predefined Tags**

| Tag | Description |
|---|---|
| Collectors (Full Set) | Each defined collector is a tag value. Each object is assigned to the tag value for the collector that it uses when you add the object to VMware Aria Operations. The default collector is VMware Aria Operations Collector. |
| Applications (Full Set) | Each defined application is a tag value. When you add a tier to an application, or an object to a tier in an application, the tier is assigned to that tag value. |
| Maintenance Schedules (Full Set) | Each defined maintenance schedule is a tag value, and objects are assigned to the value when you give them a schedule by adding or editing them. |
| Adapter Types | Each adapter type is a tag value, and each object that uses that adapter type is given the tag value. |
| Adapter Instances | Each adapter instance is a tag value, and each object is assigned the tag value for the adapter instance or instances through which its metrics are collected. |
| Object Types | Each type of object is a tag value, and each object is assigned to the tag value for its type when you add the object. |

Table 13-1. Predefined Tags (continued)

| Tag | Description |
| --- | --- |
| Recently Added Objects | The last day, seven days, 10 days, and 30 days have tag values. Objects have this tag value as long as the tag value applies to them. |
| Object Statuses | Tag value assigned to objects that are not receiving data. |
| Collection States | Tag value assigned to indicate the object collection state, such as collecting or not collecting. |
| Health Ranges | Good (green), Warning (yellow), Immediate (orange), Critical (red), and Unknown (blue) health statuses have tag values. Each object is assigned the value for its current health status. |
| Entire Enterprise | The only tag value is Entire Enterprise Applications. This tag value is assigned to each application. |
| Licensing | Tag values are License Groups found under **Home > Administration > Mangement > Licensing. Objects** are assigned to the license groups during VMware Aria Operations installation. |
| Untag | Drag an object to this tag to delete the tag assignment. |

## Add an Object Tag and Assign Objects to the Tag

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

### Prerequisites

Become familiar with the predefined object tags.

### Procedure

1 From the left menu, click **Environment** and then, click **Inventory**.

2 Click the **Manage Tags** icon above the list of tags.

3 Click the **Add New Tag** icon to add a new row and type the name of the tag in the row.

   For example, type **`Cloud Objects`** and click **Update**.

4 With the new tag selected, click the **Add New Tag Value** icon to add a new row and type the name of the value in the row.

   For example, type **`Video Cloud`** and click **Update**.

5 Click **OK** to add the tag.

6 Click the tag to which you want to add objects to display the list of object tag values.

For example, click **Cloud Objects** to display the Video Cloud object tag value.

7 Drag objects from the list in the right pane of the Inventory onto the tag value name.

You can press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.

For example, if you want to assign data centers that are connected through the vCenter Adapter, type **`vCenter`** in the search filter and select the data center objects to add.

### Use a Tag to Find an Object

The quickest way to find an object in VMware Aria Operations is to use tags. Using tags is more efficient than searching through the entire object list.

Tag values that can also be tags are Applications and Object Types. For example, the Object Types tag has values for each object that is in VMware Aria Operations, such as Virtual Machine, which includes all the virtual machine objects in your environment. Each of these virtual machines is also a tag value for the Virtual Machine tag. You can expand the tag value list to select the value for which you want to see objects.

#### Procedure

1 From the left menu, click **Environment** and then, click **Inventory**.

2 In the tag list in the center pane, click a tag for an object with an assigned value.

When you click a tag, the list of values expands under the tag. The number of objects that is associated with each value appears next to the tag value.

A plus sign next to a tag value indicates that the value is also a tag and that it contains other tag values. You can click the plus sign to see the subvalues.

3 Select the tag value.

The objects that have that tag value appear in the pane on the right. If you select multiple tag values, the objects in the list depend on the values that you select.

| Tag Value Selection | Objects Displayed |
|---|---|
| More than one value for the same tag | The list includes objects that have either value. For example, if you select two values of the Object Types tag, such as Data Center and Host System, the list shows objects that have either value. |
| Values for two or more different tags | The list includes only objects that have all of the selected values. For example, if you select two values of the Object Types tag, such as Data Center and Host System, and you also select an adapter instance such as vC-1 of the vCenter Adapter instance tag, only Data Center or Host System objects associated with vC-1 appear in the list. Data Center or Host System objects associated with other adapter instances do not appear in the list, nor do objects that are not Data Center or Host System objects. |

4 Select the object from the list.

## Manage Object Tags Workspace

A large enterprise can have thousands of objects. When objects are assigned to a tag, and you choose to display objects with that tag value, the objects are easier to find on the Inventory list.

### Where You Find Manage Object Tags

From the left menu, click **Environment** and then, click **Inventory**.

Click the **Manage Tags** icon above the list of tags in the middle pane.

### Manage Object Tags Options

The Manage Object Tags screen appears with previously created tags listed. In the left pane, you add tags. In the right pane, you add tag values.

- Click **Add a New Tag** and type a new tag name, or select a tag to delete.
- For the selected tag, click **Add a New Tag Value** and type a new tag value name, or select a tag value to delete.
- For the GEO Location tag, tag values are identified with a location on a world map. Select the tag value and click **Manage Location** to display the **Manage Location** map and pick a geographical location. Objects assigned to that tag value appear in that geographical location on the Inventory : Geographical Map of Objects.

## Manage Object Type Tags Workspace

Every object in your environment is of a particular object type. You use Manage Object Type Tags to control the object type tags displayed.

### How Manage Object Type Tags Works

For every adapter instance installed, VMware Aria Operations discovers objects in your environment and starts collecting data from those objects.

### Where You Find Manage Object Type Tags

From the left menu, click **Environment** and then, click **Inventory**. Click the **Manage Object Type Tags** icon above the list of tags.

### Manage Object Type Tags Options

Depending on the number of adapters installed , there may be hundreds of object type tags. The Manage Object Type Tags options allow you to turn on or off the tags listed.

- Type a filter word to show the object type tags with the word.
- Name lists all the object type tags.
- To toggle the display of an object type tag, select the check box in the Show Tag column of its row.

## Inventory : List of Objects

VMware Aria Operations discovers objects in your environment for each adapter instance and lists them. From the complete list of all the objects in your environment, you can quickly access and configure any object. For example, you can check if a datastore is connected or providing data, or you can power on a virtual machine.

### How the List Works

Objects appear in a data grid. To find a particular object, you can sort a column in the grid or search for a filter word. In addition to sorting and searching, assigning objects to object tags makes it easier to find objects and metrics.

### Where You Find the List

From the left menu, click **Environment** and then, click **Inventory**. The system lists all the objects in your environment.

### Inventory List Options

The center pane includes object tag options. The right pane includes toolbar options for all of the objects in your environment.

Table 13-2. Object Tag Options

| Option | Description |
| --- | --- |
| Collapse all | Closes all the tag group selections. |
| Deselect All | Tags remain selected until deselected. Use this option to deselect all tags. |
| Manage Tags | Add a tag or tag value. See Manage Object Tags Workspace. |
| Manage Object Type Tags | There might be many object type tags. Use this option to choose the object type tags to display. See Manage Object Type Tags Workspace. |

Use the toolbar options to manage objects.

- Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.
- Select the object to manage from the list. If an object tag is selected, only objects of the selected tag value are listed. Column headings help you to identify the object. See Object List Widget.

### Table 13-3. Inventory Toolbar Options

| Option | Description |
| --- | --- |
| Action | Perform an action on the selected object. Available actions depend on the object type. For example, Power on VM applies to the selected virtual machine. See Actions in VMware Aria Operations |
| Open in external application | If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. For example, Open Virtual Machine in a vSphere Client or Search for VM logs in VMware Aria Operations for Logs. |
| Start Collecting | Turn on data collection for the selected object. |
| Stop Collecting | Do not collect data for the selected object. When data collection stops, VMware Aria Operations retains metric data for the object in case data collection starts at a later time. |
| Perform Multi-Collecting | If an object collects metrics through more than one adapter instance, select the adapter instance or instances for data collection. Does not apply to objects that do not use the adapter instance. |
| Edit object | Edit the selected object. For example, add or change the maintenance schedule for a virtual machine. If multiple objects of the same type are selected, common identifiers for the object type are editable. For example, change the VM entity name of multiple datastores with a single edit. See Manage Objects Workspace. |
| Add object | VMware Aria Operations discovers objects for most adapters. For adapters that do not support autodiscovery for all objects, the objects are manually added. See Manage Objects Workspace. |
| Discover Objects | Perform an IP scan to discover objects associated with a particular adapter. See Discover Objects Workspace. |
| Delete object | Remove the object from the list. |
| Start maintenance | Take the object offline for maintenance. See Manage Maintenance Schedules for Your Object Workspace. |
| End maintenance | Terminate the maintenance period and put the selected object back online. |
| Clear Selections | Clear all object selections. |
| Select All | Select all objects displayed. |
| Show Detail | Display the **Summary** tab of the selected object. |
| Per page | The number of objects to list per page. |

## Manage Objects Workspace

To collect data from an object, you might need to add an object or edit an existing object in your environment. For example, you might need to add objects for an adapter that does not support autodiscovery, or change the maintenance schedule of an existing object.

### Where You Find Manage Objects

From the left menu, click **Environment** and then, click **Inventory**. Click the plus sign to add an object or the edit icon to edit the selected object.

Items that appear in the window depend on the object that you are editing. Not all options can be changed.

### Table 13-4. Manage Objects Add or Edit Options

| Options | Description |
| --- | --- |
| Display name | Name of the object. Use only letters and numbers. Do not use nonalphanumeric characters or spaces. |
| Description | (Optional) For informational purposes only. |
| Adapter Type | If you are editing an object, you cannot change the adapter type. |
| Adapter Instance | If you are editing an object, you cannot change the adapter instance. |
| Object Type | If you are editing an object, you cannot change the object type. More configuration options might appear, depending on the object type. |
| Collection Interval | The collection interval for an object influences the collection status for the object. The collection interval for the adapter instance determines how often to collect data.<br>For example, if the collection interval for an adapter instance is set to five minutes, setting the collection interval for an object to 30 minutes prevents the object from having the No Data Receiving collection status after five collection cycles or 25 minutes.<br>The default value is 5 for adapter instances. You can increase this value, but not decrease it. To decrease it, contact VMware.<br>In cases of adapter instances such as VMwareAriaOperationsMgrAPI and HttpPost that push data to VMware Aria Operations through the REST API, when data is no longer pushed, the status of the adapter instance is changed to Down after five collection intervals. For example, if the process pushes data every ten minutes and is stopped, the status of the adapter instance is changed to Down after 50 minutes. This behavior is expected for these adapter instance types. |
| Dynamic Thresholding | On by default, to activate dynamic thresholding and early warning smart alerts. See VMware Aria Operations Dynamic Thresholds |

### Discover Objects Workspace

If VMware Aria Operations does not discover objects after an adapter instance is configured, use manual discovery. Discovering objects is more efficient than adding objects individually.

**Note** You use discovery to define objects for embedded adapters. VMware Aria Operations discovers objects that use external adapters.

### Where You Find Discover Objects

From the left menu, click **Environment** and then, click **Inventory**. Click **Discover Objects** in the List tool bar.

### Discover Objects

The Discoveries section of the `describe.xml` file for the adapter might include parameters for discovery information. The `describe.xml` file is in the `conf` sub folder of the adapter, for example `xyz_adapter3/conf/describe.xml`.

| Options | Description |
|---|---|
| Collector | Collector that VMware Aria Operations uses to discover objects. Only the VMware Aria Operations Collector is added during installation. |
| Adapter Type | Adapter type for the objects to discover. |
| Adapter Instance | Adapter instance of the selected adapter type. |
| Discovery Info | Selection depends on the adapter type. For example, for a vCenter adapter, the Discovery Info selection adds an option to discover objects of a particular object type. |
| Only New Objects | On by default, to omit objects that are already discovered. |

### Object Type

Based on your selection of Adapter Type, Adapter Instance, and Discovery Info, the Object Type values change. The available adapter types are NSX-T, vCenter Adapter, vSAN Adapter.

NSX-T Object Types:

- Edge Cluster
- Firewall Section
- Groups
- Load Balancer Pool
- Load Balancer Service
- Load Balancer Virtual Server
- Logical Router
- Logical Switch

- Router Service
- Transport Node
- Transport Zone

vCenter Object Types:

- Cluster Compute Resource
- Datastore
- Datacenter
- Folder
- Host System
- Resource Pool
- Virtual Machine

vSAN Object Types:

- Cache Disk
- Capacity Disk
- vSAN Cluster
- vSAN Disk Group
- vSAN Fault Domain
- vSAN Witness Host

### Discovery Results List

When you use the Discover Objects feature to manually discover objects in your environment, VMware Aria Operations lists the objects of the specified object type. You can choose the objects to monitor.

### Where You Find Discovery Results

From the left menu, click **Environment** and then, click **Inventory**. Click **Discover Objects** in the List tool bar.

After you make selections in the Discover Objects Workspace, click **OK**. With the default setting, VMware Aria Operations displays only newly discovered objects. See .

### Table 13-5. Object Types

| Options | Description |
| --- | --- |
| Object Type | Discovered object types of the Object Type selected on the Discover Objects Workspace. |
| Object Count | Number of objects of the object type. |
| Import | When selected, imports the object type. Option is active and selectable for newly discovered object types. |

Table 13-5. Object Types (continued)

| Options | Description |
| --- | --- |
| Collect | When selected, imports the object type and starts collecting data. Option is active and selectable for newly discovered object types. |
| Credential | If the object type requires a login credential to collect data from the object., the value is **True**. |

Double-click the Object Type to display a list of objects to monitor.

Table 13-6. Objects

| Options | Description |
| --- | --- |
| Object | Objects of the selected type that exist in the environment for the adapter. For example, the vCenter adapter discovers objects in the vCenter Server system. |
| Import | When selected, imports the object but does not start collecting data. Option is active and selectable for newly discovered objects that do not exist in the VMware Aria Operations environment . |
| Exists | Indicates that the object exists in the VMware Aria Operations environment. |
| Collect | When selected, imports the object and starts collecting data. Option is active and selectable for newly discovered objects that do not exist in the VMware Aria Operations environment. |

## Manage Maintenance Schedules for Your Object Workspace

You use maintenance mode to take an object offline. Many objects in your environment might be intentionally taken offline. For example, you might deactivate a server to update software. If VMware Aria Operations collects metrics when the object is offline, it might generate incorrect alerts that affect the data for the object's health. When an object is in maintenance mode, VMware Aria Operations does not collect metrics from the object and does not generate alerts for it.

### How Maintenance Schedules Work

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object into maintenance mode from midnight until 3 a.m. every Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can put an object in maintenance mode or take it out of maintenance mode, even if it has an assigned maintenance schedule.

### Where You Find Manage Maintenance Schedules

From the left menu, click **Environment** and then, click **Inventory**. Click **Start Maintenance** in the List tool bar.

Table 13-7. Manage Maintenance Schedules Options

| Options | Description |
| --- | --- |
| I will come back and end maintenance myself. | Maintenance mode starts for the selected object when you click **OK**. You must manually end maintenance mode for this object. |
| End maintenance in | Type the number of minutes that the object is in maintenance mode. |
| End maintenance on | Click the calendar icon, and select the date that maintenance mode ends. |

### Define Custom Property Workspace

In VMware Aria Operations, you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign custom properties to any subset of objects irrespective of the adapter kind and resource kind. You can use a mouse click, search filter, or a tag selector to select the correct object.

#### Where You Find Add/Edit Custom Property

From the left menu, click **Environment** and then click **Inventory**. Click **Add/Edit Custom Property** in the List tool bar.

Table 13-8. Add/Edit Custom Property

| Options | Description |
| --- | --- |
| Property Name | Select or enter a property name. |
| Type | Select the property type from the drop-down menu. |
| Value | Enter a value for the property. |

You can assign the custom properties defined in this page to the Custom Object Groups and New Groups.

For more information, see Custom Object Groups Workspace to Create a New Group.

## Inventory : Geographical Map of Objects

VMware Aria Operations discovers objects in your environment for each adapter. Objects that are assigned a GEO Location tag appear on a geographical map. You can use this map to quickly locate your objects in the world.

### How the Geographical Map Works

Objects with the GEO Location tag appear on a map of the world.

- To create a GEO Location tag, see Manage Object Tags Workspace.
- To assign objects to the tag, see Creating and Assigning Tags.

### Where You Find the Geographical Map

From the left menu, click **Environment** and then, click **Inventory**. Click the **Geographical** tab.

### Geographical Map Options

Use the plus sign to zoom in. Use the minus sign to zoom out. Click and drag to pan the map to the left or right.

# Managing Custom Object Groups in VMware Aria Operations

A custom object group is a container that includes one or more objects. VMware Aria Operations uses custom groups to collect data from the objects in the group, and report on the data collected.

## Why Use Custom Object Groups?

You use groups to categorize your objects and have the system collect data from the groups of objects and display the results in dashboards and views according to the way you define the data to appear.

You can create static groups of objects, or dynamic groups with criteria that determine group membership as VMware Aria Operations discovers and collects data from new objects added to the environment.

VMware Aria Operations provides commonly used object group types, such as World, Environment, and Licensing. The system uses the object group types to categorize groups of objects. You assign a group type to each group so that you can categorize and organize the groups of objects that you create.

## Types of Custom Object Groups

When you create custom groups, you can use rules to apply dynamic membership of objects to the group, or you can manually add the objects to the group. When you add an adapter, the groups associated with the adapter become available in VMware Aria Operations.

- Dynamic group membership. To dynamically update the membership of objects in a group, define rules when you create a group. VMware Aria Operations adds objects to the group based on the criteria that you define.
- Mixed membership, which includes dynamic and manual.
- Manual group membership. From the inventory of objects, you select objects to add as members to the group.
- Groups associated with adapters. Each adapter manages the membership of the group. For example, the vCenter Server adapter adds groups such as datastore, host, and network, for the container objects in the vSphere inventory. To modify these groups, you must do so in the adapter.

Administrators of VMware Aria Operations can set advanced permissions on custom groups. Users who have privileges to create groups can create custom groups of objects and have VMware Aria Operations apply a policy to each group to collect data from the objects and report the results in dashboards and views.

When you create a custom group, and assign a policy to the group, the system uses the criteria defined in the applied policy to collect data from and analyze the objects in the group. VMware Aria Operations reports on the status, problems, and recommendations for those objects based on the settings in the policy.

**Note** Only custom groups defined explicitly by users can be exported from or imported to VMware Aria Operations. Users are able to export or import multiple custom groups. Once an import function has been executed, the user must check to determine if a policy or policies should be associated with the imported group. Export-import operations are available for user defined (created explicitly by user) custom groups only.

## How Policies Help VMware Aria Operations Report On Object Groups

When you apply a policy to an object group, VMware Aria Operations uses threshold settings, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that you activated in the policy to collect data from the objects in the group, and report the results in dashboards and views.

When you create a new object group, you have the option to apply a policy to the group.

- To associate a policy with the custom object group, select the policy in the group creation wizard.
- To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

VMware Aria Operations applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, VMware Aria Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria Operations associates the highest ranking policy with that object.

## User Scenario: Creating Custom Object Groups

As a system administrator, you must monitor the capacity for your clusters, hosts, and virtual machines. VMware Aria Operations monitors them at different service levels to ensure that these objects adhere to the policies established for your IT department, and discovers and monitors new objects added to the environment. You have VMware Aria Operations apply policies to the object groups to analyze, monitor, and report on the status of their capacity levels.

To have VMware Aria Operations monitor the capacity levels for your objects to ensure that they adhere to your policies for your service levels, you categorize your objects into Platinum, Gold, and Silver object groups to support the service tiers established.

You create a group type, and create dynamic object groups for each service level. You define membership criteria for each dynamic object group to have VMware Aria Operations keep the membership of objects current. For each dynamic object group, you assign the group type, and add criteria to maintain membership of your objects in the group. To associate a policy with the custom object group, you can select the policy in the group creation wizard.

### Prerequisites

- Know the objects that exist in your environment, and the service levels that they support.
- Understand the policies required to monitor your objects.
- Verify that policies are available to monitor the capacity of your objects.

### Procedure

1. To create a group type to identify service level monitoring, from the left menu click **Environment** and then, click the horizontal ellipsis and select **Group Types** from the pop-up menu.
2. On the Group Types toolbar, click **Add** and type **`Service Level Capacity`** for the group type.

   Your group type appears in the list.
3. From the left menu, click **Environment** and then, click **Custom Groups**.
4. To create a new object group, click **Add**.

   The New Group workspace appears where you define the data and membership criteria for the dynamic group.

   a. In the Name text box, enter a name for the object group, such as **`Platinum_Objects`**.

   b. In the **Group Type** drop-down menu, select **Service Level Capacity**.

   c. (Optional) In the **Policy** drop-down menu, select your service level policy that has thresholds set to monitor the capacity of your objects.

      To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

   d. Select the **Keep group membership up to date** check box so that VMware Aria Operations can discover objects that meet the criteria, and add those objects to the group.

5 Define the membership for virtual machines in your new dynamic object group to monitor them as platinum objects.

   a From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Virtual Machine**.

   b From the empty drop-down menu for the criteria, select **Metrics**.

   c From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.

   d From the conditional value drop-down menu, select **is less than**.

   e From the **Metric value** drop-down menu, type `10`.

6 Define the membership for host systems in your new dynamic object group to monitor them as platinum objects.

   a Click **Add another criteria set**.

   b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Host System**.

   c From the empty drop-down menu for the criteria, select **Metrics**.

   d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.

   e From the conditional value drop-down menu, select **is less than**.

   f From the **Metric value** drop-down menu, type `100`.

7 Define the membership for cluster compute resources in your new dynamic object group.

   a Click **Add another criteria set**.

   b From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Cluster Compute Resources**.

   c From the empty drop-down menu for the criteria, select **Metrics**.

   d From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **capacityRemaining**.

   e From the conditional value drop-down menu, select **is less than**.

   f From the **Metric value** drop-down menu, type `1000`.

   g Click **Preview** to determine whether objects already match this criteria.

8 Click **OK** to save your group.

   When you save your new dynamic group, the group appears in the Service Level Capacity folder, and in the list of groups on the **Groups** tab.

9 Wait five minutes for VMware Aria Operations to collect data from the objects in your environment.

### Results

VMware Aria Operations collects data from the cluster compute resources, host systems, and virtual machines in your environment, according to the metrics that you defined in the group and the thresholds defined in the policy that is applied to the group, and displays the results about your objects in dashboards and views.

### What to do next

To monitor the capacity levels for your platinum objects, create a dashboard, and add widgets to the dashboard. See Chapter 14 Dashboards in VMware Aria Operations.

## Object Group Types in VMware Aria Operations

An object group type is an identifier that you apply to a specific group of objects in your environment to categorize them. You can add new group types, and apply them to groups of objects so that VMware Aria Operations can collect data from the object group and display the results in the dashboards and views.

### How the Group Types Work

Use group types to categorize your objects so that the system can apply policies to them to track, and display specific status, such as alerts, workload, faults, risk, and so on.

When you create a new group type, VMware Aria Operations adds it to the existing list of group types, and creates a new folder with the name of your group type in the Environment Custom Groups list.

When you create a new group of objects, you assign a group type to that group of objects. You add objects from the inventory trees to your custom group, then create your dashboard, add widgets to the dashboard, and configure the widgets to display the data collected from the objects in the group. You can then monitor and manage the objects.

You can apply a group type to a group of objects that you create manually, or to object groups that you cannot modify, such those added by adapters. Each adapter that you add to VMware Aria Operations adds one or more static groups of objects to group the data received from the adapter sources.

The list of group types appears in the Content area under Group Types. The custom object groups appear in the Environment area under Custom Groups.

### Where You Create and Modify a Group Type

From the left menu, click **Environment** and then, click **Custom Groups**. Click **Group Types** next to custom groups. You can add, edit, delete, and select groups from the group types page.

### Group Type Options

You can add, edit, or delete group types. You cannot edit group types that are created by adapters.

## Groups Tab on the Environment Overview Pane

Groups are containers that can contain any number and type of objects in your environment. VMware Aria Operations collects data from the objects in the group and displays the results in dashboards and views that you define.

### How Groups Work

Groups are installed with VMware Aria Operations, created by an adapter, or created by a user. Based on the group criteria, you can use groups to organize your environment and monitor all objects in the group together. You can also assign policies to groups and make group membership dynamic.

For example, if you have a set of vSphere hosts and you do not want to generate alerts when the host goes into maintenance mode, you can put the vSphere hosts in a group and assign a policy that includes a maintenance schedule setting. During the maintenance period, VMware Aria Operations ignores any metrics for those objects and does not generate any alerts. After the maintenance period ends, VMware Aria Operations returns to monitoring the objects and generates alerts if an outage occurs.

### Where You Find Custom Groups

From the left menu, click **Environment** and then, click **Custom Groups**.

### Custom Group Options

Click **ADD** button to add a group. You can only edit, clone, or delete a user-created group. You cannot modify groups installed with VMware Aria Operations or by an adapter.

You can click the **Horizontal Ellipses** to import or export the custom group. The Groups data grid displays an overview of the state of each group. You can use the All Filters option to sort the custom groups based on Name, ID, Group Type, and Description columns.

To sort the list of custom groups based on columns, click the column heading of the following columns:

- Name
- Health
- Risk
- Efficiency
- Description
- Members Count

Table 13-9. Group Data Grid Options

| Option | Description |
| --- | --- |
| Name | Select the group name to display a summary of the group. Select to the right of the name to edit, clone, or delete the group. |
| Summary | Criticality of the health, risk, and efficiency of any group. Click a group with a red, orange, or yellow criticality to get more details about potential problems with objects in the group. |
| Members Count | Displays the number of members in the selected group. |
| Policy | Displays the policy associated with the selected group. |
| Dynamic Membership | Displays whether the group is static or dynamic. The available options are true and false. |
| Defined by | Displays who has defined the attributes of the group. The available options are:<br>■ System<br>■ User Defined<br>■ Management Pack |

## Custom Object Groups Workspace

You can create and edit custom groups of objects to have VMware Aria Operations collect data from the objects and display the results in the dashboards and views so that you can monitor your objects and take action on them when problems occur.

### How the Custom Groups Workspace Works

When you create a new object group, you define a meaningful group name, and select the group type. To associate the custom object group with a policy for analysis, you select the policy in the group creation wizard. You can leave the policy selection blank to not associate a policy with the object group. When the policy selection is blank, the custom object group is associated with the policy that is designated as the default policy.

You select the object types, and determine whether membership in the object group is static, dynamic, or a combination of static and dynamic membership.

- To create a static object group, you add objects to the group. You do not include criteria for object membership.
- To create a dynamic object group that VMware Aria Operations updates based on specific criteria, you select the object type and define membership criteria for the group based on metrics, relationships, and properties.

When you add objects to a custom object group, a new folder appears in the Custom Groups navigation pane on the left, and includes the member objects.

### Where You Create and Modify Object Groups

To create or modify static or dynamic object groups, or object groups that have a combination of static and dynamic membership, click **Environment > Custom Groups**. The **Custom Groups** tab displays a list of custom object groups, and the object groups for adapters added to VMware Aria Operations.

To edit existing groups, select a group and click the edit icon on the **Custom Groups** tab.

### Custom Object Groups Workspace to Create a New Group

You can create a new object group, define custom properties, assign a group type and objects to the group. When you create the group, you can assign a policy, or leave the policy selection blank to apply the default policy. VMware Aria Operations collects data from the objects in the group based on the settings in the policy that is associated with the group. The results appear in the dashboards and views.

### Where You Assign Custom Group Type, Policy, and Membership

To assign the group type, policy, and membership, click **Environment**, click **Custom Groups**, and click **Add** to add a new group. In the New Group workspace, you can define the membership criteria, and select the objects to include or exclude.

To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

### Table 13-10. New Group Workspace

| Option | Description |
|---|---|
| **Custom Group Tab** | |
| Name | Meaningful name of the object group. |
| | Meaningful description for the object group. |
| Group Type | Categorization for the object group. New custom groups appear in a dedicated folder in the Custom Groups navigation pane on the left. |
| Policy | Assigns a policy to one or more groups of objects to have VMware Aria Operations analyze the objects according to the settings in your policy, trigger alerts when the defined thresholds are violated, and display the results in dashboards, views, and reports. You can assign a policy to the group when you create the group, or you can assign it later from the edit custom group wizard or from the policies area. |
| Keep group membership up to date | For dynamic object groups, VMware Aria Operations can discover objects that match the criteria for the group membership according to the rules that you define, and update the group members based on the search results. |
| **Define Membership Criteria Tab** | |

### Table 13-10. New Group Workspace (continued)

| Option | Description |
| --- | --- |
| Membership Criteria | Defines the criteria for a dynamic object group and has VMware Aria Operations keep the object membership of the group current.<br>■ Object Type drop-down menu. Selects the type of objects to add to the group, such as virtual machines.<br>■ Metrics, Relationship, and Properties criteria drop-down menu. Defines the criteria for VMware Aria Operations to apply to collect data from the selected objects.<br>  ■ Metrics. An instance of a data type, or attribute, that varies based on the object type. A metric is used as measurement criteria to collect data from objects. For example, you can select system attributes as a metric, where an attribute is a type of data that VMware Aria Operations collects from objects.<br>  ■ Relationship. Indicates how the object is related to other objects. For example, you can require a virtual machine object to be a child object that contains a certain word in the vSphere Hosts and Clusters navigation tree.<br>  ■ Properties. Identifies a configuration parameter for the object. For example, you can require a virtual machine to have a memory limit that is greater than 100KB.<br>  ■ Object Name. Specifies the name of the object. For example, you can create custom group which contains object with specific names.<br>  ■ Tag. Specifies a tag for the object. For example, you can create membership criteria based on specific tags.<br>■ You can use the filtering option to define the membership criteria for a group. The available options are:<br>  ■ is - Displays the result of the typed text. Example - "!", displays all results equal to the symbol "!"<br>  ■ is not - Displays the result of the typed text. Example - "1_node", displays all results not equal to "1_node"<br>  ■ contains -Displays the result of the typed text. Example - "An", displays all results that contain "An"<br>  ■ does not contain - Displays the result of the typed text. Example - "An", displays all results that does not contain "An"<br>  ■ starts with - Displays the result of the typed text. Example - "!!", displays all results that starts with "!!"<br>  ■ ends with - Displays the result of the typed text. Example - "N", displays all results that ends with "N"<br>  ■ does not start with - Displays the result of the typed text. Example - "N", displays all results that does not start with "N"<br>  ■ does not end with - Displays the result of the typed text. Example - "S", displays all results that does not end with "S"<br>  ■ matches regular expression - Displays the result if it matches the regular expression. Example .*8\.\d*<br>  ■ does not match regular expression Displays the result if it does not match the regular expression. Example .*8\.\d*<br>■ Add. Includes another metric, relationship, or property for the object type.<br>■ Remove. Deletes the selected object type from the membership criteria, or delete the selected metric, relationship, or property type from the criteria for the object type. |

### Table 13-10. New Group Workspace (continued)

| Option | Description |
| --- | --- |
| | ■ Add another criteria set. Adds another object type to add to the group. For example, you might want to create a single object group to track vCenter Server instances and Host Systems. |
| **Objects to Include/Exclude Tab** | |
| Objects To Always Exclude | Determine which objects to exclude from the group every time VMware Aria Operations collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of VMware Aria Operations, these objects were called a denylist.<br>■ **Select Objects Pane > Inventory Tab**. Displays the list of available object groups. Click on an object group to view the objects in the right pane. You can select some or all the objects in the object group and drag and drop it to the left pane for exclusion from the custom group.<br>■ **Select Objects Pane > Custom Group Tab**. Displays the list of available custom groups. Click on a custom group and drag and drop it to the left pane for exclusion from the custom group.<br>■ From the left pane, select the objects you want to delete and click **Remove** to delete the objects that you have added for exclusion from the custom group. |
| Objects To Always Include | Determine which objects to include in the group every time VMware Aria Operations collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of VMware Aria Operations, these objects were called an allowlist.<br>■ **Select Objects Pane > Inventory Tab**. Displays the list of available object groups. Click on an object group to view the objects in the right pane. You can select some or all the objects in the object group and drag and drop them to the left pane for inclusion into the custom group.<br>■ **Select Objects Pane > Custom Group Tab**. Displays the list of available custom groups. Click on a custom group and drag and drop it to the left pane for inclusion into the custom group.<br>■ From the left pane, select the objects you want to delete and click **Remove** to delete the objects that you have added for inclusion to the custom group. |
| **Assign Custom Properties Tab** | |
| Assign Custom Properties | In VMware Aria Operations, you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign the newly defined custom properties to new groups or existing groups.<br>■ Property Name. Select or specify a name for the custom property.<br>■ Type. Select the type of custom property from the drop-down menu.<br>The custom property can either be a string or a numeric.<br>■ Value. Specify a custom property value, which should be assigned to this custom property when an object is added to the group.<br>■ Reset Value. Specify a custom property value, which should be assigned to this custom property when an object leaves the group.<br>■ **Reset**. Resets the custom property to a non-zero value.<br>■ **Remove**. Removes the custom property from the group.<br>■ **Add Another Custom Property**. Adds another custom property to the group. |

- Use the **Previous** and **Next** buttons to navigate between the tabs.
- Use the **Next** button to move to the next tab.
- Use **Create** button to save and create the custom group.
- Use the **Preview** button to preview the list of objects in the group and to verify that the criteria you defined is applicable to the group of objects. If the criteria that you defined is valid, the preview displays applicable objects. If the criteria is not valid, the preview does not display any objects.

# Managing Application Groups

An application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. VMware Aria Operations builds an application to determine how your environment is affected when one or more components in an application experiences problems, and to monitor the overall health and performance of the application. Object membership in an application is not dynamic. To change the application, you manually modify the objects in the container.

## Reasons to Use Applications

VMware Aria Operations collects data from components in the application and displays the results in a summary dashboard for each application with a real-time analysis for any of the components. If a component experiences problems, you can see where in the application the problems arise, and determine how problems spread to other objects.

**Note** VMware Aria Operations provides for calendar periodicity. If your application includes work performed on a specific day of the month, for example, the 15th of the month or the last day of the month, this calendar function identifies the pattern after six cycles of the application. Once the pattern is recognized, the system can forecast accurately into the future. Because the system acquires its information from the input data, you do not have to give any details about how you schedule periodical work.

## Applications Tab on the Environment Overview Pane

Applications are groups of related objects in your environment that mimic an application in your business. Use the summary to track the health of objects in the application and help troubleshoot performance issues.

### How Applications Work

In VMware Aria Operations, each application contains one or more tiers and each tier contains one or more objects. The tier is a convenient way to organize objects that perform a specific task in an application. For example, you can group all of your database servers together in a tier.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the application.

Construct an application to view a particular segment of your business. The application shows how the performance of one object affects other objects in the same application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data for your business, you see a yellow, orange, or red status if the application health is degrading. Starting with the application summary dashboard, you can investigate which server is causing or exhibiting the problem.

### Where You Find Applications

In the menu, click **Environment**, then click the **Applications** tab.

Applications defined in a previous release of VMware Aria Operations appear after an upgrade.

### Application Options

Select an application to edit or delete, or click the **ADD** button to add an application.

The Applications data grid displays an overview of the state of each application.

Table 13-11. Application Data Grid Options

| Option | Description |
| --- | --- |
| Name | Select the application name to display a summary of the application. Select to the right of the name to edit or delete the application. |
| Summary | Criticality of the health, risk, and efficiency of any application. Click an application with a red, orange, or yellow criticality to see more details about potential problems with objects in the application. |

## User Scenario: Adding an Application

As the system administrator of an online training system, you must monitor components in the Web, application, and database tiers of your environment that can affect the performance of the system. You build an application that groups related objects together in each tier. If a problem occurs with one of the objects, it is reflected in the application display and you can open a summary to investigate the source of the problem further.

In your application, you add the DB-related objects that store data for the training system in a tier, Web-related objects that run the user interface in a tier, and application-related objects that process the data for the training system in a tier. The network tier might not be needed. Use this model to develop your application.

### Procedure

1. From the left menu, click **Environment**, then click **Applications** in the left pane.
2. Click **ADD**.
3. Click **Basic n-tier Web App** and click **OK**.

   The Application Management page that appears has two rows. Select objects from the bottom row to populate the tiers in the top row.

4 Type a meaningful name such as **`Online Training Application`** in the Application text box.

5 For each of the Web, application and database tiers listed, add the objects to the Tier Objects section.

   a Select a tier name. This is the tier that you populate.

   b To the left of the object row, select object tags to filter for objects that have that tag value. Click the tag name once to select the tag from the list and click the tag name again to deselect the tag from the list. If you select multiple tags, objects displayed depend on the values that you select.

      You can also search for the object by name.

   c To the right of the object row, select the objects to add to the tier.

   d Drag the objects to the Tier Objects section.

6 Click Save to save the application.

#### Results

The new application appears in the list of applications on the Environment Overview Applications page. If any of the components in any of the tiers develops a problem, the application displays a yellow or red status.

#### What to do next

To investigate the source of the problem, click the application name and see Evaluate Object Information Using Badge Alerts and Summary Tab .

To investigate the source of the problem, click the application name and evaluate the object summary information. See the *VMware Aria Operations User Guide* .

## Add Application

When you add an application to an environment, you select from a list of predefined templates or create your own custom template, to group the objects to monitor in your application.

### Where You Find Add Application

From the left menu, click **Environment** and then, click **Applications** tab. From the Applications page, click **Add**.

### Add Applications Options

Each predefined template provides you with a list of suggested tiers designed to help you group related objects that perform a specific task in your application. After you select an option, you can alter the selection and number of tiers on the Application Management page.

| Option | Description |
| --- | --- |
| Basic n-tier Web App | Use this template for any basic application. |
| Advanced n-tier Web App | Use this template for an application that monitors more physical devices, such as the devices that VMware Aria Operations discovers when you add a network-related Management Pack or Management Packs. |
| Legacy non-Web App | Use this template for an application that has no Web-related objects. |
| Network | Use this template for an application that has only network-related objects. |
| Custom | Select this option to build your own application topology. |

## Application Management Dialog Box

You use Application Management to select the objects for your application. The objects you select are grouped in tiers and help you to track the health of your application.

### Where You Find Application Management

From the left menu, click **Environment** and then, click **Applications**. On the **Applications** tab, click **Add**. After you select an application template, click OK.

### Application Management Options

At the top of the screen, enter a new application name or use the default name from the Add Application page. The application name must be unique.

Below the name, the page is divided into the tier row and the objects row. On each row, selections in the pane on the left filter the selections in the pane on the right.

The tier row is where you select the tiers to populate with objects to monitor for the application.

#### Table 13-12. Tier Row

| Option | Description |
| --- | --- |
| Tiers pane | Select the tier where you want to place your objects. You can add or delete tiers to fit your application. |
| Tier Objects pane | Add or remove objects that serve a common function and to monitor. For example, to monitor all the virtual machines that are database servers for the application, put them in the database tier. |

The object row is where you select objects to add to the tiers.

#### Table 13-13. Object Row

| Option | Description |
| --- | --- |
| Object Tags pane | Expand a tag to see a group of objects with that tag value. For example, if Adapter Types is an object tag, the tag values include vCenter Adapter, and an object is an adapter instance. Objects are not displayed. The tag filters the object pane. To select a tag value, click once. To deselect a tag value, click twice. Tag values remain selected until they are deselected. |
| Objects pane | Drag an object with the object tag value to add to the Tier Objects pane. To find an object, search by name. Each object listed includes identifier information to help distinguish between objects of similar names. **Add All Objects To Parent** adds all the objects to a tier. |

# Dashboards in VMware Aria Operations

14

Dashboards present a visual overview of the performance and state of objects in your virtual infrastructure. You use dashboards to determine the nature and timeframe of existing and potential issues with your environment. You create dashboards by adding widgets to a dashboard and configuring them.

VMware Aria Operations collects performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems. The data and analysis are presented through alerts, in configurable dashboards, on predefined pages, and in several predefined dashboards.

- You can start with several predefined dashboards in VMware Aria Operations.
- You can create extra ones that meet your specific needs using widgets, views, badges, and filters to change the focus of the information.
- You can clone and edit the predefined dashboards or start from scratch.
- To display data that shows dependencies, you can add widget interactions in dashboards.
- You can provide role-based access to various dashboards for better collaboration in teams.

Table 14-1. Features

| Features | Description |
| --- | --- |
| Manage | You can also manage dashboards by clicking **Visualize > Dashboards**. From the Dashboards panel, click **Manage**. |
| Create | Use this option to create a dashboard. See Create and Configure Dashboards. |
| Search | You can search for a dashboard across the **Favorites**, **Recents**, and **All** folders in the **Dashboards** panel. |
| Favorites | You can mark a dashboard as a favorite using the **Favorite** icon at the top of each dashboard. All the dashboards that you have marked as a favorite, are listed under the Favorites folder in the **Dashboards** panel. |

### Table 14-1. Features (continued)

| Features | Description |
| --- | --- |
| Recents | The dashboards are listed in the order in which you select them, with the most recent dashboard that you selected, appearing at the top. Up to ten dashboards can be displayed as Recent dashboards.<br>If you do not pin the dashboard and log out of the user interface, on logging back in, the dashboard is removed from the Recents folder. |
| Shared | If you have shared the dashboard, the shared icon is displayed against the dashboard name. |
| All | Lists the dashboard folders and the dashboards that are activated. You can use this menu for quick navigation through your dashboards. When you navigate to a dashboard using the **Visualize > Dashboards** option, the dashboards are listed in the left pane of the **Dashboards** panel under **All**. You can also search for dashboards using keywords and letters. |
| Actions | Available dashboard actions, such as edit, delete, remove dashboard from the menu, set as dashboard landing page, and set as the Home landing page. These actions are applied directly to the dashboard that you are on.<br>To remove the dashboard as the Home landing page, from the dashboard that has been set as the Home landing page, select **Actions > Reset from Home landing page**.<br>To remove the dashboard as the dashboard landing page, from the dashboard that has been set as the landing page, select **Actions > Reset from Dashboards landing page**. |
| Dashboard Time | The dashboard time panel is activated by default on all predefined and user-created dashboards. Using this option, you can select a time for the widgets in the dashboard. The default time is 6 hours. The pre-defined time/day options in the panel are 1 hour, 6 hours, 24 hours, or 7 days. You can also set a customized time option.<br>To activate widgets to use the dashboard time, select ***Date Controls/Time Range*** **> Dashboard Time** from the widget toolbar. Some widgets have **Dashboard Time** as the default option. For example, **Metric Chart**, **View**, **Rolling View**, **Sparkline**, **Health Chart**, and **Mashup Chart** widgets.<br>Dashboard time persists if:<br>■ You activate a widget in a dashboard to use the dashboard time and then log out and log back in, or<br>■ You activate a widget in a dashboard to use the dashboard time, and you export and then import the dashboard into another instance of VMware Aria Operations. |

Read the following topics next:

- Accessing Predefined Dashboards

- Types of Dashboards
- Create and Configure Dashboards
- Manage Dashboards
- Dashboards Actions and Options


# Accessing Predefined Dashboards

You can access some of the useful, predefined dashboards from the **Dashboards** home page.

To access these dashboards, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Home**.

The dashboards are categorized as follows: Availability, Configuration, Inventory, Performance, Capacity, and Cost. To easily access some of the useful, predefined dashboards under these categories, click on the drop-down button against the selected category and click on the specific dashboard.

# Types of Dashboards

You can use the predefined dashboards or create your own custom dashboard in VMware Aria Operations.

See Chapter 15 Predefined Dashboards for more information.

## Custom Dashboards

You can create dashboards that meet your environment needs in VMware Aria Operations.

For information about creating a dashboard, see Create and Configure Dashboards.

# Create and Configure Dashboards

To view the status of all objects in VMware Aria Operations, create a dashboard by adding widgets or views. You can create and modify dashboards and configure them to meet your environment needs.

### Procedure

1. From the left menu, click **Visualize > Dashboards**.
2. From the **Dashboards** panel on the left, click **Create**.

3 Complete the following steps to:

   a Enter a name for the dashboard.

     Dashboard Name

   b Add widgets or views to the dashboard.

     Widget or View List Details

   c Configure widget interactions.

     Widget and View Interactions Details

   d Create dashboard navigation.

     Dashboard Navigation Details

4 Click **Save**.

5 Click **Actions > Edit Dashboard** to modify the dashboard.

## Dashboard Name

The name and visualization of the dashboard as it appears on the VMware Aria Operations Home page.

### Where You Add a Name in a Dashboard

To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. Enter a name in the **New Dashboard** field.

To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit Dashboard**.

If you use a forward slash while entering a name, the forward slash acts as a group divider and creates a folder with the specified name in the dashboards list if the name does not exist. For example, if you name a dashboard **`clusters/hosts`**, the dashboard is named `hosts` under the group `clusters`.

## Widget or View List Details

VMware Aria Operations provides a list of widgets or views that you can add to your dashboard to monitor specific metrics and properties of objects in your environment.

### Where You Add Widgets or Views to a Dashboard

To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard.

To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit Dashboard**.

## How to Add Widgets or Views to a Dashboard

In the widgets list panel, you see a list of predefined VMware Aria Operations widgets or views that are most commonly used. Click **Show More** to view all the widgets and views. To add a widget or view to the dashboard workspace in the upper panel, you can:

- Drag the widget or view to the dashboard workspace in the upper panel, or
- Click on a widget or view from the widgets list panel, view a brief description of the widget or view from the pop up box and then click **Add to Dashboard**, or
- Double-click on the widget or view.

To locate a widget or view, you can enter the name or part of the name of a widget or view in the **Filter > Name** option. For example, when you enter `top`, the list is filtered to display the Top Alerts, Top-N, and Topology Graph widgets. You can then select the widget you require. You can also filter by either widgets or views by selecting **Filter > Show** to add a widget or view to the dashboard. Drag the widget or view to the dashboard workspace in the upper panel.

---

**Note** The following widgets are deprecated. Usage of these widgets is discouraged as they will be phased out in future releases. Deprecated widgets are marked with a yellow triangle.

- Current Policy
- Weather Map
- Anomalies
- DRS Cluster Settings
- Efficiency
- Environment Status
- Risk
- Environment
- Container Overview
- Faults

---

Most widgets or views must be configured individually to display information. For more information about how to configure each widget, see Chapter 16 Widgets in VMware Aria Operations.

## How to Arrange Widgets or Views in a Dashboard

You can modify your dashboard layout to suit your needs. By default, the first widgets or views that you add are automatically arranged horizontally wherever you place them.

- To position a widget or a view, drag the widget or view to the desired location in the layout. Other widgets and views automatically rearrange to make room.
- To resize a widget or a view, drag the bottom-right corner of the widget or the view.

- To maximize or minimize a widget or a view, use the maximize and minimize options in the top-right corner.

## Widget and View Interactions Details

You can connect widgets and views so that the information they show depends on each other.

### Where You Create Widget and View Interactions

To create interactions for widgets or views in a dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. From the toolbar, click **Show Interactions**.

To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit Dashboard**.

Internal interactions are interactions defined between widgets in a single dashboard. External interactions are interactions defined between a widget in one dashboard and widget/s in another dashboard.

### How to Create and Remove Widget Interactions

The list of available interactions depends on the widgets or views in the dashboard. Widgets and views can provide, receive, and can both provide and receive interactions at the same time.

To create interactions, click **Show Interactions**. Click a provider plug and drag to the receiver. You can also apply interactions from receiver to provider plugs. For more information about how interactions work, see Widget Interactions.

To remove interactions, click on the interaction line and select **Remove Interaction**. You can also click the provider plug and select **Remove Interaction > <widget name>**.

**Note** You can create up to 25 widget interactions in a dashboard.

### How to Easily View Widget Interactions

To view widget interactions from a dashboard, click on an object from a provider widget that has widget interactions defined. A window appears with the external and internal widget interaction details. Clicking on an external interaction takes you to the external dashboard. Click on the internal interaction to view the details in the receiver widget.

## Dashboard Navigation Details

You can apply sections or context from one dashboard to another. You can connect widgets and views to widgets and views in the same dashboard or to other dashboards to investigate problems or better analyze the provided information.

### Where You Add Another Dashboard

To create dashboard navigation to a dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. In the dashboard workspace, click **Show Interactions**. From the **Select Another Dashboard** drop-down menu, select the dashboard to which you want to navigate.

To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit Dashboard**.

### How Dashboard Navigation Works

You can create dashboard navigation only for provider widgets and views. The provider widget or view sends information to the destination widget or view. When you create dashboard navigation, the destination widgets or views are filtered based on the information type they can receive.

### How to Add Dashboard Navigation to a Dashboard

The list of available dashboards for navigation depends on the available dashboards and the widgets and views in the current dashboard. To add navigation, you can drag from a sender widget interaction plug to a receiver widget interaction plug. You can select more than one applicable widget or view.

**Note** If a dashboard is unavailable for selection, it is unavailable for dashboard navigation.

The Dashboard Navigation icon ( ) appears in the top menu of each widget or view when a dashboard navigation is available.

After you have set widget interaction in the provider dashboard, the widget and menu bar are highlighted and two arrows appear in the top-left corner of the widget. After you have set widget interaction, clicking the object in the provider widget takes you to the receiver widget of the navigated dashboard.

## Manage Dashboards

You can select dashboards individually or as a group and perform several actions.

To manage your dashboards, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Manage**. Use the options from the horizontal ellipsis next to the **Add** option.

All the dashboards are listed on this page. You can filter the dashboards based on the name of the dashboard, the dashboard folder, activated dashboards, shared dashboards, or the dashboard owner. You can click **Add** to create a dashboard. For information about creating a dashboard, see Create and Configure Dashboards.

You can select a dashboard from the list, click the vertical ellipsis against each dashboard, and select the various options such as edit, delete, clone, and deactivate a dashboard. You can also change ownership of dashboards and export the dashboard. By default, the list of dashboards is sorted by name and all the columns can be sorted.

**Note** A wrench icon appears when the data in an imported dashboard depends on the existence of one or more adapters that are currently not present. The wrench icon disappears if the required data in an imported dashboard appears in VMware Aria Operations after configuration.

Imported dashboards regardless of used data, remain stuck and include a wrench icon if the dashboard that is stuck (with the wrench icon), already exists.

### Datagrid Options

| Column Names | Description |
|---|---|
| Name | Displays the name of the dashboard. |
| Folder | Lists the folder to which each dashboard belongs. |
| Description | Displays the description of the dashboard. |
| Activated | Displays whether the dashboard is activated or not. You can also activate and deactivate the dashboard. |
| URL | Displays whether the dashboard is shared externally. For dashboards that have been shared, click to view the shared links. |
| Shared | Displays whether the dashboard is shared internally. Click the icon to view and edit the groups to which the dashboard has been shared from the **Group Sharing** dialog box.<br>From the **Group Sharing** dialog box, to share dashboard edit privileges with the group, you can click the **Allow Editing** check box or click the pencil icon if a dashboard is shared with a group.<br>**Note** Users who are part of the user group to whom the dashboard has been shared, will see an unlocked or locked icon in the **Shared** column. The icon is unlocked when the user with whom the dashboard is shared can edit it, and the icon is locked when the dashboard cannot be edited. |
| Editable | Displays if the dashboard is editable. |
| Owner | Displays the owner of the dashboard. |
| Report Usage | Displays the number of reports where the dashboard is used. Click the number in the column to view the name of the report/s. |

| Column Names | Description |
| --- | --- |
| Last Modified | Displays the date the dashboard was last modified. |
| Modified By | Displays the user who last modified the dashboard. |

You can select more than one dashboard and perform a set of options by clicking the horizontal ellipsis next to the **Add** option.

### Table 14-2. Dashboards Options

| Option | Description | Usage |
| --- | --- | --- |
| Delete | Deletes a dashboard. | |
| Activate | Activates a dashboard that was previously deactivated. | |
| Deactivate | Deactivates a dashboard. | |
| Change Ownership | Assigns a new owner to the dashboard. | After you assign a dashboard to a new owner, the dashboard is no longer displayed as one of your dashboards.<br>When you transfer a dashboard that was previously shared with user groups, information about the shared user groups and group hierarchy is retained. |
| Export | When you export a dashboard, VMware Aria Operations creates a dashboard file in JSON format. | You can export a dashboard from one VMware Aria Operations instance and import it to another.<br>To export a dashboard, select the dashboard that you want to export, and click **Export** from the horizontal ellipsis. |
| Import | A PAK or JSON file that contains dashboard information from VMware Aria Operations. | You can import a dashboard that was exported from another VMware Aria Operations instance.<br>To import a dashboard:<br>1 Click the **Import** option from horizontal ellipsis.<br>2 Click **Browse** and select a Dashboard ZIP, PAK, or JSON file to import.<br>3 Select if you want to **Overwrite** or **Rename** the file in case of a conflict.<br>4 Click **Import** to import the dashboard, and click **Done**. |
| Auto-rotate Dashboards | Changes the order of the dashboard tabs on VMware Aria Operations home page. | You can configure VMware Aria Operations to switch from one dashboard to another. For more information, see Auto-Rotate Dashboards. |
| Manage Summary Dashboards | Provides you with an overview of the state of the selected object, group, or application. | You can change the **Summary** tab with a dashboard to get information specific to your needs. For more information, see Manage Summary Dashboards |

Table 14-2. Dashboards Options (continued)

| Option | Description | Usage |
| --- | --- | --- |
| Manage Dashboard Folders | Groups dashboards in folders. | You can create dashboard folders to group the dashboards in a way that is meaningful to you. For more information, see Manage Dashboard Folders. |
| Manage Dashboard Sharing | Makes a dashboard available to other users or user groups. | You can share a dashboard or dashboard template with one or more user groups. For more information, see Share Dashboards with Users. |

The dashboard list depends on your access rights.

## Manage Summary Dashboards

The **Summary** tab provides you with an overview of the state of the selected object, group, or application. You can change the **Summary** tab with a dashboard to get information specific to your needs.

### Where You Configure a Summary Tab Dashboard

To manage the summary dashboards, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and select **Manage Summary Dashboards**.

### How You Manage the Summary Dashboards

Table 14-3. Manage Summary Dashboards Toolbar Options

| Option | Description |
| --- | --- |
| Use Default | Click to use VMware Aria Operations default **Summary** tab. |
| Assign a Dashboard | Click to view the Dashboard List dialog box that lists all the available dashboards. |
| Adapter Type | Adapter type for which you configure a summary dashboard. |
| Filter | Use a word search to limit the number of adapter types that appear in the list. |

To change the Summary tab for an object, select the object in the left panel, click the **Assign a Dashboard** icon. Select a dashboard for it from the All Dashboards dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard that you have associated to the object type when you navigate to the **Summary** tab of the object details page.

## Auto-Rotate Dashboards

You can change the order of the dashboard tabs on your home page. You can configure VMware Aria Operations to switch from one dashboard to another. This feature is useful if you have

several dashboards that show different aspects of your enterprise's performance and you want to look at each dashboard in turn.

### Where You Configure Auto-Rotation of a Dashboard

To reorder and configure a dashboard switch, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Manage**. Select **Auto-rotate Dashboards** from the horizontal ellipsis next to the **Add** option.

### How You Reorder the Dashboards

The list shows the dashboards as they are ordered. Drag the dashboards up and down to change their order on the home page.

### How You Configure an Automatic Dashboard Rotation

1 Double-click a dashboard from the list to configure.

2 From the Rotation drop-down menus, select **On**.

3 Select the time interval in seconds.

4 Select the dashboard to switch and click **Update**.

5 Click **Save** to save your changes.

On the home page, the current dashboard will switch to the dashboard that is defined after the specified time interval.

## Manage Dashboard Folders

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

### Where You Manage Dashboard Folders

To manage the dashboard folders, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Folders**.

### How You Manage the Dashboard Folders

Table 14-4. Manage Dashboard Folders Options

| Option | Description |
| --- | --- |
| Dashboards List | A list with all available dashboards. |
| Folders | A hierarchy tree with all the available group folders. |

To create a dashboard folder, click **New Folder** in the **Folders** pane and enter the name of the folder. If you want to create a folder under another folder, select a parent folder under which you want to create the child folder, then click **New Folder**. To add a dashboard, drag one from the dashboards list to the selected folder in the **Folders** pane.

You can delete folders and/or detach dashboards from a folder, by selecting one or more folders and dashboards from the **Folders** pane and by clicking **Actions > Delete**.

You can rename a folder by selecting a single folder from the **Folders** pane and by clicking **Actions > Rename**.

## Share Dashboards with Users

You can share a dashboard with one or more user groups. When you share a dashboard, it becomes available to all the users in the user group that you select. The dashboard appears the same to all the users who share it. If you edit a shared dashboard, the dashboard changes for all users. Other users can only view a shared dashboard. They cannot change it.

### Where You Share a Dashboard From

To share a dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Sharing**.

Table 14-5. Dashboard Sharing Options

| Option | Description |
| --- | --- |
| All Dashboards | Link to view all the available dashboards that you can share. The dashboards are displayed on the right side in the dashboards list. |
| User Groups | Lists the available user groups that you can share a dashboard with. The list includes the **Everyone** group. |
| Dashboard List | List of shared dashboards with the selected user group or all the available dashboards that you can share, if no user group is selected. |

### Manage Dashboard Sharing

To share a dashboard, navigate to the dashboard in the list of dashboards and drag it to the group to share it with, on the left.

To stop sharing a dashboard with a group, click that group on the left panel, navigate to the dashboard in the right panel, and click **Stop Sharing** above the list.

Table 14-6. Datagrid Options

| Column Name | Description |
| --- | --- |
| Name | Displays the name of the dashboard. |
| Folder | Displays the folder where the dashboard is located. |
| Description | Displays the description of the dashboard. |
| Activated | Displays whether the dashboard is activated for viewing. |
| Shared | Displays if the dashboard is shared with user groups. |

Table 14-6. Datagrid Options (continued)

| Column Name | Description |
| --- | --- |
| Editable | Click the check box to share dashboard edit privileges with the specific group. |
| Owner | Displays the owner of the dashboard. |
| Last Modified | Displays the date the dashboard was last modified. |
| Modified By | Displays who modified the dashboard. |

# Dashboards Actions and Options

You can change the order of the dashboard tabs, configure VMware Aria Operations to switch from one dashboard to another, create dashboard folders to group the dashboards in a way that is meaningful to you, share a dashboard or dashboard template with one or more user groups, and transfer selected dashboards to a new owner.

## Options for Sharing Dashboards

You can share predefined or custom dashboards by assigning and unassigning a dashboard to specific user groups and by exporting the dashboard configuration details.

Dashboard sharing can only be applied to Groups with a Standard Edition license.

### Where You Can Access the Options to Share Dashboards

From the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click on an existing dashboard and then click the **Share Dashboard** icon in the top-right corner.

Table 14-7. Options in the Share Dashboard Dialog Box

| Option | Description |
| --- | --- |
| Groups | Allows you to assign and unassign a dashboard to specific user groups.<br>■ Select the group to which you want to grant dashboard access from the drop-down menu.<br>■ To share dashboard edit privileges with the group, select the **Allow Editing** check box if a dashboard is shared with a group.<br>■ Click **Include**. You can include more than one group.<br>■ From the label, select the cross mark to unassign the dashboard.<br>Log in to VMware Aria Operations to view the shared dashboard. |
| Export | Allows you to export the dashboard configuration details.<br>Log in to VMware Aria Operations to export/import a dashboard. |

## Manage Widgets in Dashboards

You can replicate widgets multiple times in a dashboard by using the copy and paste functionality.

Navigate to the dashboard from which you want to copy widgets. Select **Actions > Edit Dashboard**. Select one or more widgets that you want to copy by clicking the title of the widget and then select **Actions > Copy Widget(s)**. Click **Actions > Paste Widget(s)** to paste one or more widgets in the same dashboard.

To paste one or more widgets into another dashboard, exit the edit screen of the dashboard by selecting **Cancel**. Navigate to the dashboard to which you want to paste one or more widgets and select **Actions > Edit Dashboard** and then **Actions > Paste Widget(s)**.

# Predefined Dashboards

15

VMware Aria Operations includes a broad set of simple to use, but customizable dashboards to get you started with monitoring your VMware environment. The predefined dashboards address several key questions including how you can troubleshoot your VMs, the workload distribution of your hosts, clusters, and datastores, the capacity of your data center, and information about the VMs. You can also view log details.

Each set of dashboards is complemented with a series of out-of-the-box customizable alerts and reports to assist with your operational awareness. Alerts, reports, and dashboards, each have a purpose with minimal overlap. Several activities that are carried out using alerts should be carried out using dashboards. Reports should be kept to a minimum as they are not interactive and do not provide timely information.

The following table details how alerts, dashboards, and reports are complimentary.

| | Alerts | Dashboards | Reports |
|---|---|---|---|
| Nature | Reactive | Proactive | Passive.<br>For those with no access to<br>VMware Aria Operations<br>and VMware Aria Operations for Logs. |
| Suitability | Exception<br>(something went wrong) | Exception<br>Big Picture<br>Details Analysis | Big Picture<br>Exception (but not urgent)<br>No analysis as it is not interactive |
| Use Case | Troubleshooting (the start) | Monitoring<br>Troubleshooting (the actual) | FYI (optional)<br>Export for further analysis (spreadsheet) |
| Time & Urgency | Urgent (minutes) &<br>Important | Regular (daily, SOP) | Not urgent (monthly) & optional<br>No point in a daily report.<br>For daily reports, login for<br>interactivity |
| Access Requirement | Online. Desktop.<br>1280 * 1024 pixel | Online. Desktop.<br>280 * 1024 pixel | Offline or mobile. Small resolution.<br>Email. Laptop or tablet. |
| Scope/Area | Availability<br>Performance<br>Compliance<br>Configuration (?)<br>Capacity (less relevant,<br>unless it is an emergency) | Availability<br>Performance<br>Capacity<br>Compliance<br>Configuration<br>Inventory | Same as dashboards, but:<br>• without interactivity<br>• time bound (e.g. calendar month)<br>• No performance report, covered in Capacity |
| Roles | Operations Team | Operations Team<br>Architect Team | IT Management (not hands-on)<br>Auditor (compliance) |

### Insight vs Alerts

VMware Aria Operations dashboards support a concept we call insight. Insight complements alerts but does not replace it. Alerts miss the larger picture and only see what is triggered. For one object that reaches the threshold, there might be many just beneath the threshold. The objects below the threshold are called insight.

Alerts might auto-close if the symptoms disappear. Managing alerts is not the same as minimizing alerts. Minimizing alerts is about preventing alerts.

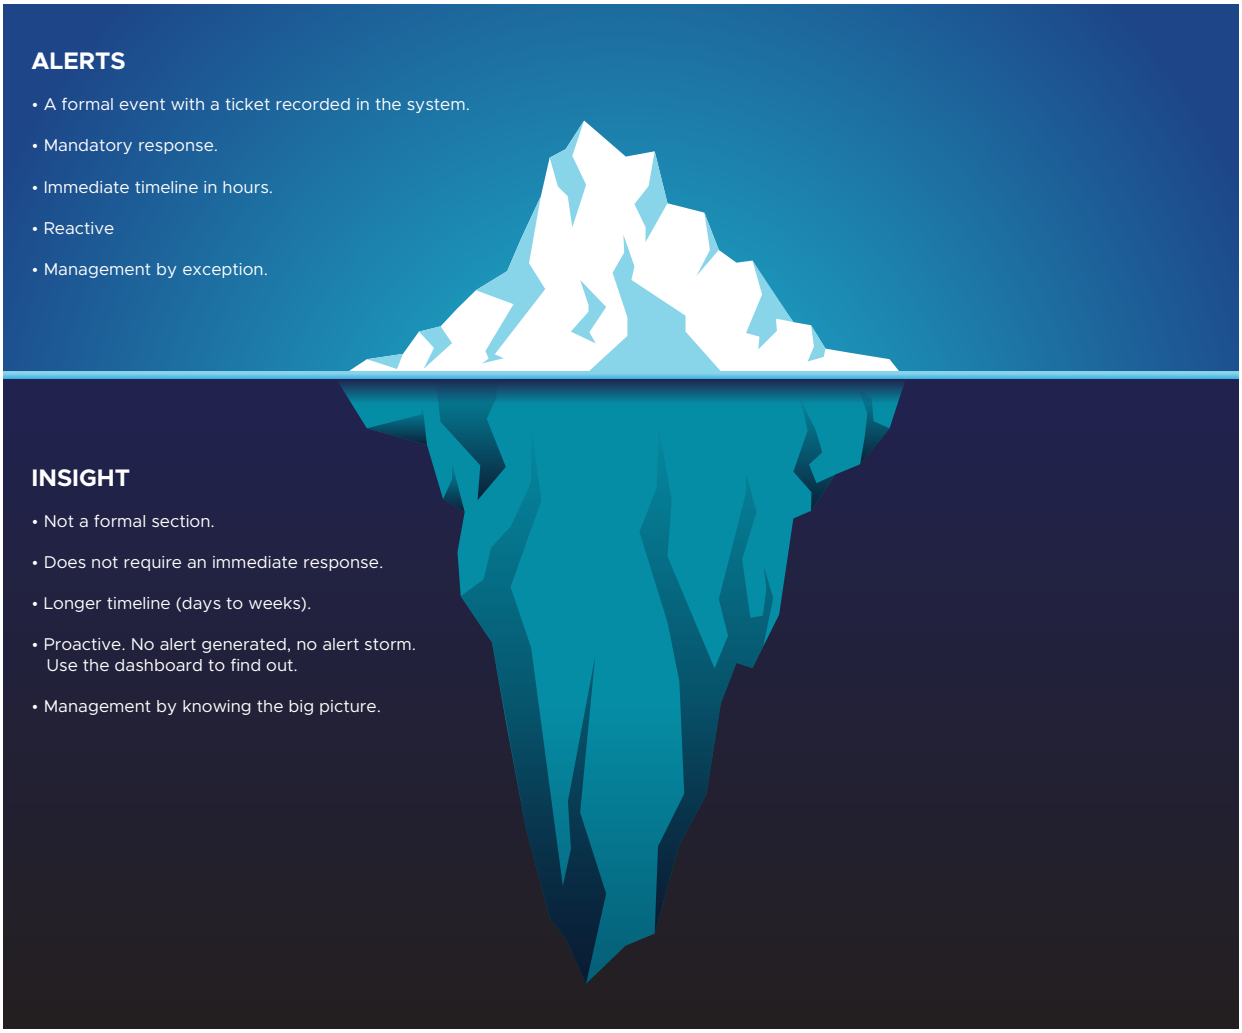


### Working with Predefined Dashboards

The default dashboard that appears when you click **Dashboards** from the left menu is the **Getting Started** dashboard. You can close a dashboard from the left pane by selecting the dashboard and clicking the **X** icon. The dashboard you last opened is displayed the next time you navigate to **Dashboards** in the menu. If there is only one dashboard left in the left pane, you cannot close it.

To access the predefined dashboards, from the left menu, click **Visualize > Dashboards**.

To access the deprecated dashboards, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select **All > Deprecated**.

You can customize dashboards and widgets if you have VMware Aria Operations Advanced edition or higher. Any customization you make is overwritten during upgrade and as a result, it is recommended that you back up your dashboards before an upgrade.

Read the following topics next:

- Application Monitoring Dashboards

- Availability Dashboards
- Capacity Dashboards
- Configuration Dashboards
- Cost Dashboards
- Performance Dashboards
- Sustainability Dashboards
- Dashboard Library
- Software Defined Wide Area Network Dashboard
- VMware Aria Automation Dashboards
- VMware Aria Operations Dashboards
- Service Discovery Dashboards
- Inventory Dashboards
- Microsoft Azure Dashboards
- AWS Dashboards
- Dashboards in VMware Cloud on AWS
- Google Cloud VMware Engine Dashboards
- Dashboards in NSX Management Pack
- Cloud Federation Adapter Dashboards
- Dashboards in VMware Infrastructure Health
- Dashboards in Google Cloud Platform


# Application Monitoring Dashboards

Use the **Linux OS discovered by Telegraf** dashboard and the **Windows OS discovered by Telegraf** dashboard to view the Linux and Windows OS objects discovered on all types of machines such as Virtual Machines, Azure Virtual Machine, EC2 Instance, and physical servers.

## Linux OS discovered by Telegraf Dashboard

Use the **Linux OS discovered by Telegraf** dashboard to view detailed information related to Linux OS objects. You can monitor CPU, disk, and network metric usage. You can also view object relationships.

### Table 15-1. Widgets in the LInux OS discovered by Telegraf Dashboard

| Widget | Description |
|---|---|
| Linux Operating Systems | Displays a list Linux OS objects that are discovered on all types of machines. Click on the Linux OS object for which you want to view details. The other widgets are populated with relevant details. |
| Process | Displays Linux Process metrics for the selected Linux OS object. |
| CPU | Displays CPU metrics for the selected Linux OS object. |
| Memory | Displays Memory metrics for the selected Linux OS object. |
| Memory Swap File | Displays Memory Swap File metrics for the selected OS object. |
| CPU: Individual CPU Performance | Displays all CPU usage metrics for the selected Linux OS object. |
| Disk: Individual Partition Capacity | Displays the instance name, total capacity, free capacity, and percentage of capacity that is used. |
| Disk: Individual Disk Performance | Displays the Individual disk performance metrics for the selected Linux OS object. |
| Related Objects | Displays a detailed object relationship of the Linux OS objects. |

## Windows OS discovered by Telegraf Dashboard

Use the **Windows OS discovered by Telegraf** dashboard to view detailed information related to Windows OS objects. You can monitor CPU, disk, and network metric usage. You can also view object relationships.

### Table 15-2. Widgets in the Windows OS discovered by Telegraf Dashboard

| Widget | Description |
|---|---|
| Windows Operating Systems | Displays a list Windows OS objects that are discovered on all types of machines. Click on the Windows OS object for which you want to view details. The other widgets are populated with relevant details. |
| Associated VM | Displays the VM version, CPU, memory, and disk space details. |
| System | Displays System metrics for the selected Windows OS object. |
| CPU | Displays CPU metrics for the selected Windows OS object. |
| Memory | Displays Memory metrics for the selected Windows OS object. |
| Windows: Memory Rate | Displays Memory Rate metrics for the selected Windows OS object. |

Table 15-2. Widgets in the Windows OS discovered by Telegraf Dashboard (continued)

| Widget | Description |
| --- | --- |
| CPU: Individual CPU Usage | Displays all CPU usage metrics for the selected Windows OS object. |
| Disk: Individual Partition Capacity | Displays the instance name, total capacity, free capacity and percentage of capacity that is used. |
| Network: Individual Adapter Usage | Displays the individual network adapter usage metrics for the selected Windows OS object. |
| Related Objects | Displays a detailed object relationship of the Windows OS objects. |

# Availability Dashboards

Availability covers the uptime of the object now and the uptime trend over time. The availability of hybrid clouds should be tracked at both the provider and consumer layers to understand the availability of the environment. These dashboards show the current uptime and the uptime percentage over the past month.

## VM Availability Dashboard

Use the **VM Availability** dashboard to calculate the availability of the Guest OS. The availability of the Guest OS is calculated because the Guest OS might not be running even when the VM is powered on. There are two layers of Availability, that is, the Consumer layer and the Provider layer. This dashboard covers the Consumer layer. You can view VMs in the selected data center, uptime trend for a selected cluster, and so on.

### Design Considerations

The **VM Availability** dashboard helps you check the availability (uptime in percentage) of VMs, as availability is typically part of the services provided by the IaaS provider.

This dashboard does not check the application uptime because it is possible that the application such as, a database, or a web server, is down while the underlying Windows or Linux is up. Generally, the service provided by the IaaS team is only for Windows or Linux. For information on the application, use the network ping or application-specific agent such as application monitoring.

### How to Use the Dashboard

- In the **Datacenters** widget, click any data center from the list.
  - To view the overall information, click the **vSphere World** object.
  - The other widgets are automatically updated once you click any data center.

  - Create a filter that reflects your class of service for this widget. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads and you can focus on the important VMs. You can achieve this by creating a VMware Aria Operations custom group for each class of service.

- The **VMs by Uptime in the last 30 days** widget displays the average uptime of VMs grouped by their availability. The bucket distribution helps you cater to a wide array of environments. If you are monitoring only production VMs where the uptime is expected to be near 100% all the time, edit the bucket to meet your operational needs.
  - The VMs in the **Selected Datacenter** widget display all the VMs that are currently deployed to the data center. The average uptime is displayed for the last month. For a production VM, expect this number to be 100% or closer to 100%.

    **Note** The Services column will be blank unless Service Discovery is activated and the services/processes are discovered on a specific virtual machine.

  - The VMs column includes all VMs including the powered off VMs.
- Click any VM in the **VMs by Uptime in the last 30 days** widget to view the details in the **VM in the Selected VM Powered On Status**, **Selected VM Uptime Trend**, and **Selected Cluster Uptime Trend** widgets.
  - The **Selected VM Uptime Trend** widget displays the selected VM's Guest Tool Uptime (%) across the last 30 days.
- The **Guest OS: Services** widget displays the service state over time and the process or services running inside the Guest OS. If Guest OS services or processes are discovered inside a VM, their availability is analyzed. This requires the Service Discovery.
- The **ESXi Host(s) where the VM has run** widget displays the historical migration of the VM. This can be useful in determining the cause of a VM downtime.

## Points to Note

- The metric only tracks the availability of VMware Tools and not the entire Guest OS. If VMware Tools is not up, it assumes the Guest OS to be down. You can check that this is not a false negative by adding a few line charts that display the evidence of activity. A good counter is IO counters such as Disk IOPS, Disk Throughput, and Network Transmit Throughput, because IO requires CPU processing. CPU usage is not a reliable counter as the work by VMkernel on the VM is charged to the CPU counters.
- VMware Aria Operations exhibits a new ping adapter. This allows you to enhance the accuracy of the uptime measurement by creating a super metric that adds the ping information or by checking the process using an agent, such as application monitoring.
- Add a property widget that lists the selected VM properties to give you more context about the VM. In a large environment, the VM name alone might not provide enough context.

# vSphere Availability Dashboard

There are two layers of Availability, that is, the Consumer layer and the Provider layer. The **vSphere Availability** dashboard covers the Provider layer. This dashboard includes a cluster and not an ESXi host because the cluster is operationally a single compute provider. This dashboard considers the N+1 design, where the cluster can withstand one host failure. Logically, a cluster with fewer hosts has a higher risk.

## Design Considerations

The **vSphere Availability** dashboard helps you analyze and report the uptime, as availability is typically part of the official business SLA. It is also often required in the monthly operational summary report.

This dashboard is not designed for live monitoring of the uptime. A NOC style of dashboard is better suited for those use cases. VMware Tools such as VMware Aria Operations for Logs must be leveraged as the fault is typically preceded with soft errors.

## How to Use the Dashboard

- The **Clusters** widget lists all the clusters in the environment. It is sorted by the lowest uptime so that the cluster with the lowest uptime in the last one month is displayed.
  - The **Running Hosts** column is color-coded as logically a smaller cluster has a higher risk. A single host failure results in a relatively higher capacity degradation.
  - The **vSAN?** column is hyper-converged, which means both the compute and the storage part is considered.
  - The **Admission Control Policy** column is based on the Cluster Configuration \ DAS Configuration \ Active property. The mapping between the code to name is:
    - -1 : Disabled
    - 0 : Cluster Resource Percentage
    - 1 : Slot Policy (Powered-on VMs)
    - 2 : Dedicated Failover Hosts
  - In a large environment, creating a filter for the list of clusters can make it more manageable. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, you can easily view your gold clusters.
- Click any cluster from the **Clusters** widget.
  - The cluster uptime is automatically plotted in the **Selected Cluster Uptime Trend** widget. It uses 99%, 99.%, and 99.99% as the threshold for red, orange, and yellow colors respectively.
  - The ESXi host details in **ESXi in the Selected Cluster** widget are automatically updated. For more context, you can add a property widget that lists the selected ESXi host properties.

  - In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.
- The **Datastores not available** widget lists only the datastores with powered off status. This covers both local and shared datastores. To add context, consider adding an extra column such as the data center where it resides, and the datastore types such as NFS and VMFS.
- The **Port Group Availability** widget lists port groups that currently have an uptime of less than 100%. To add context, consider adding an extra column such as the data center where it resides, number of used ports, and the maximum number of ports.
- For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.
- In a large environment, you can create a filter for this dashboard. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads.
- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.

### Points to Note

- You can add vCenter Server and NSX components availability. This requires the VMware SDDC Health Monitoring Solution.

## Ping Overview Dashboard

Use the Ping Overview dashboard to configure the ping functionality and verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN. You can view ping adapter details like, latency distribution and packet loss distribution in this dashboard.

### Customizations Available for Your Use

For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.

**Note** The FQDN names are checked for validity, the FQDN validation relies on RFC1034 and RFC1123, and only top level domains of the internet are validated. The `.local` domain is not supported as it does not fall into the list of top-level domains in the Domain Name System (DNS) of the Internet.

### Widget Information

- Latency distribution - You can use this widget to see the objects that are experiencing high latency.
- Packet Loss Distribution – You can use this widget to see the objects that are experiencing high packet loss.
- Ping Targets – You can use this widget to view the list of ping targets grouped by their FQDN. Latency and packet loss information is also displayed for the ping objects.
- Breakdown by Source Initiator – You can use this widget to view the List of ping statistics by the source (ping initiator). You can ping the target from multiple locations, to determine if the issue is network-related or server-related.

## Capacity Dashboards

Capacity quantifies the resources used, resources remaining, and opportunities to reclaim unused resources. Projections of the demand provide a proactive view of capacity. The **Capacity Dashboards** display capacity in terms of time remaining before capacity is projected to run out, the amount of capacity remaining, the number of VMs that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

Capacity management is about balancing demand and supply. It is about meeting demand with the lowest possible cost.

For IaaS or DaaS, capacity management begins before the hardware is deployed. It begins with a business plan that defines what class of service will be provided. Each class of service, for example, gold, silver, bronze is differentiated by the quality of service and covers the availability, for example, 99.99% uptime for Gold, 99.95% uptime for Silver. It also covers performance, for example, 10 ms disk latency for Gold, 20 ms disk latency for Silver, and security or compliance.

The quality incurs cost and in turn drives price. Gold VM is higher per vCPU and per GB of RAM because it has a higher quality of service. A proper pricing model must be planned. If you want your customers to rightsize in advance, then a 64 vCPU VM has to be more than 64x the price of 1 vCPU VM. If the pricing model is a simple straight line, there is no incentive to go small and no penalty if it is over provisioned. In this case, you end up forcing rightsizing in production, which is a costly and time consuming process.

Demand is more than the active load that is consuming your capacity. Since capacity based on utilization is incomplete by itself, the principles displayed in the following figure are considered.

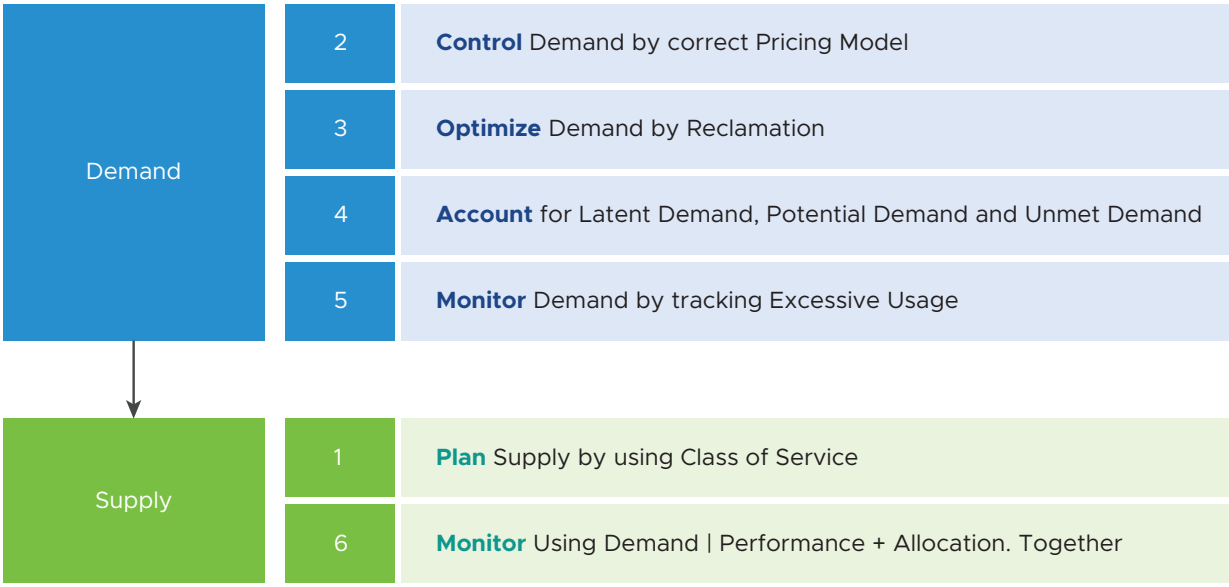


- Latent demand. Many critical VMs are protected with Disaster Recovery. During a Disaster Recovery drill or an actual disaster, this load is consumed.
- Potential demand. Many newly provisioned VMs take time to reach their expected demand. It takes time for the database to reach the full size, the user base to reach the target, and the functionalities to complete. When this is achieved, it results in the increase in demand.
- Unmet demand occurs when the VM or Kubernetes Pod is undersized. The load is running nearly 100% most of the time.
- Excessive demand can wreak havoc in a shared environment. A group of highly demanding VM can collectively impact overall performance of the cluster or datastore.

## Cluster Capacity Dashboard

The **Cluster Capacity** dashboard includes the ESXi host and resource pools as they impact cluster capacity.

### Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

### How to Use the Dashboard

The **Cluster Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

### Overall Analysis

- The three bar charts which are **Clusters by Capacity Remaining**, **Clusters by Time Remaining**, **Clusters by VM Remaining**, summarize the overall situation. The first two charts can be used together to identify when you need to add capacity to address growth. Time remaining uses historical growth in a cluster to forecast when more capacity is needed. This allows you to operate more efficiently by making sure you have enough capacity currently and proactively plan for adding capacity. The third bar chart which is **Clusters by VM Remaining**, provides complete contexts, as different clusters can have different VM sizes.

  For a large environment, a heat map is helpful. The three heat maps are Time Remaining, Capacity Remaining, and VM Remaining. If your cluster sizes are not standardized, create another heat map, and use the number of ESXi hosts to show the size difference.

### Cluster Analysis

- The **Clusters Capacity** widget provides a table with details. The number of ESXi hosts are color coded as smaller clusters have a relatively higher overhead. Select a cluster from the table to view the capacity details that are automatically displayed.

  #### Performance

  Ensure that the performance of the cluster meets your SLAs.

  #### Utilization

  The next two charts are Memory Workload (%) and CPU Workload (%), that show values relative to your usable capacity. Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average, so you can focus on the overall trend. For memory, the focus is on consumed memory and not active memory.

  #### Allocation

  You can view the trend of the three which are CPU, disk, and memory components together on the **Overcommit Ratio** chart. In general, your CPU overcommit should be the highest, followed by the disk (because of thin provision). Memory overcommit tends to be near one due to its nature as cache.

  Use the line chart in the **Allocation** widget, to see the trend. The data is averaged hourly.

  In the **VM Count** widget, the trend line of the number of VMs over time is important to spot if there are many newly provisioned VM. If you see that the VMs are increasing but demand remains low, it indicates a sign of potential demand in the future.

  #### Reservation

  Reservation can impact the efficiency of your cluster. Your cluster could be low on capacity because of real workload or just reservation. If your cluster size varies, complement the reservation number by showing a relative value. Once you have a standardized number, you can visualize them on a heat map.

- **ESXi Analysis**

Good cluster capacity does not indicate that there is no issue at the ESXi level. Unbalance is a common problem, especially in large clusters and stretched clusters.

The **ESXi Hosts in a Cluster** table displays all the member ESXi hosts. You can see the unbalance clearly, thanks to the color code. The color code reflects the unbalance.

The **99th percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric.

Select an ESXI host to view the details. Both the **CPU Workload (%)** and the **Memory Workload (%)** trend line charts display if there is a steady demand, cyclical demand, rising demand, or declining demand. The trend is as important as the present value. View trends over a longer time. Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average. The focus is on memory consumed and not memory active. Memory consumed includes the total memory consumed, so it includes the memory consumed by VMkernel. Both total and usable utilization in terms of memory and CPU are displayed and provides the absolute amount of capacity.

- **VM Analysis**

  Use the **VMs in the selected Cluster or Host** table to analyze the cause of the low capacity remaining and which VMs are impacting the infrastructure resources, such as, CPU, memory, and disk space. The table lists either the VMs in the cluster or host. When you select one of the VMs, additional relevant information is displayed.

  If there are many large VMs running low on capacity you can stop provisioning until you upsize the existing VMs first.

## Datastore Capacity Dashboard

The **Datastore Capacity** dashboard complements out of the box capacity pages and dashboards. It focuses on storage, provides an overall picture, and highlights the datastores that need attention.

### Design Considerations

See Capacity Dashboards for common design consideration among all the dashboards for capacity management.

### How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

- **Overall Analysis**

  The summary banner answer basic questions such as number of datastores, the capacity, number of VMs, and the running VMs.

  The distribution charts are **Shared Datastores by Capacity Remaining** and **Shared Datastores by Time Remaining**.

There are three heat maps, the primary being the **Remaining Capacity** heat map. The two other heat maps cover used capacity. One of them is designed for an environment that uses datastore clusters. Each box represents a datastore. If you have many datastores, the heat map will group them. You can further see the members. The larger the datastore, the larger the box.

The **Shared Datastores** table lists the shared datastores. The table provides a summary, displaying all the datastores at a glance. They are grouped by data center. By default, the table is sorted by the least capacity remaining. There are three reclamation opportunities: powered off VM, snapshot, and orphaned VMDK.

- **Datastore Analysis**

  Select a datastore from the **Summary** table. The capacity details are automatically displayed. A snapshot that lasts beyond a few days should be investigated. Orphaned VMDK are the ones that are not associated to any VM.

  For disk space, the total capacity, allocated capacity, and the actual capacity used are displayed.

- **VM Analysis**

  To analyse at the VM level, review the **VMs in the datastore** table. Click on the VM you want to investigate further to see usage over time.

- **Local Datastores**

  The **Local Datastores Capacity** table appears at the end of the dashboard. Avoid running VMs on local data stores, unless the storage requirements can be met with a local disk and does not need vMotion.

### Points to Note

If the underlying LUN is also thin provisioned, add visibility into the physical array. The dashboard does not have datastore clusters. If your environment uses datastore clusters, modify this dashboard or create a new one. In a large environment with many datastores and datastore clusters, add a View List to list the datastore clusters, so you get summary information. Alternatively, create a heat map, listing the datastore clusters.

## ESXi Capacity Dashboard

The **ESXi Capacity** dashboard supports the **Cluster Capacity** dashboard and is also required for the non-clustered ESXi.

### Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

### How to Use the Dashboard

The **Summary** heat map provides an overall view of the ESXi Host capacity, grouped by their clusters.

- Each ESXi host is represented by a box, displaying their capacity remaining.
- The ESXi host size is made constant for ease of use. If your ESXi sizes are not standardized, consider using the number of physical cores or Total CPU GHz to display the difference in sizes. Ensure that the smallest ESXi is not too small.
- Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

The **ESXi Hosts Capacity** widget lists all the ESXi hosts in your environment, grouped by their parent cluster.

- The standalone ESXi are displayed at the bottom under No Group.
- In a large environment with many data centers, you can zoom into a specific vCenter or data center. You can also filter or search for specific ESXi hosts matching certain names.
- The **99th Percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric. To rule out the outlier, the worst performance (which is equivalent to 100th percentile) is not considered. Also, the performance threshold is set to be stringent.

Select one of the ESXi hosts from the **ESXi Hosts Capacity** widget. All the three line charts automatically display the trend of selected ESXi host.

- Displays both total and usable utilization in terms of RAM and CPU.
- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.

### Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.
- A technology refresh is often used to address the capacity shortage. Consider adding a property widget that displays the hardware model and specification to help you determine the age of the hardware.

## VM Capacity Dashboard

The **VM Capacity** dashboard helps you analyze the capacity of all the VMs with the ability to analyze each VM.

### Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

### How to Use the Dashboard

- **Overall Analysis**

  The **Datacenters** table lists all your data centers. vSphere World is included so you can see all the VMs from all the data centers. Unlike infrastructure objects, there are potentially tens of thousands of VMs. The charts will take longer to refresh if you select vSphere World.

  The **VMs by Capacity** widget, groups the VMs by capacity remaining, while the **VMs by Time Remaining** widget groups the VMs by time remaining. Ideally, you want all of them to be low on capacity remaining, but high on time remaining.

  The heat map provides an additional view by grouping them by cluster. The heat map helps you identify which cluster is at risk (where most of the VMs need more capacity) and which cluster provides extra resources (where most of the VMs are not using their capacity).

  Review the **VMs by Capacity Remaining** heat map. The heat map provides the next level of detail by grouping the VMs by clusters, so you can see which clusters need attention. The VM size is standardized for better visualization.

- **VM Analysis**

  Review the **VMs Capacity in the Selected Datacenter** table that lists all the VMs in the selected data center. The list is sorted by the VM with the least capacity remaining. You can also sort by Time Remaining.

  Select a VM from the table. The capacity details are displayed. You can view both the CPU and memory trend over time. Three months data is displayed and is averaged to hourly so you can see the overall trend.

  Disk details are displayed using the Guest OS partition. Avoid using VM virtual disk as there may not be a direct mapping to the actual partition.

  The recommendation for right-sizing is displayed for both CPU and memory. Unlike a physical server, it is important to right-size a VM. For CPU, the CPU Usage counter is used instead of Demand. For disk, Guest OS partition level is displayed. There is no overall capacity at the VM level because different partitions have different capacity.

- **Relevant Configuration**

  Relevant configuration is automatically displayed to provide context to the VM. Information such as VM owner and business units can be useful in the analysis.

## Reclamation Dashboard

The **Reclamation** dashboard helps you manage various types of reclamation that can be carried out on VMs and datastore. It is designed for both the Capacity team and the Operations team.

### How to Use the Dashboard

- **Overall Analysis**

  The scoreboard provides a summary of the total reclamation.

You can select either a data center, a cluster, or a datastore. The datastore level is required as orphaned disks do not have VM association, hence it is not related to any cluster. The **Datastore** table only drives the snapshot table.

**Powered-off VMs Distribution Size**, **Idle VMs Distribution by Memory Footprint**, and **Snapshots Distribution by Size** charts display summary information.

Adjust the bucket size in the charts to suit your operational requirements. The reclamation potentials are presented as three bar charts, each corresponds to an area you can reclaim:

- Powered off VMs that are no longer needed contribute to wasted disk usage. Consider deleting them to free up space or moving them to archival storage.
- Idle VMs are running, but not being used actively. These VMs consume memory that may be used by active VMs. Consider removing these VMs to reduce memory contention.
- Snapshots are meant to be temporary and can cause performance issues and waste disk space if not deleted after a few days.

When reviewing each of the three tables observe that they are sorted by the largest reclamation opportunities. This allows you to get the greatest benefit from the least amount of effort. For example, focus on snapshots first, as it does not involve changing the VM. For the powered off VMs, you may want to consider VMs that have been powered off longer and these are more likely to be unneeded. Idle VMs can be challenging to reclaim since they are still running, so you might prioritize powered off VMs before attempting to reclaim idle VMs.

The **99P CPU Usage** column displays the CPU usage at the 99th percentile during the time period. It is an easy way to check if it is indeed idle.

- **VM Analysis**

  To analyze VMs for reclamation opportunities, select a VM from one of the three tables (Powered Off VM, Idle VMs, or VM Snapshots). The selected VM will populate the widgets with the following details:

  - Powered off over time displays the amount of time a VM has been powered off.
  - CPU Usage over time provides insights into the aggregate CPU usage, including peak usage periods. This way you can validate that an idle VM has not had any brief usage.
  - VM Snapshot over time provides you with an understanding of the age and growth of snapshots on the VM. Watch for fast growing snapshots, as these can quickly consume disk space.
  - Contexts of selected VM is a summary of the VM configuration information.

## Points to Note

To organize your reclaim efforts, it is helpful to create custom groups to make it easier to filter by department or VM owner. This can make it easier to seek approvals and communicate with anyone who may be impacted.

# vSAN Capacity Dashboard

The **vSAN Capacity** dashboard complements the vSphere **Cluster Capacity** dashboard by displaying capacity related to vSAN. To manage vSAN capacity, use both dashboards.

## Design Considerations

The dashboard focuses on vSAN specific metrics, but does not list non-vSAN clusters.

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

## How to Use the Dashboard

- **Overall Analysis**

  The **Clusters by Capacity Remaining** and **Clusters by Time Remaining** bar charts focus on vSAN disk space and not compute and network.

  Use the **vSAN Clusters** table to view vSAN specific metrics.

- **Cluster Analysis**

  Select a vSAN cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  Like a physical array, there can be hot spots and imbalance. The **Is Disk Group Space Utilization Balanced** heat map displays individual disk groups.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.

  You can view the **Dedupe and Compressed** scoreboard for more details in this area.

- **Disk Group Analysis**

  If there is imbalance, you can analyze each disk group. The **Disk Groups Selected in Selected vSAN Cluster** displays all the disk groups in the cluster. Their usage may not be similar, but should not deviate drastically. To view the disk group usage trend, click on a disk group.

- **VM Analysis**

  You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

# vSAN Stretched Clusters

The vSAN Stretched Clusters dashboard provides an overview of the cluster resources used across vSAN fault domains. Using the stretched clusters dashboard you can monitor the resource

consumption at the site level for Preferred Sites and Secondary Sites. You can create custom dashboards for specific vSAN stretched cluster metrics.

### Where to View vSAN Stretched Cluster Objects

On the menu, click **Dashboard > Capacity and Utilization > vSAN Stretched Clusters**.

You can also view the vSAN stretched cluster objects from **Environment > VMware vSAN > vSAN and Storage Devices > vSAN Clusters**, if the vSAN cluster is a stretched cluster.

The vSAN Stretched Clusters dashboard provides information about CPU Capacity, Cores, Memory Capacity, and Disk Capacity for the Preferred Site and the Secondary Site. You can identify the vSAN stretched clusters running out of capacity looking at the utilization metrics.

## vSAN ESA Capacity Dashboard

The vSAN ESA ( Express Architecture) Capacity dashboard provides an overview of the capacity available across all the vSAN clusters. The dashboard displays the time remaining before capacity is projected to run out, the amount of capacity remaining, the number of vSAN clusters that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

### How to Use the Dashboard

- **Overall Analysis**

  The **All Clusters by Capacity Remaining** and **All Clusters by Time Remaining** bar charts focuses on the number of vSAN disk against the used disk capacity percentage across vSAN clusters.

- **Cluster Analysis**

  Select a vSAN ESA cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  The **Is Storage Pool Space Utilization Balanced** heat map displays individual storage pools.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.

- The **Storage Pools in Selected vSAN Cluster** widget displays the utilization for Disks and Disk Space.
- Like a physical array, there can be hot spots and imbalance. The **Is Physical Disk Space Utilization Balanced** heat map displays individual physical disks.
- **VM Analysis**

You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

## vSAN OSA Capacity Dashboard

The vSAN OSA (Original Storage Architecture) Capacity dashboard provides an overview of the capacity available across all the vSAN clusters. The dashboard displays the time remaining before capacity is projected to run out, the amount of capacity remaining, the number of vSAN clusters that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

### How to Use the Dashboard

- **Overall Analysis**

  The **All Clusters by Capacity Remaining** and **All Clusters by Time Remaining** bar charts focuses on the number of vSAN disk against the used disk capacity percentage across vSAN clusters.

- **Cluster Analysis**

  Select a vSAN OSA cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  Like a physical array, there can be hot spots and imbalance. The **Is Disk Group Space Utilization Balanced** heat map displays individual disk groups.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.

- **Disk Group Analysis**

  If there is imbalance, you can analyze each disk group. The **Disk Groups Selected in Selected vSAN Cluster** displays all the disk groups in the cluster. Their usage may not be similar, but should not deviate drastically. To view the disk group usage trend, click on a disk group.

- **VM Analysis**

  You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

# Configuration Dashboards

As an operations management software, VMware Aria Operations focuses on the impact to day-to-day operations a product has, rather than the feature of the product itself. Products under

monitoring, such as vSphere and vSAN, can have features that are related, but have a different impact on operations. For example, vSphere provides Limits, Reservation, and Shares for the VM.

Limits, Shares, and Reservation. As a feature, they are closely related, and appear in the same dialog box and must be learned as one. However, they impact operations differently. The following table describes that in more detail.

| | | | |
|---|---|---|---|
| VM Limit | Impacts the VM | • Should not be used. Right-size instead.<br>• Results in unpredictable performance of the Guest OS. | Incorrect Configuration |
| VM Reservation | Impacts the Infrastructure | • Keep the total amount low, and relative to the total capacity of the cluster.<br>• Absolute value. A 2-GHz reservation is in fact a 2-GHz reservation.<br>• Results in suboptimal infrastructure capacity, as overcommit is not possible. | Suboptimal Configuration |
| VM Share | Impacts the Infrastructure | • Keep the number of variations to below three. One for each class of service.<br>• Relative value. 2000 worth of reservations depends on the value of other VM reservations. Be careful when you move the VM to another cluster, as the relative value changes.<br>• Results in complex operations. It is harder to troubleshoot performance when the dynamic entitlements of each VM fluctates more. | Complex Configuration |

VMware Aria Operations follows the principle that there are different impacts on operations and applies a methodology for looking at configuration. It does not group the settings by features or objects. Rather, it begins with the impact and prioritizes what can be done.

| Correct it? | Update it? | Simplify it? | Optimize it? |
|---|---|---|---|
| Setting is wrong.<br><br>Correct it. | It is correct, but on an older version.<br><br>Upgrade or update it. | It is correct and up to date, but it complicates your IaaS operations.<br><br>Minimize it. | It is correct, up to date, does not complicate your IaaS operations, but it is not as efficient.<br>You can lower cost.<br><br>Optimize it. |
| • Settings that are meant to be enabled are disabled.<br><br>• Insecure settings.<br><br>• Snapshot running for weeks.<br><br>• VMs with Limits. | • Running an older version of VMware Tools, Guest OS, ESXi, VMFS, distributed switch, and Telegraf agents.<br><br>• vSAN: Take the older version. | • Large configuration. Example is a monster VM.<br><br>• Clusters with too many permutations of shares or VM settings, that override cluster-wide HA/DRS settings. | • Small clusters. HA overhead is relatively higher.<br><br>• Clusters with lots of reservations.<br><br>• Powered off ESXi.<br><br>• Under-sized or over-sized VMs have suboptimal capacity, not configuration management. |

Each operation is unique and as a result, customers run operations differently. What is right for other customers, might not be right for you. Even in the same environment, what is right for a development environment might not be appropriate for a production environment.

The following table lists some of the areas for improvement for the operations in your environment:

## Areas of Improvement

| | Correct it? | Update it? | Simplify it? | Optimize it? |
|---|---|---|---|---|
| **IaaS Consumer:**<br>• **Process**<br>• **Applications**<br>• **Guest OS**<br>• **Container**<br>• **VM** | • Java JVM or Database \ memory config too large relative to Guest OS<br>• Guest \ Metric not collecting<br>• Guest \ High TX Broadcast packets<br>• VM \ Tools not installed<br>• VM \ Tools not running<br>• VM \ CPU Limit<br>• VM \ Memory Limit<br>• VM \ Old Snapshot<br>• VM \ On local Datastore | • Guest OS \ Tools<br>• Guest OS \ Windows<br>• Guest OS \ Linux<br>• Guest OS \ Telegraf agent<br>• VM \ Hardware (vmx) | • VM \ Large VM (CPU, RAM, Disk)<br>• VM \ lots of disks, NIC card<br>• VM \ lots of IP address.<br>• VM \ with RDM<br>• VM \ on multiple datastores<br>• VM \ Fault Tolerant<br>• VM \ SRM protected<br>• VM \ Hot Add\|Remove \ CPU<br>• VM \ Hot Add\|Remove \ RAM | • Java JVM or Database \ memory config too small relative to Guest OS<br>• Guest OS \ no visibility<br>• Container \ smaller than the parent VM<br>• VM \ Tools unmanaged<br>• VM \ bigger than the whole ESXi cores.<br>• VM \ bigger than CPU socket.<br>• VM \ Large Snapshot<br>• VM \ Reservation. |
| **IaaS Provider:**<br>• **Telegraf**<br>• **ESXi**<br>• **Cluster**<br>• **Datastore & Cluster**<br>• **Switch and Port Group**<br>• **Hardware**<br>• **NSX**<br>• **vSAN** | • ESXi \ vMotion disabled<br>• ESXi \ Disconnected from vCenter<br>• ESXi \ Maintenance Mode<br>• ESXi \ NTP disabled<br>• ESXi \ Standalone<br>• Cluster \ Admission Control disabled<br>• Cluster \ HA disabled<br>• Cluster \ HA Failover %<br>• Cluster \ DRS disabled<br>• Cluster \ DRS manual<br>• Cluster Inconsistency<br>  • BIOS, ESXi: version<br>  • BIOS, ESXi: Power Management<br>  • ESXi Storage Path<br>  • ESXi Hardware<br>• Datastore Cluster inconsistency<br>  • Capacity<br>  • Performance<br>• Datastore \ single path<br>• Datastore \ no path. This is unlikely.<br>• NSX \ no redundancy for Controller, Manager | • ARC \ server<br>• ARC \ agent<br>• ESXi \ hardware<br>• ESXi \ vSphere<br>• ESXi \ 1 Gb NIC.<br>• Server \ not on warranty<br>• vCenter \ version<br>• Datastore \ VMFS version<br>• vSAN \ version<br>• Switch \ version<br>• NSX \ version | • ESXi \ Too many variations. No standard<br>• Cluster \ Many VM Shares (CPU)<br>• Cluster \ Many VM Shares (RAM)<br>• Cluster \ Resource Pools<br>• Cluster \ Stretched compute + storage<br>• Cluster \ 32 nodes or more<br>• Cluster \ VM to Host affinity<br>• Cluster \ Too many storage paths<br>• Datastore \ Shared by >1 cluster WLP uses this<br>• Datastore \ Many paths<br>• Network \ LBT?<br>• Network \ MAC Address change | • ESXi \ low CPU cores count<br>• ESXi \ low RAM size<br>• ESXi \ Powered Off<br>• ESXi \ HT Disabled<br>• ESXi \ 4 socket or higher.<br>• Cluster \ small clusters \ host especially for vSAN<br>• Cluster \ small clusters \ CPU<br>• Cluster \ small clusters \ RAM<br>• Cluster \ EVC Mode<br>• Cluster \ High Reservation<br>• Cluster \ DRS Automation Level<br>• Cluster \ DPM disabled<br>• vSAN \ All Flash: Dedupe disabled<br>• vSAN \ All Flash: Compressed disabled<br>• Datastore \ small<br>• Datastore \ low VM count<br>• Datastore \ no ESXi<br>• Distributed Switch \ unused |

## Design Considerations

The dashboards display configurations that need immediate attention, before displaying the overall configuration. This helps you take measures toward optimizing configuration.

Operations vary among customers, and as a result, it is not possible to design one dashboard to meet every customer's operational needs. A configuration that is important for one customer might not be relevant for another customer. Tailor the dashboard to your unique environment. You can collapse or expand the widgets to allow relevant data to be displayed.

The overall layout is designed to balance ease of use, performance (loading time of the dashboard page), and completeness of configuration check. As a result, not all configuration settings are displayed. Lack of screen real estate is another consideration behind the design.

## Cluster Configuration Dashboard

Use the **Cluster Configuration** dashboard to view the overall configuration of vSphere clusters in your environment, especially configurations that need attention.

### Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

## How to Use the Dashboard

The **Cluster Configuration** dashboard is organized into sections for ease of use.

- The first section of the dashboard consists of three bar-charts. They correspond to the three main features of vSphere clusters, namely High Availability (HA), Dynamic Resource Scheduler (DRS), and Distributed Power Management (DPM).
  - HA: The best practice is to activate HA admission control. You can specify the admission control policy in the vCenter Server and the threshold for failover shares.
  - DRS: The best practice is to activate DRS. Envision the vSphere cluster as a single logical computer that balances within itself.
  - DPM: The best practice is to activate DPM in an environment where environmental concern is the top priority or the high peak rarely occurs as most of the time you run very low utilization.
- The second section of the dashboard consists of eight pie-charts. They show the relative distribution of key configurations.
  - Two of the bar-charts cover admission control. You must activate admission control. The pie-charts displays the policy code instead of the policy name, as it is based on the property: `Cluster Configuration | Das Configuration | Active Admission Control Policy`. The mapping between the code to name is:
    - -1 = Disabled
    - 0 = Cluster Resource Percentage
    - 1 = Slot Policy (Powered-on VMs)
    - 2 = Dedicated Failover Hosts
  - There are two bar-charts that cover the HA Failover Share. One for CPU, and one for memory.
  - The next two bar-charts cover DRS settings. You might want to fully automate DRS, which means that there is no operator intervention required for both initial VM placement and subsequent load balancing, but with a moderate migration threshold (value = 3.0). The value ranges from 1.0 to 5.0.
  - There are two pie-charts that show reservation. One for CPU and one for memory. Minimize the total reservation value as it prevents overcommit of resources and hence results in less optimal utilization. Memory reservation can remain and occupy the memory space of the ESXi host, even though the VM does not use the memory anymore. Consider the analogy of unused files that you have not opened for months in the c:\ drive of your laptop. They still take up space on the hard disk. Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service.

- The third section of the dashboard consists of two bar-charts. They show the absolute distribution of the clusters.
  - The first bar-chart displays the cluster grouped by the number of ESXi hosts. Small clusters, defined as having a lower number of ESXi hosts, have a higher overhead while large clusters have a higher risk if there are cluster-wide outages. Performance risk is lower, because there are more nodes that DRS can tap on, but if there is an actual problem, troubleshooting can be tougher, because there are more nodes to analyze. For large clusters, have a disaster recovery plan as an unexpected cluster-wide outage can impact many VMs.
- The fourth section of the dashboard lets you drill-down to an individual cluster.
  - A table lists all the clusters with their key configuration. You can export this list as a spreadsheet for further analysis or reporting.
  - Select a cluster. The list of ESXi hosts under the cluster, with shares and resource pool information, is automatically filled up.
  - Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service. Avoid providing different services to individual VMs as that increases the complexity of the cluster performance.
  - Keep the number of resource pools minimal.
  - Some of the columns are color coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

### Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

## ESXi Configuration Dashboard

Use the **ESXi Configuration** dashboard to view the overall configuration of the ESXi hosts in your environment, especially the configurations that need attention.

## Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

## How to Use the Dashboard

The **ESXi Configuration** dashboard is organized into sections for ease of use.

- The upper section of the dashboard displays basic ESXi configurations that should be standardized for ease of operations.
  - There are six pie-charts that are displayed as one set because there is a relationship between their values. There should be a correlation between them. Ideally, the ESXi version, the ESXi build, and the BIOS must be identical across all ESXi hosts in a cluster. Keep the variations of the hardware model, NIC speed, and storage path minimal. The more complex the pie-chart, the more variants you have. This results in complex operations, that potentially results in higher operating expenses.
  - The configurations should reflect your current architecture standard. Each pie-chart counts the occurrences of a particular value. A large slice signifies that the value is the most common value, and if that is not your current standard, then you must address it.
- The second section of the dashboard displays configurations that are potentially suboptimal.
  - The three bar-charts display various size dimensions of the ESXi hosts. The bar-charts are designed to be seen as one set. Ensure that there are a minimal number of variations to reduce complexity.
  - Smaller ESXi hosts have a relatively higher overhead, and are limited in running larger VMs. If they have a low core count, they might be using an outdated CPU. Small ESXi hosts are more expensive on a per core, per GB, per rack unit basis than larger ones if they occupy the same space. However, a 4-CPU socket ESXi host is likely to be too large, resulting in a concentration risk (too many VMs in a single ESXi host). Maintain a good balance that balances your budget and risk constraints.
  - Adjust the distribution chart bucket size to fit your environment.
- The third section of the dashboard displays configurations that you might want to avoid.
  - The six bar-charts focus on security, availability, and capacity settings that you can set as a standard. For example, you should consider activating the NTP daemon for a consistent time, which is critical for logging and troubleshooting.
  - The three tables list the actual ESXi hosts that are in a non-productive state. They can be on maintenance mode, powered off, or in a disconnected state.
- The last section of the dashboard displays all the ESXi hosts in your environment.
  - You can sort the columns and export the results into a spreadsheet for further analysis.

  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

### Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result, there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- For complete visibility, consider adding physical server monitoring by using the appropriate management pack. For more information, see the following page.

## Network Configuration Dashboard

Use the **Network Configuration** dashboard to view the overall configuration of vSphere distributed switches in your environment, especially for the areas that need your attention.

### Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

### How to Use the Dashboard

The dashboard is organized into two sections for ease of use.

- The first section displays network configurations that need your attention.
  - There are five bar-charts that focus on critical security settings.
  - The last bar-chart displays the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section provides overall configuration information, with the ability to drill down to a specific switch.
  - Click the row to select a switch from the list.
  - The ESXi hosts, port groups, and the VMs on the switch are displayed.
  - Review each of the tables. For the ESXi host table, ensure that the settings are consistent.
  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
  - You can sort the columns and export the result into a spreadsheet for further analysis.

### Points to Note

- `No data to display` does not imply that there is something wrong with data collection by VMware Aria Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- For complete visibility, consider adding physical network device monitoring by using the appropriate management pack. For more information, see the following page.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

## VM Configuration Dashboard

Use the **VM Configuration** dashboard to view the overall configuration of virtual machines in your environment, especially for the areas that need attention.

### Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

### How to Use the Dashboard

- Click the row to select a data center from the data center table.
  - In a large environment, loading thousands of VMs increases the web page loading time. As a result, the VM is grouped by data center. In addition, it might make sense to review the VM configuration per data center.
  - For a small environment, vSphere World is provided, so you can view all the VMs in the environment.

  The **VM Configuration** dashboard is organized into three sections for ease of use. All the three sections display the VM configuration for the selected data center.
- The first section covers limits, shares, and reservations.
  - Their values can easily become inconsistent among VMs, especially in an environment with multiple vCenter Servers.
  - Shares should be mapped to a service level, to provide a larger proportion of shared resources to those VMs who pay more. This means that you should only have as many shares as your service levels. If your IaaS provides gold, silver, and bronze, then you should have only three types of shares.
  - Value of the shares and reservation is relative. If you move a VM from one cluster to another (in the same or different vCenter Server), you might have to adjust the shares.

  - Reservation impacts your capacity. Memory reservation works differently from CPU reservation, and it is more permanent.
- The second section covers VMware Tools.
  - VMware Tools is a key component of any VM, and should be kept running and up to date.
- The third section covers other key VM configurations.
  - Keep the configurations consistent by minimizing the variants. This helps to reduce complexity.
  - **VM Network Cards** widget. If you suspect that your environment might have a VM with no NIC, consider adding it as a dedicated bucket.
- The last section of the dashboard is collapsed by default.
  - You can view all the VMs with their key configurations.
  - You can sort the columns and export the results into a spreadsheet for further analysis.

### Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.
- The pie-chart and bar-chart cannot drive other widgets. For example, you cannot select one of the pie-slices or buckets, and expect it to act as a filter to a list or a table.
- You can apply a specific color in a pie-chart or distribution chart for a specific numeric value, but not string value. For example, you cannot apply the color red to the value `Not Installed`.

## vSAN Configuration Dashboard

The **vSAN Configuration** dashboard provides overall configuration details and is useful in large clusters with many vSANs, where you have to follow a certain standard configuration.

### Design Considerations

See Configuration Dashboards for common design considerations among all the dashboards for configuration management.

### How to Use the Dashboard

The **vSAN Configuration** dashboard is organized into three sections for ease of use.

- The first section displays six pie-charts.
  - There are five bar-charts that focus on critical security settings.
  - The last bar-chart shows the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section displays three bar-charts.
  - The three bar-charts together provide a good overview of the vSAN key capacity configuration. By analyzing the distribution, you can identify if you have capacity configuration that is outside your expectation.
- The last section of the dashboard displays all the vSAN clusters with their key configuration.
  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
  - You can sort the columns and export the result into a spreadsheet for further analysis.

### Points to Note

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

## Workload Management Configuration Dashboard

This dashboard provides a quick configuration summary of all the key objects associated with workload management such as Supervisor Clusters, Namespaces, vSphere Pods and Tanzu Kubernetes clusters. It is essential that the configuration is consistent across all the objects. Configuration drifts may result in inconsistent performance or availability of the applications leveraging workload management Kubernetes constructs.

Use the dashboard to ensure that the configuration is consistent across all objects.

You can view the following widgets in the dashboard.

- **Environment Summary**
- **Supervisor Cluster Versions**
- **Cluster Status**
- **Pod Data**
- **Supervisor Cluster Configuration Summary**

- **Pod Configuration Summary**
- **Kubernetes cluster Configuration Summary**
- **Namespace Configuration Summary**

## Consumer \ Correct it? Dashboard

The **Consumer \ Correct it?** dashboard complements the main VM configuration dashboards by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

### Design Considerations

The dashboard is designed to focus on VMs that need attention. Lists are used to keep it simple, and show actual objects. The lists can be tailored using the filter and the custom group. The lists can also be exported for an offline discussion.

The dashboard is extendable, reflecting the reality that different customers have a different set of settings to verify. Since the dashboard layout is a collection of tables (List View), you can extend it by adding more tables. You can add more List View widgets to verify the VM configurations that your operations require.

### How to Use the Dashboard

The **Consumer\Correct it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Tools Widgets:
  - Using VMware Tools has multiple benefits. For the list of benefits, refer to KB 340.
  - VMware Aria Operations uses VMware Tools to retrieve Guest OS metrics. Without this, right-sizing VM memory can be inaccurate, because the hypervisor metrics (VM Memory Consumed and VM Memory Active) are not designed to measure Windows or Linux memory utilization. ESXi VMkernel does not have visibility into the Guest OS for security reasons.
  - Independent software vendor (ISV) support is the most common reason that VMware Tools is not installed. The ISV vendor might claim that no additional software is installed in their appliance unless they have certified it. For more information about VMware Tools, see the VMware Tools documentation.
  - If VMware Tools is installed, there might be reasons why the application team deactivates it. The Infrastructure team should inform and educate their application team, and document the technical recommendations about why VMware Tools is recommended to be running all the time.

- CPU Limits and Memory Widgets:
  - It is recommended that you do not use memory and CPU limits as it can result in an unpredictable performance. The Guest OS is not aware of this restriction as it is at the hypervisor level. It is recommended that you shrink the VM instead.
- Guest OS Counters Missing Widget:
  - There is no visibility into the Guest OS performance counters because the requirements are not met. The memory counter is especially important as VM Consumed and VM Active are not replacements for Guest OS counters. See KB 55675 for more details.
- Old Snapshot Widget:
  - Ensure that the snapshot is removed within one day after the change request. If not, it might result in a large snapshot and impact the performance of the VM.

### Points to Note

- Add a banner summary to the top of this dashboard so that you can verify if there is an incorrect confirmation. Add a scoreboard and select the World object and then collapse all the tables below. Create a super metric for each summary and apply it to the World object.
- In a large environment, create a filter for this dashboard to help you to focus on a segment of the environment. Group it by a class of service such as, gold, silver, and bronze. Default the selection to gold, your most important environment. In this way, your monitoring is not cluttered with less critical workloads.
- There are other VM configurations that maybe relevant to your environment. Review the list of VM settings that you might want to add to this dashboard.
- For context, add a property widget that lists the selected VM properties. In this way, you can check the property of your interest without leaving the screen. Multiple List View widgets can drive the same property widget, so you do not have to create one property widget for each List View.
- If your operations require it, add a list of VMs that do not have these three key performance counters: CPU Run Queue, CPU Context Switch, and Disk Queue Length.

## Consumer \ Optimize it? Dashboard

The **Consumer \ Optimize it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Optimize it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities. A suboptimal configuration might not impact performance or increase complexity, but it can be more expensive.

### Design Considerations

The **Consumer \ Optimize it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How To Use the Dashboard

The **Consumer \ Optimize it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- VM Reservation:
  - VM reservation causes a positive impact on the VM, but a negative impact on the cluster. Total reservation cannot exceed cluster capacity. This creates a suboptimal cluster as VMs do not use the entire assigned memory at the same time.
  - VM reservation places a constraint on the DRS placement and HA calculation. Avoid using reservation as a means to differentiate performance SLA among all the VMs in the same cluster. It is difficult to correlate CPU Ready with CPU Reservation. A VM CPU Ready does not improve two times because you increase its CPU reservation by two times. There is no direct correlation.
- Guest OS visibility:
  - Since your workloads are sharing resources and are over-committed, your operations are easier if you know what is running inside. This helps with monitoring and troubleshooting, resulting in optimal operations.
  - For critical VMs, consider logging the Guest OS, such as Windows and Linux, to capture errors that do not surface as metrics. These errors typically appear as events in the log files or in the event database in the case of Windows. Use VMware Aria Operations for Logs to parse Windows events into log entries that can be analyzed.
- Snapshot:
  - Old snapshots tend to be larger. They consume more space and have a higher chance of impacting performance.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result shares limitations and customization ideas.

## Consumer \ Simplify it?

The **Consumer \ Simplify it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere

administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Simplify it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

## Design Considerations

The **Consumer \ Simplify it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

## How to Use the Dashboard

The **Consumer \ Simplify it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Large VMs (CPU, Memory, and Disk):
  - A large VM, relative to the underlying ESXi host and datastore, requires more careful planning (Day 0) and monitoring (Day 2).
  - Ensure that the VM size does not exceed the size of the underlying ESXi host. If your ESXi host has CPU hyper-threading, do not count the logical processor. Instead, count the physical core. For best performance, keep it within a (non-uniform memory access) NUMA boundary.
  - During monitoring, verify if the VM is highly utilized. If the VM vCPU count is equal to the ESXi cores, and the VM is running at almost full capacity, you might not be able to run other VMs. Large VMs can impact the performance of other VMs, especially if it is given higher shares. Only when the large VM is under-utilized, can the ESXi hosts run other VMs.
  - If the number of configured vCPUs on a VM is higher than the number of cores per socket on the ESXi, the VM can experience the NUMA effect. If the ESXi has more than one physical CPU (socket), cross-NUMA access negatively impacts performance.
  - The larger the VM, the longer the time required to vMotion, Storage vMotion, and backup.
  - For disk space, if the disk is thin-provisioned and under-utilized, you can deploy other VMs in the same datastore. Ensure that the snapshot is tracked closely, as the risk of capacity running out is higher for a large virtual disk.
- VMs with many virtual disks:
  - It is simpler to have a 1:1 mapping between Guest OS partitions and the underlying virtual disk (VMDK or RDM).
  - For performance and capacity, evaluate the disks and partitions. Each virtual disk must be monitored in terms of IOPS, throughput, and latency. Having multiple virtual disks increases the monitoring and troubleshooting need.

  - If the reason for having many virtual disks is performance, identify which counter serves as proof that multiple virtual disks are required. It is possible that the performance required is met by a single virtual disk.
- VM with many IP addresses or NICs:
  - A VM might need multiple networks, such as production, back up, and management. It is recommended that you route the network interfaces through the NSX-Edge VM. A VM that has multiple network interfaces can bridge the network, causing security risks or network problems.
  - A VM that is part of multiple networks can do so with just a single NIC. A single NIC can be configured to access multiple networks, with each interface having their own IP configuration.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

## Consumer \ Update it? Dashboard

The **Consumer \ Update it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Update it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

### Design Considerations

The **Consumer \ Update it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboard

The **Consumer \ Update it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated Tools Widget:
  - Lists all the VMware Tools versions that are still supported. Tailor the filter to fit your operational needs.

- Outdated VM Hardware Widget:
  - Lists all the VM vmx versions that are not 13, 14, 15, or 16. Tailor the filter to fit your operational needs.
- Outdated Windows and Red Hat Widgets:
  - Lists all the Windows client versions that are not version 10.
  - Lists all the Windows server versions that are not versions 2016 and 2019.
  - Lists all the RHEL versions that are not version 7 or 8.
  - If you run other operating systems like Ubuntu, clone the widget. You can also repurpose the widget if you do not run RHEL and Windows.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

## Provider \ Correct it? Dashboard

The **Provider \ Correct it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

### Design Considerations

The **Provider \ Correct it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations.
  - A cluster is the smallest logical building block for compute. Consider it as a single computer with physically independent components. As a result, consistency matters.
  - Clusters with DRS set to manual. This means that DRS initiated vMotion does not take place unless the administrator manually approves it. Since DRS calculates every five minutes, your quick approval is required to prevent a change of condition.
  - Clusters with HA disabled. Without high availability provided by the infrastructure, each application must protect itself from an infrastructure failure.

  - Clusters with DRS disabled. DRS focuses on performance and capacity, while HA focuses on availability. Without DRS, you must build a buffer on every ESXi host to cope with peak demand.
  - Clusters with Admission Control disabled. Reservation is respected only when Admission Control is activated.
- The second section covers the ESXi host configurations.
  - ESXi with Network Time Protocol disabled. Logs are a critical component of operations, and are the main source of information in troubleshooting. While troubleshooting performance across objects, the sequence of logs determines which event is the likely root cause, as the oldest event starts the chain of events.
  - A disconnected ESXi host indicates that the ESXi host is not participating in HA and you cannot migrate any VM on it.
  - An ESXi host that is in maintenance mode does not contribute resources to the cluster or the data center if there is a standalone ESXi.
- The third section covers ESXi host configurations that must be consistent within a cluster.
  - BIOS version and ESXi versions.
  - BIOS Power Management, ESXi: Power Management. Ideally, should be set to OS controlled. The ESXi level should be set to balance level.
  - ESXi Storage Path. Ensure that the number of paths and the path policies are identical.
  - ESXi hardware specifications. Different specifications can result in inconsistent performances experienced by the VM.

### Points to Note

- See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- If you have a standalone ESXi, and you plan to replace it with a clustered ESXi host, add a table to list them.
- Based on your security settings, add a table to check the Distributed Switch and Port Group to ensure that security settings such as promiscuous mode, are used correctly.

## Provider \ Optimize it? Dashboard

The **Provider \ Optimize it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Optimize it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

## Design Considerations

The **Provider \ Optimize it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

## How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations:
  - A small cluster has a higher HA overhead when compared to a large one. For example, a three-node cluster has 33% overhead while a 10-node cluster has 10%. For vSAN, a low number of hosts limits the availability option. Your choice of FTT is relatively more limited.
  - Many small clusters result in silos of resources. As a cluster behaves like a single computer, ensure that it has enough CPU cores, CPU GHz, and Memory. For ESXi in 2020, it is typical to have 512 GB of RAM. This results in 12 TB of RAM for a 12-node cluster, which is enough for DRS to place many VMs as it balances them.
  - If there is a lot of reservation, add a list for clusters with a relatively high reservation. If your clusters are of different sizes, use a super metric to convert the reservation value to a percentage.
- The second section covers ESXi host configurations.
  - Small ESXi. A small host faces scalability limits in running a larger VM. While a 2-socket, 32-cores, 128 GB memory ESXi can run 30 vCPU, 100 GB RAM VMs, the VM experiences a non-uniform memory access (NUMA) effect.
  - ESXi powered off. You can mark the ESXi hosts for decommissioning using the custom property feature of VMware Aria Operations. You can then create a separate list, so they are not overlooked.
- The third section cover storage and network.
- 
- Unused network (distributed port group). This is a potential security risk as you might not monitor it.

## Points to Note

- See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- For CPU cores, a change in vSphere licensing means that the ideal core is 32-cores per CPU socket. This maximizes the software license. For more information, see the vSphere Pricing Model.

# Provider \ Simplify it? Dashboard

The **Provider \ Simplify it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Simplify it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

## Design Considerations

The **Provider \ Simplify it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

## How to Use the Dashboard

- Click the row in the **Clusters** widget to select one of the clusters from the table.
  - A cluster is more complex to operate when it has resource pools, shares, and limits.
- Review the list of resource pools:
  - Ensure that the number of VMs in each resource pool reflects the intended settings for the VM. The resource pool value is divided and shared among the VMs. The more the VMs, the lesser the resources allotted to each VM.
  - Verify if there are VMs who are siblings to the resource pools.
  - Verify if the resource pools are further split into subresource pools.
- Review the CPU Share and Memory Shares pie-charts:
  - Multiple combinations of shares, especially both CPU and memory, makes troubleshooting difficult.
  - Each share must map to exactly one class of service, such as one for gold and one for silver as the shares define the class of service. Shares are also relative, meaning the value depends on the value of sibling objects, such as, resource pool or VM. Ensure that the values are consistent across clusters to avoid unintended consequences while moving the VM to another cluster.
- Review the CPU Reservation and Memory Reservation tables:
  - High total reservation, especially both CPU and memory, complicates the cluster operations as it impacts the HA slot calculation, and limits the DRS choice of placement.
- Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

## Provider \ Update it? Dashboard

The **Provider \ Update it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Update it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

As part of operations best practices, keep the infrastructure up to date. Running outdated components that are too far behind the latest version, can cause support problems or upgrade problems. It is common that the fix for the problem is only available in the later versions. Outdated hardware can also result in higher operating costs. Outdated hardware might cost more data center footprint, such as rack space, cooling, and UPS. Refreshing your technology and consolidation are two common techniques to optimize cost.

### Design Considerations

The **Provider \ Update it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? Dashboard dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboards

The **Consumer \ Update it?** dashboard is a collection of tables (List View) that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated vSphere Components Widgets:
  - Lists all the vCenter Servers versions that are not 6.7 or 7.0.
  - Lists all the ESXi host versions that are not 6.5, 6.7, or 7.0.
  - Lists all the vSAN ESXi host versions that are not 6.7 or 7.0. A more stringent filter is applied for vSAN because of a relatively higher maturity in the latest release. From VMware Aria Operations and VMware Aria Operations for Logs, there are more counters, properties, and events that improve monitoring and troubleshooting.
  - Lists all the vSphere distributed switches, regardless of the version.
  - You should tailor the filter to fit your operational needs.

- Outdated Server BIOS Widget:
  - Lists all the ESXi hosts regardless of the BIOS version. Edit the widget and tailor the filter to fit your operational needs.
- Other than customizing the existing widgets, consider adding the following checks:
  - ESXi hosts with outdated hardware, using a filter based on your environment.
  - ESXi hosts that are no longer on warranty. Create a custom property to capture the end of warranty.
  - Physical storage arrays with outdated firmware, model, and an expiring warranty.
  - Physical network switch with an outdated OS version and hardware model

  **Note** Install the relevant management pack for the last two points.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? Dashboard dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

# Cost Dashboards

The dashboards in the cost category cater to cloud administrators who are responsible for managing the expenses related to your cloud infrastructure. Using Cost dashboards, you can compare the cost of VMware cloud infrastructure with other cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently.

## Consumer Layer

The consumer layer dashboards of VMware Aria Operations helps you to know how a customer can do a deeper analysis of the Return on Investment from the consumer perspective.

The available dashboards for consumers are:

- Chargeback VM Price Dashboard
- Showback VM Cost Dashboard
- Showback vSphere Pod Cost Dashboard

### Chargeback VM Price Dashboard

The chargeback VM price dashboard lets you know how much you must spend to run a VM on behalf of your customer. In VMware Aria Operations, you can configure the cost drivers and let the system automatically determine how much a VM costs based on your infrastructure requirement. Cost Drivers cover server hardware, storage, licenses, application, maintenance, labor, network, facilities, and additional costs configured within VMware Aria Operations.

Price is what you charge your customer for running their VM. The price of a VM can be based on the cost of the VM or based on a rate card that you define. Prices can include up charges, service charges, and others.

### How to Use the Dashboard

- Select a Group widget displays the price of the group.
- Price Summary of Selected Group shows the month to date price of the group.
- VM Price Distribution (Top 100) shows the most expensive VMs in the group.
- Powered Off VMs shows reclaimable VMs and their potential savings.
- Idle VMs shows reclaimable VMs and their potential savings.
- VMs with Snapshots shows reclaimable snapshots and their age.
- Price of VMs in the Selected Group shows the price and configuration of each VM in the selected group.

## Showback VM Cost Dashboard

The Showback VM Cost dashboard provides a quick Showback of the cost associated with the VMs in a group. Based on the Showback you can improve the accuracy of the costs by editing the cost drivers. Cost drivers that are not customized use reference cost, cost driver customization is available only in Advanced or Enterprise edition of VMware Aria Operations.

### How to Use the Dashboard

- Select an object in the Select a Group widget to view the cost of the group.
- Cost Summary (This Month) shows the month to date cost, potential savings, and projected cost of the group.
- VM Cost Distribution (Top 100) shows the most expensive VMs in the group.
- Potential Savings (Top 10) shows the VMs ranked by their potential savings.
- Members of the Group (Select to View Trend) shows the cost and configuration of each VM in the selected group.
- Cost Trend of Selected VM shows the trend of the VMs cost over time.

## Showback vSphere Pod Cost Dashboard

The Showback vSphere Pod Cost Dashboard provides a quick Showback of the cost associated with the vSphere Pods in a group. Based on the Showback you can improve the accuracy of the costs by editing the cost drivers. Cost Drivers that are not customized use reference cost, cost driver customization is available only in Advanced or Enterprise edition of VMware Aria Operations.

### How to Use the Dashboard

- Select an object in the Select a Group widget to view the cost of the group.

- Cost Summary (This Month) shows the month to date cost and projected cost of the group.
- vSphere Pod Cost Distribution (Top 100) shows the most expensive vSphere Pods in the group.
- Idle vSphere Pods shows the vSphere Pods that have been identified as potentially idle.
- Members of the Group (Select to View Trend) shows the cost and configuration of each vSphere Pod in the selected group.
- Cost Trend of Selected vSphere Pod shows the trend of the vSphere Pod's cost over time.

## Provider Layer

The provider layer dashboards of VMware Aria Operations, helps you to know how a customer can analyze the Return on Investment for the virtual infrastructure used in the customer’s environment.

The available dashboards for providers are:

- Assess Cost Dashboard
- Datacenter Cost Drivers Dashboard
- Server Hardware Depreciation Dashboard
- Base Rate Analysis Dashboard
- VM Cost versus Price Dashboard
- Reclaimable Hosts Dashboard

### Assess Cost Dashboards

The **Assess Cost** dashboard provides an overview of the scale of your infrastructure in terms of physical capacity available.

#### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

#### Widget Information

- You can view the total cost of ownership per month for the infrastructure and the savings opportunities details, if any, for the infrastructure.
- You can view the details of the division of infrastructure investments across all data centers. The dashboard provides the magnitude of each data center in terms of the number of physical servers and virtual machines. It also provides details about the amount of savings that can be achieved from each of these data centers.
- The dashboard displays data about how you invest across clusters of different quality offered across all vCenter Servers.

## Base Rate Analysis Dashboard

The **Base Rate Analysis** dashboard helps you analyze the cost efficiency of your data center.

### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

### Widget Information

- The total cost of ownership is the cost required to run your data center per month. This is derived from the cost drivers.
- The average cost per VM is derived by considering the cost of all the VMs in your environment. The cost of each VM depends on the base rate of the cluster the VM is placed on and its utilization. The base rate of the cluster is computed based on the total cost of ownership and the expected utilization levels of the cluster. Storage base rates are directly obtained from cost drivers.
- If the cluster is running on an allocation-based capacity model, the base rate is derived from the total cost of the cluster and the over-commit ratio. The base rate is indicative of how costly a resource is, on a given cluster.
- A base rate is derived from the total cost and the expected utilization of the cluster.
- A deeper analysis of the base rates can be performed using the CPU, memory, or storage-related widgets, which help rank clusters and datastores relative to their base rates.

## Datacenter Cost Drivers Dashboard

The **Datacenter Cost Drivers** dashboard provides the cost of different data centers in a private cloud.

### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

### Widget Information

- You can select individual data centers to view summary and trends. The summary of the data center costs is grouped into two:
  - Compute. Covers all the costs that are spent on compute related hardware, software, and services.
  - Non-Compute. Covers storage and network.
- Expense trends provide cost variations over a period which indicate infrastructure additions or removal to the data center.

- Cluster expenses indicate the component clusters of a data center that consume the costs. Datastores that represent the storage part of the data center cost are listed alongside.

  **Note** Network costs are mapped directly to ESXi hosts and hence are costed under compute as well, as of today. This might change in the future.

- When you select a cluster, you can view the component hosts that the cluster is made up of and their monthly depreciated costs. It also provides details on the purchase cost of the server and how many months until it depreciates completely.

**Note** Server costs can be suggested out-of-the-box by the system, or can be customized by the user. Depreciation information is not available for servers when the server costs are suggested out-of-the-box by the system. Depreciation information is available for those servers when the server cost is customized by the user.

## Reclaimable Hosts Dashboard

The Reclaimable Hosts dashboard helps you to identify clusters with reclaimable hosts and the potential cost savings from reclaiming the hosts. Reclaimable hosts are identified from the Total Recommended Capacity generated by the AI powered capacity engine in VMware Aria Operations.

### Widget Information

- The Reclaimable Hosts Cost Pie Chart displays the reclaimable host cost distribution for individual clusters in your virtual environment.
- Potential Savings graph depicts the total cost savings (potential) for all the clusters in your virtual environment for a given period.
- The Top 10 Clusters with Reclaimable Hosts displays the number of reclaimable hosts.
- The Top 10 Clusters with Reclaimable Hosts by Cost displays the top 10 reclaimable hosts by cost.

## Server Hardware Depreciation Dashboard

The Server Hardware Depreciation Dashboard helps you to calculate the depreciation value for server hardware which is marked as owned in Cost Drivers. You can configure the depreciation cost settings as per your business requirement.

### Widget Information

- The Server Purchase Cost is the total purchase price of all servers as entered in Cost Drivers.
- Accumulated Depreciation is the amount of server purchase costs that have been depreciated according to purchase date and depreciation settings.
- Remaining Depreciation is the amount of server purchase costs left to be depreciated.

- Number of Fully Depreciated Servers identifies servers that have been fully depreciated. These servers may exhibit higher failure rates or have lower capacity. Use What-If scenarios to model the cost and capacity impact of replacing these servers.

### VM Cost vs. Price Dashboard

The VM Cost vs. Price Dashboard helps you to analyze the relationship between cost and price for virtual machines. You can use this dashboard to ensure the price of VMs for chargeback is sufficient to cover the cost of running virtual machines.

#### How to Use the Dashboard

- Select a Group allows selection of a group of VMs to analyze.
- Summary (Month to Date) shows the month to date price and cost.
- Members of the Group (Select to View Trend) shows all VMs in the selected group with their Month to Date Cost, Today's Cost, Month to Date Price, and Today's Price.
- Daily Cost and Daily Price trend chart shows both cost and price over time.

## Cost Optimization

The cost optimization dashboards are as follows:

- Cost Optimization Dashboard
- Potential Cost Savings Dashboard
- Realized Cost Savings Dashboard
- Total Cost of Ownership Dashboard
- VM Rightsizing Details Dashboard

### Cost Optimization Dashboard

The Cost Optimization dashboard helps you to measure the return on investment if you use VMware Aria Operations to manage your virtual infrastructure. You can track the total cost of ownership of the entire environment along with potential savings and realized savings from recommendations provided, the dashboard helps you quantify the cost efficiency and cost savings over time.

Potential Savings is a summary of all cost savings opportunities identified by VMware Aria Operations. Realized Savings is a summary of cost savings from actions performed that are related to recommendations provided by VMware Aria Operations.

#### Widget Information

- Total Cost of Ownership provides details of the monthly cost of server hardware, licenses, maintenance, facilities, labor, network, storage, and additional costs.

- Average Cost per VM is a good indicator of cost efficiency over time. It is natural for the cost per VM to go up when new capacity is added and trend downwards as additional capacity is consumed. The goal is to reduce the average cost per VM over time.
- Realize Savings Breakdown shows the cost of reclaimed resources from the VM identified by VMware Aria Operations.
- Potential Savings covers the cost savings opportunities identified by VMware Aria Operations.

## Potential Cost Savings Dashboard

The Potential Cost Savings Dashboard helps you to measure the cost saving as reported by VMware Aria Operations. You can evaluate the potential savings to track recommendations and improve cost efficiency over time. The dashboard shows both cost savings and capacity savings for idle VMs, powered off VMs, VM snapshots, orphaned disks, oversized VMs, and reclaimable hosts.

### Widget Information

- The Cost Savings Breakdown widget displays potential savings and reclaimable capacity for idle VMs, powered off VMs, VM snapshots, orphaned disks, oversized VMs, and reclaimable hosts. You can also view the allocation changes for the oversized VMs.
- Reclaimable widget provides the metric details for the reclaimable vCPU, reclaimable memory, and reclaimable disk space.
- The Optimization Opportunities Breakdown widget covers the projected costs to improve performance as identified by VMware Aria Operations.
- Allocation Changes for Undersized VMs shows the number of vCPUs and GB of memory to add to undersized VMs.

## Realized Cost Savings Dashboard

The Realized Cost Savings Dashboard helps you to quantify the realized cost savings from actions performed that are related to recommendations provided by VMware Aria Operations. You can analyze the realized savings to track improvements to cost efficiency over time. Realized savings covers powered off VMs that were flagged as idle, deleted VMs that were flagged as idle or powered off, deleted snapshots that were flagged as reclaimable, deleted disks that were flagged as orphaned, oversized VMs that were rightsized, and deleted hosts that were flagged as reclaimable.

### Widget Information

- Realized Savings covers the cost savings from reclamation opportunities recommended by VMware Aria Operations.
- The Reclaimed Capacity shows the amount of capacity that was reclaimed based on recommendations from VMware Aria Operations.

- Allocation Changes for Oversized VMs shows the number of vCPUs and GB of memory removed from formerly oversized VMs.
- Cost of Deleted VMS shows the cost of all deleted VMs in the past 30 days, shows the cost of all deleted VMs (by cluster) for the past 30 days, and shows the year-to-date cost of all deleted VMs.

### Total Cost of Ownership Dashboard

The Total Cost of Ownership dashboard helps you to understand the total cost of ownership of your environment from multiple perspectives. You can use this dashboard to learn how cost drivers, capacity, and data centers affect the total cost of ownership.

#### Widget Information

- Cost Driver Breakdown widget shows how cost drivers affect the total cost of ownership.
- Cost of Capacity Used and Capacity Remaining widgets shows cost breakdown by the cost of capacity used and the cost of capacity remaining.
- Cost per Datacenter widget shows how the costs broken down per data center.

### VM Rightsizing Details Dashboard

The VM Rightsizing Details dashboard provides an overview of the rightsizing recommendations for Undersized VMs and Oversized VMs. Rightsizing is defined as changing the amount of resources allocated to a VM based on the Recommended Size for a VM. Recommended Size is the maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining.

#### How to Use the Dashboard

- Select a Cluster, Datacenter, or World Object.
- Select an Undersized VM to view the recommendations.
- Select an Oversized VM to view the recommendations.
- Search for a VM to view the recommendations.

## Performance Dashboards

Performance is about ensuring workloads get the necessary resources. Key Performance Indicators (KPI) can be used to identify performance problems related to workloads. Use these KPIs to define SLAs associated with tiers of service. These dashboards use KPIs to display the performance of workloads at the consumer layer and the aggregate performance of workloads at the provider layer.

SLA is the formal business contract that you have with your customers. Typically, SLA is between the IaaS provider (the infrastructure team) and the IaaS customer (the application team or business unit). Formal SLA needs operational transformation, for example, it requires more than technical changes and you might need to look at the contract, price (not cost), process,

and people. KPI covers SLA metrics and additional metrics that provide early warning. If you do not have an SLA, then start with Internal KPI. You must understand and profile the actual performance of your IaaS. Use the default settings in VMware Aria Operations if you do not have your own threshold, as those thresholds have been selected to support proactive operations.

The following graphics depict the above relationship.

| Reactive | Internal KPI | Formal SLA |
| --- | --- | --- |
| Complaint-based Operations.<br><br>Blamestorming.<br><br>IaaS Performance Measures based on Business Impact. | IaaS Performance is Quantified and Measured.<br><br>Performance is based on the Production Environment.<br><br>Default Policy. | SLA is part of the business contract. Customers can track their SLA using the self-service portal.<br><br>One policy for each SLA. |

## The Three Processes of Performance Management

In performance management, there are three distinct processes.

- Planning. Set your performance goals. When you architect a vSAN, you must know how many milliseconds of disk latency you want. 10 milliseconds measured at the VM level (not the vSAN level) is a good start.
- Monitoring. Compare the plan with the actual. Does the reality match what your architecture was supposed to deliver? If not, you must fix it.
- Troubleshooting. When the reality is not according to the plan, you must fix it proactively and not wait for issues and complaints.

To understand what is not healthy for performance management consider the following areas in the given order.

1. Contention: This is the primary indicator.
2. Configuration: Check the version incompatibilities.
3. Availability: Check for soft errors. vMotion stun time, lock up. This requires Log Insight.
4. Utilization: Check this in the end. If the first three parameters are good, you can skip this.

## The Three Layers of Performance Management

There are three main realms of enterprise applications. Each of these realms has its own set of teams. Each team has a set of unique responsibilities and requires the associated skill set. The three realms comprise of Business, Application, and IaaS. Refer to the graphic below to understand the three layers and the typical questions asked on each layer.

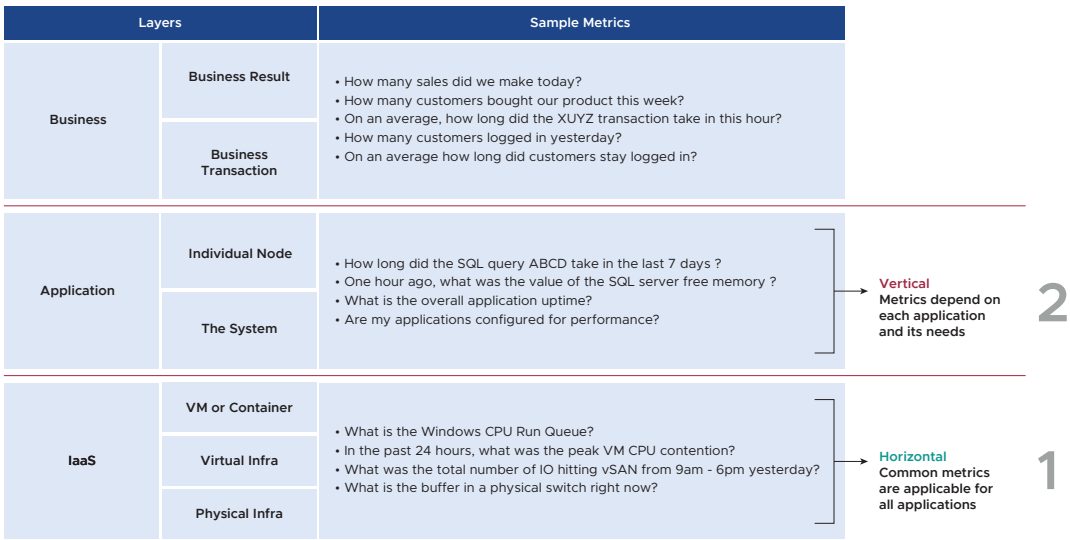


Performance Management is largely an exercise in elimination. The methodology slices each layer and determines if that layer is causing the performance problem. Hence it is imperative to have a single metric to indicate if a particular layer is performing or not. This primary metric is aptly named Key Performance Indicator (KPI).

The upper layer depends on the layer below it, and hence the infrastructure layer is typically the source of contention. As a result, focus on the bottom layer first, as it serves as the foundation for the layer above it. The good part is this layer is typically a horizontal layer, providing a set of generic infrastructure services, regardless of what business applications are running on it.

## The Two Metrics of Performance Management

The primary counter for performance is contention. Most look at utilization, because they fear something wrong might happen if utilization is high. That something is contention. Contention manifests in different forms like, queue, latency, dropped, canceled, and context switch.

However, do not confuse ultra-high utilization indicators as a performance problem. If your ESXi host experiences ballooning, compression, and swapping, it does not mean that your VM has a performance problem. You measure the performance of the host by how well it serves its VMs. While performance is related to the ESXi host utilization, the performance metric is not based on the utilization, instead it is based on contention metrics.

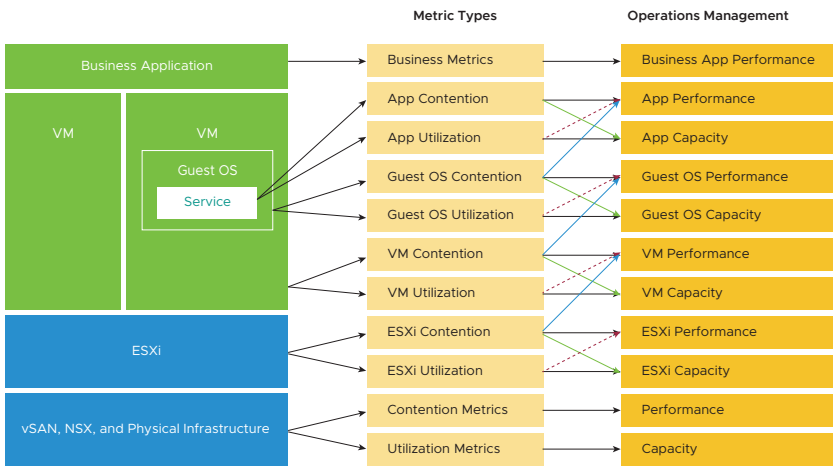


It is possible for VMs in the cluster get affected from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at an individual consumer (VM). The following table shows various possible reasons.

| Infra Configuration | VM and Guest OS Configuration |
| --- | --- |
| ESXi Settings<br>■ Host and BIOS power management causes Frequency to drop.<br>■ HT enabled. It looks like twice the capacity, but it is actually 1.25 X throughput.<br>■ ESXi - HW compatibility. Driver and firmware are two areas that can impact the performance.<br>■ Mismatch of queue depths along the various storage stacks. Must calibrate all the way to the physical array.<br>■ vMotion too slow or high stunned time. | VM: Limit, Share, and Reservation<br>■ Make sure that no limit is set. CPU ready includes limit.<br>■ Make sure that the shares are consistent (as per what the VMs want or you agree to.)<br>■ Avoid reservation if possible. This impacts the net available resources for the other VMs. |
| Network<br>■ MTU mismatch.<br>■ Hops. Especially horse-shoe, or going through multiple ESXi. | Size: NUMA effect. VM spanning NUMA nodes. |
| Cluster Settings<br>■ Inconsistent configuration among hosts in a cluster. EVC Mode can play a part if the hosts are from different generations.<br>■ Resource Pool<br>  ■ Make sure the shares match the number of VMs.<br>  ■ Make sure that no VM is siblings to RP.<br>■ VM- Host Affinity.<br>■ DRS Setting. | Snapshot. IO is processes 2x.<br>VM drivers. |
| vSAN<br>■ The host where the storage was having performance issues. | Windows or Linux process ping pong, process runaway, and OS level queue. |

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

## VM Performance

Since VM is the most important object in vSphere, it warrants an extra explanation. The graphic below lists the counters you should look at.

| | CPU | RAM | Network | Disk |
|---|---|---|---|---|
| **Inside Guest OS** (Linux, Windows) Need VMware Tools | Run Queue<br>Context Switch | Paging Rate (MB/s)<br>Committed % | OS Output Queue Length<br>Driver Queue | OS Queue<br>Driver Queue |
| | Utilization | In Use<br>Modified + Standby | Throughput (Mbps)<br>Latency | Latency |
| **Outside Guest OS** (Guest OS can't control) | Run I Used<br>System + VMX +MKS | Active, Consumed, Granted, Swapped-in | Throughput | IOPS,<br>Throughput (Large Block) |
| | Ready + Co-Stop + Overlap<br>IO Wait + Swap Wait | Contention | TX Dropped Packet<br>Normalized Latency | Outstanding IO<br>Latency |

The KPI counters can get technical for some users, so VMware Aria Operations includeS a starting line to get them started. You can adjust the threshold, once you profile your environment. This profiling is a good exercise, as most customers do not have a baseline.The profiling requires an advanced edition.

| | Metric | Green | Yellow | Orange | Red |
|---|---|---|---|---|---|
| **Guest OS Contention** | Total CPU Run Queue | 0-5 | > 5 | > 10 | > 20 |
| | CPU Context Switch Rate | 0 - 5K | < 25K | < 100K | > 100K |
| | Total Disk Queue Length | 0 - 25 | > 25 | > 50 | > 100 |
| **Guest OS Usage** | RAM Free (MB) | > 512 MB | > 256 | > 128 | ≤ 128 |
| | RAM Page-in Rate (KB/s) | 0 - 25K | > 25 K | > 50 K | > 100K |
| **VM Contention** | CPU Co-Stop (%) | 0 - 2.5% | > 1 | > 3 | > 5 |
| | [SLA] CPU Ready (%) | 0 - 2.5% | > 2.5 | > 5 | > 7.5 |
| | Total CPU Overlap (ms) at VM level | 0 - 1000 | > 1000 | > 2500 | > 5000 |
| | CPU IO Wait | 0 - 1000 | > 1000 | > 2500 | > 5000 |
| | [SLA] RAM Contention (%) | 0 - 1% | > 1 | > 2 | > 4 |
| | [SLA] Disk Latency (ms) | 0 - 10 ms | > 10 | > 20 | > 40 |
| | [SLA] Network TX Dropped Packet | 0 | > 0 | > 1 | > 2 |
| **VM Usage** | CPU Usage (%) | 0 - 85% | > 85 | > 90 | > 95 |

## Performance Metrics

VMware Aria Operations uses the following threshold for internal KPI.

| IaaS | VM Counter | Threshold |
|---|---|---|
| CPU | Ready | 2.5% |
| RAM | Contention | 1% |
| Disk | Latency | 10 ms |
| Network | TX Dropped Packet | 0 |

The table is an example of a stringent threshold. A high standard for performance is used because it is an internal KPI for the consumption of the infrastructure team. It is not an external formal SLA that is confirmed with the customers. There must be a buffer between the internal KPI and the external SLA so that the operations team receive early warnings and has the time to react before the external SLA is breached. A high standard also works from the mission critical point to view to the development environment. If the standard is set to the least performing environment, then it cannot be applied to the more critical development.

A single threshold is used to keep the operations simple. This means that the performance in production is expected to have a higher score than the development environment. The development environment performance is expected to be worse than the production environment, while everything else is equal. A single threshold helps to explain the difference in Quality of Service (QoS) provided by a different class of service. For example, if you pay less, you get a poor performance and if you pay half the price, expect to get half the performance.

The four elements of IaaS (CPU, RAM, Disk, and Network) as mentioned in the table, are evaluated on every collection cycle. The collection time is set at five minutes as it is an appropriate balance for monitoring. If SLA is based on one minute, it is too close and results in either cost increase or reduction in threshold.

## Design Considerations

All the performance dashboards share the same design principles. They are intentionally designed to be similar, as it is confusing if each dashboard looks different from one another, considering they have the same objective.

The dashboards are designed with separate two sections: summary and detail.

- The summary section is typically placed at the top of the dashboard to provide the overall picture.
- The detail section is placed below the summary section. It lets you drill down into a specific object. For example, you can get the detailed performance report of any specific VM.

In the detail section, use the quick context switch to check the performance of multiple objects during performance troubleshooting. For example, if you are looking at the VM performance, you can view the VM-specific information and the KPIs without changing screens. You can move from one VM to another and view the details without opening multiple windows.

The dashboard uses progressive disclosure to minimize information overload and ensure the webpage loads fast. Also, if your browser session remains, the interface remembers your last selections.

Many of the performance and capacity dashboards share a similar layout since there is a shared commonality between these pillars of operations.

## Guest OS Performance Profiling Dashboard

Use the **Guest OS Performance Profiling** dashboard to know the actual performance of your environment.

Some counters directly impact the performance of Windows or Linux, the operating systems running inside the VM. These KPIs are outside the control of the hypervisor.

Modern operating systems such as Linux and Windows use memory as cache, since it is faster than a disk. Some counters directly impact the performance of Windows or Linux. These KPIs are outside the control of a hypervisor, which means that the ESXi VMkernel cannot control the increase or decrease of the KPI values. The KPI visibility also requires an agent, such as VMware Tools. As a result, they are typically excluded in performance monitoring.

Since they are closer to the applications, it is critical to know their values and establish an acceptable range. The acceptable level of these KPIs among all the VMs in your environment varies. By profiling the actual performance across time and from all VMs, you can establish a threshold that is supported by facts. Since there are 8766 instances of 5 minutes in a month, profiling 1000 VM over a month means you are analyzing 8.8 million datapoints.

## Design Considerations

The dashboard uses progressive disclosure to minimize information overload and ensures that the webpage loads fast.

In a large environment, loading thousands of VMs increases the loading time of VMware Aria Operations. As a result, the VM is grouped by data center. For a small environment, vSphere World is provided so you can see all the VMs in the environment.

## How to Use the Dashboard

Select data center from the data centers list. The three tables listing CPU, memory, and disk will show the VMs in the selected data center or vSphere world. Each table shows the highest value in the last one week (2016 datapoints based on five minutes collection cycles), and hence uses the term max as a prefix, for example Max Page-Out/sec or Max Guest OS Disk Queue.

Select any of the VMs in any of the tables. The three line charts are displayed. They are showing data from the same VM to facilitate correlation.

- CPU table widget:
  - The Max CPU Queue column shows the highest number of processes in the queue during the given period. As a best practice, keep the queue below three for each queue. A VM with eight CPUs has eight queues, hence keep this number below 24.
  - The CPU Hyperthreading gives twice the queue as it should as both threads are interspersed in the core pipeline.
  - CPU Context Switch. There is a cost associated with the context switch. There is no guidance for this number, and it varies widely.
- Memory list widget:
  - In memory paging, the modern operating systems (Linux and Windows) use memory as cache, it is much faster than a disk. It proactively pre-fetches pages and anticipates future needs (Windows calls this Superfetch). The rate pages that are being brought in and out can reveal memory performance abnormalities. A sudden change, or one that has sustained over time, can indicate page faults. Page faults indicate that pages are not readily available and must be brought in. If a page fault occurs too frequently, it can impact application performance. While there is no concrete guidance, as it varies by application, you can view a relative size. operating systems typically use 4 KB or 2 MB page sizes.

- Disk list widget:
  - Disk queues are queued IO commands that are not sent to the VM. They have been retained inside the Guest OS (either at a kernel level or a driver level). A high disk queue in the guest OS, accompanied by low IOPS at the VM, can indicate that the IO commands are stuck waiting on processing by Windows/Linux. There is no concrete guidance regarding these IO commands threshold as it varies for different applications. You should view this with the Outstanding Disk IO at the VM layer.

### Points to Note

- These Guest OS widgets do not appear unless the vSphere pre-requisites are met. For more information, see KB article 55697.
- Once you determine an acceptable threshold for your environment, consider adding thresholds to the table so you can easily view the VMs that exceed a threshold.
- The CPU queue is the sum from all virtual CPUs. A larger VM can tolerate a higher queue as it has more processors. If you want to compare VMs of different sizes, create a super metric that calculates the queue per vCPU. For more information, see Create a Super Metric.
- Group the VM by clusters of the same class (for example, Gold), so you can see the profile for each environment.
- For a smaller environment, consider changing the table from listing data centers to listing clusters.

## Network Top Talkers Dashboard

Use the **Network Top Talkers** dashboard to monitor network demand in your IaaS. In a shared environment, a few VMs generating excessive activity can impact the entire data center. While a single VM might not cause a serious problem, a few of them can.

### Design Considerations

The **Network Top Talkers** dashboard helps you analyze how hard these VMs hit your IaaS. It classifies the workload into two: short bursts and sustained hits. A short burst lasts for a short period, maybe for a few minutes. A sustained hit can last for an hour and cause serious problems.

The **Network Top Talker** dashboard forms a pair with the **Storage Heavy Hitter** dashboard. To understand the IO demand in your environment, use both of them concurrently.

The **Network Top Talkers** dashboard displays sustained hits that last for an hour, as they can cause serious problems in a shared IaaS environment. You can identify the villain VMs and compare their demands with the capabilities of the underlying IaaS.

### How to Use the Dashboard

The dashboard shows the current workload. This is the total network load (received and transmitted) from all the vSphere environments monitored by VMware Aria Operations. The idea is to give you an indicator on how hard the overall load is.

- Select a data center from the data centers list.
  - The columns show the number of clusters, ESXi hosts, and VMs for each data center. The VM count includes the powered off VM. To only see the running VM count, edit the widget.
  - If you want to see information from all the data centers, select the vSphere world row.
  - Upon selection, the Total Demand Line chart and the Top Talkers tables fill up.
- Total Demand Line Chart
  - The total throughput (received and transmitted) in the selected data center.
  - Displays both, the five minute peak and the hourly average in one line chart. You can click the metric name to hide it.
- Top Talkers Table
  - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important. For example, an ESXi with 2 x10 GB port can theoretically handle 20 GB TX + 20 GB RX as its full duplex.

### Points to Note

- Understanding high demand helps you monitor IaaS and plan your capacity. IaaS provides four services, CPU, memory, disk, and network. While CPU, memory, and disk are bound, an active VM can consume all your network bandwidth, packet per second capacity, and the storage IOPS capcity. A VM with 4 vCPU and 16 GB memory cannot consume more than this amount, the same applies to disk space. A VM configured with 100 GB disk space cannot consume more than that.
- Network throughput, disk throughput, and disk IOPS can spike as their physical limits are very high per VM. This means that IaaS has enough capacity for all workloads and performs well until the VMs start consuming abnormally high amounts of network and disk bandwidth.

## Storage Heavy Hitters Dashboard

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard. To understand the IO demands in your environment, use both of them together. If you are using ethernet-based storage, storage traffic runs over the same physical network as your ethernet-based network traffic.

### Design Considerations

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard, so they share a consideration behind their design. For more information, see Network Top Talkers Dashboard.

### How to Use the Dashboard

- See the **Network Top Talkers** dashboard as they have the same design.
  - The main difference between **Storage Heavy Hitters** and **Network Top Talkers** is that the storage IO has two dimensions: IOPS and throughput.
  - Network IO does not have the IOPS dimension as the packet size is identical (1500 bytes being the standard packet, and 9000 bytes being the jumbo frames).
  - Storage IOPs and throughput are related, so use both to gain insight, they should display a similar pattern. If not, that indicates varying block sizes. For example, a throughput spike without an accompanying IOPs spike indicates large block sizes.
- Which VMs hit the storage the hardest.
  - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important, because different classes of SSD have different IOPS and throughput capabilities.

    After identifying the villain VM, talk to the VM owners if the numbers are excessive during peak hours and identify the reasons behind the excessive usage. You must ensure that they do not create a hot spot. For example, vSAN cluster with > 100 disk can handle numerous IOPS but if the VM objects are only on a few disks, those disks can become a hot spot.

### Points to Note

- Interpreting IOPs and throughput metrics depends on your underlying physical storage. For visibility into this hardware layer, add physical storage metrics to the dashboard.

## VM Contention Dashboard

The **VM Contention** dashboard is the primary dashboard for VM performance. It is designed for VMware administrators or architects. It can be used for both, monitoring, and troubleshooting. Once you determine that there is a performance issue, use the **VM Utilization** dashboard to see if the contention is caused by high utilization.

### Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

For understanding the performance concept of the selected counters and their thresholds, see the Performance Dashboards

## How to Use the Dashboard

- Select a data center from the data center table.
  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

    **Note** The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.
- The two bar charts are automatically shown.
  - Use them together to get an insight about your CPU readiness and your Memory contention analysis. Analyze how the cluster serves the VMs. For each VM, it picks the worst metric in the last 24 hours. By default, VMware Aria Operations collects data every 5 minutes, so this is the highest value among 288 datapoints. Once it has the value from each VM, the bar charts puts each VM in the respective performance buckets. The threshold in the buckets considers best practices, hence they are color coded.
  - For any critical environment, expect that all the VMs are served well by the IaaS. You must see green on both distribution charts. For development purposes, you can tolerate a small amount of contention in both CPU and Memory.
- VM Performance in selected Data Center.
  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
  - The table is sorted by KPI Breach columns, directing your attention to the VMs that are not served well by the IaaS.
  - The table shows the hostnames known by Windows or Linux. This is the name that the application team or VM owner knows, as they might not be familiar with the VM name.
  - The rest of the columns show performance counters. Because the goal is proactive monitoring, the counters are the worst and not the average, during the monitoring period. Because the operations context here is performance, not capacity, the table considers the last 24 hours only. Daily use is encouraged as any activity older than 24 hours is considered irrelevant from a performance troubleshooting viewpoint.
  - The column KPI Breach counts the number of SLA breaches in any given 5 minutes. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter

varies from 0–4, with 0 being the ideal. The value 4 indicates that all 4 IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, not an external SLA. Your internal threshold should be more stringent, so that you have a reaction time.

- Select a VM from the table.
  - All the health charts show the KPI of that VM.
  - The health charts display the last value, lowest value, and the peak value. Expect that the peak is within your threshold.

### Points to Note

- This dashboard uses Guest OS counters and VM counters appropriately. The two layers are distinct layers, and they each provide a unique visibility that the other layers might not give. For example, when the VMkernel de-schedules a VM as it has to process something else (for example, other VM, kernel interrupt). The Guest OS does not know the reason. In fact, it experiences frozen time for that particular vCPU running on the physical core and experiences time jumps when it is scheduled again.
- Guest OS counters logically require VMware Tools.
- The health chart is color coded. Change the settings if it does not suit your environment. If you are unsure of what suitable numbers to set for your environment, profile the metrics. The Guest OS Performance Profiling Dashboard dashboard provides an example of how to profile metrics.
- For a smaller environment with one or two data centers, change the filter from data center to cluster. Once you are list a cluster, you can then add the cluster performance (%) metric and sort them in an ascending order. This way the cluster that needs immediate attention is on the top.
- If you have a screen real estate, group the VMs by cluster or by ESXi host. This way, you can quickly see if the problem is in a particular cluster or ESXi host.
- Change the default timeline from one week to one day as and when required to suit your operations.
- If you navigate a lot to the **VM Utilization** dashboard from this dashboard, add a connection using the dashboard to dashboard navigation feature. For more details, see Dashboard Navigation Details.

## VM Performance Dashboard

Use the **VM Performance** dashboard to find out if a VM has a performance problem. As a first step, when a VM has a problem, verify if other VMs have the same problem. If the problem is widespread the root cause is not with the VM.

## How to Use the Dashboard

The **VM Performance** dashboard is organized into sections for ease of use.

- Select a data center from the **Datacenters** widget. To find out if there is a performance problem, what the problem is, and the extent of the problem, use the following three bar charts together: **Are VMs facing CPU Ready**, **Are VMs facing Memory Contention**, **Are VMs facing Disk Latency**. Each bar chart analyzes how the VMs are served by the cluster. These bar charts indicate if the VMs are waiting for CPU resources, facing memory contention, or disk latency. For each VM, it picks the worst metric in the last 24 hours. By default, VMware Aria Operations collects data every 5 minutes, so this is the highest value among 288 datapoints (12 x 24 = 288). Once it has the value from each VM, the bar charts put each VM in the respective performance buckets. The threshold in the buckets considers best practices, and hence they are color coded. For each bar chart, you can change the time period to the period of your interest. The maximum number is then displayed. The value is the worst 20-seconds, within the 5-minute collection time period. For your mission-critical environment, you must expect that all the VMs are being served well by the IaaS. If you see green on the distribution charts, you do not have to analyze further.

  For development, you may tolerate a small amount of contention in both CPU and Memory as you need to balance cost.

  You can also change the filter from data center to cluster. If you are listing clusters, you can then add the cluster performance (%) metric and sort them in ascending order. This way the cluster that needs immediate attention is on top.

  You can click on the bar to see the list of VMs under that performance bucket. From there, you can select a VM, and it's KPI is automatically displayed on the lower section of the dashboard.

- **Multiple VM Analysis**

  When you select a data center from the **VMs Performance in selected Datacenter** widget, the table listing all the VMs in the data center is displayed.

  The table is sorted by the KPI Breached column, directing your attention to the VMs that are not served well by the IaaS. The column counts the number of SLA breaches in any given 5-minute period. It is based on the counter `Performance \ Number of KPIs Breached`. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter varies from zero through four, with zero being the ideal. The value four indicates that all four IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, and not an external SLA.

  Because the goal is proactive monitoring, as opposed to reactive troubleshooting, the counters show the worst value instead of the average of the monitoring period.

- **Per VM Analysis**

  When you select a VM from the table, the CPU, memory, disk, and network performance charts are automatically displayed, each widget showing the KPIs of that VM.

- **Alerts**

  The relevant alerts are displayed automatically. You can view the settings by editing the widget, and adjust them accordingly to fit your operational needs.

- **Virtual Disks**

  A VM can have many disks, and it is possible that these disks may have different performance levels. The table lists the individual virtual disks and their contention and utilization metrics.

- **Configuration**

  The relevant configuration of the selected VM is displayed. You can customize as appropriate.

- **Relationship**

  From the VM, you can navigate to the parent cluster or datastore. Use the **Relationship** widget to navigate and auto select the associated cluster or datastore.

## VM Utilization Dashboard

The VMware administrator uses the **VM Utilization** dashboard with the **VM Contention** dashboard for managing performance.

### Design Considerations

Use the **VM Utilization** dashboard to identify virtual machines with a high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when a queue develops inside the Windows or Linux operating systems. By default, VMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

### How to Use the Dashboard

- Select a data center from the data center table.
  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

    **Note** The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- VM Peak CPU Usage (%).
  - There is no peak memory use as it is not applicable. Memory is a form of storage, for example consider a hard disk occupied space. A 90% utilization of the total space is not slower than 10%. This means that the issue is related to capacity issue and not performance.
  - The bar chart is color coded using five colors instead of four. The color gray is introduced to convey any wastage. Resources that are hardly utilized do not signify that the performance is at its peak. It can also mean the opposite. For example, if a VM needs 1+ vCPU, configuring it with 2 CPUs results in better performance instead of configuring it with 128 CPUs.
- VM Peak Utilization.
  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
  - The table focuses on peak utilization, because the context is performance and not capacity.
- Select a VM from the table.
  - All the health charts show the KPI of that VM.
  - Compliment the free memory with the memory IOPS or the memory throughput metric. The metrics in a gigabyte measure the space, and not the speed. Memory is a form of storage, so what you must measure is the rate, for example, read-write per second.

### Points to Note

- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the VM Contention Dashboard.

## Troubleshoot an Application Dashboard

The VMware Aria Operations Application Management Pack provides discovered applications to be managed in VMware Aria Operations. Using the **Troubleshoot an Application** dashboard, users can see the applications and the relevant metrics and alerts for the selected application. The dashboard also displays its relationship to the infrastructure. In the list of metrics, select a metric to see its trend over time.

## Cluster Contention Dashboard

The **Cluster Contention** dashboard is the primary dashboard for vSphere cluster performance. It is designed for VMware administrators or architects. It can be used for both, monitoring and troubleshooting. Once you determine that there is a performance issue, use the **Cluster Utilization** dashboard to see if the contention is caused by high utilization.

## Design Considerations

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

Utilization of the cluster is not shown in the **Cluster Contention** dashboard. You must separate the two concepts: utilization and contention. Performance and capacity are different concepts managed by two separate teams. Both CPU and memory are also shown separately. You can have a problem with one, without any issue in the other. CPU is more common as memory tends to have a lower overcommit ratio.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

## How to Use the Dashboard

- Average Cluster Performance (%).
  - This is the primary KPI for your entire IaaS. It plots how your IaaS is performing every 5 minutes, giving you the trend view of the overall performance.
  - The metric itself is simply the average of the Cluster KPI / Performance (%) metric. This performance metric in turn averages the VM Performance / Number of KPIs Breached metric from all the running VMs in the cluster. Hence a value of 100% indicates that every running VM in the cluster is served well.
  - As this KPI takes into account every running VM in your environment, the number should be steady. The analogy in real life is the stock market index. While individual stocks can be volatile, overall the index should be relatively steady on a 5 minutes by 5 minutes basis.
  - The relative movement of the metric is as important as the absolute value of the metric. Your absolute number might not be as high you want it to be, but if there are no complaints for a long time, then there is no urgent business justification to improve it.
- Clusters Performance.
  - It lists all the clusters, sorted by the least performing cluster in the last one week. You can change this time period.
  - The worst performance shows the lowest number in the time period. As VMware Aria Operations collects data every 5 minutes, there are 12 x 24 x 7 = 2016 data points in a week. This column shows the worst point among these 2016 datapoints.
  - A single number among 2016 datapoints can be an outlier that needs to be complemented with another number sometimes. A logical choice is the average of these numbers. For the average performance to be low, a lot of criterias have to be low. Waiting for the average causes a delay in your operations, and rise in complaints. For performance monitoring, the 95th percentile is a better summary than the average.
  - Your cluster should function at a 100% and perform its fuctions as planned.

- Select a cluster from the table.
  - All the health charts show the KPI of the selected cluster.
  - For performance, it is important to show both the depth and breadth of the performance problems. A problem that impacts one or two VMs requires a different troubleshooting than a problem that impacts all the VMs in the cluster.
  - The depth is shown by reporting the worst among any VM counter. So the highest value of VM CPU Ready, VM Memory contention, and VM Disk Latency among all the running VMs are shown. If the worst number is good, then you do not need to look at the rest of the VMs.
  - A large cluster with thousands of VMs can have a single VM experiencing poor performance while 99.9% of the VM population is fine. The depth counter might not report that most VMs are fine. It only reports the worst. This is where the breadth counters come in.
  - The breadth counters report the percentage of the VM population that is experiencing performance problem. The threshold is set to be stringent, as the goal is to provide early warning and activate proactive operations.

### Points to Note

It is possible for VMs in the cluster to suffer from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at individual consumer (VM). The following table shows various possible reasons.

| Event | Aware? |
|---|---|
| Power Management | No |
| HT | No |
| Ready | No |
| Co-Stop | No |
| System | No |
| Steal | No |
| IO Wait | No |
| Memory Wait | No |

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

Certain settings such as DRS automation level and the presence of many resource pools can impact performance. Consider adding a property widget to show the relevant property of a selected cluster, and a relationship widget to show resource pools.

For a large environment with many clusters, add a grouping to make the list more manageable. Group it by class of service, so you can focus more on the critical clusters.

## Cluster Performance Dashboard

The **Cluster Performance** dashboard combines the functionality of the **Cluster Contention**, **Cluster Utilization**, **ESXi Contention**, and **ESXi Utilization** dashboards.

### How to Use the Dashboard

The **Cluster Performance** dashboard is organized into sections for ease of use.

- **Overall Analysis**

  The **Average Cluster Performance** health chart is green when all the clusters are performing well. If the clusters are unable to serve the VMs well, all the clusters are no longer green, with a few occurrences of red.

  As the chart displays all the clusters, it uses the vSphere World object. This object is the parent of the vCenter Server object, and so it displays all the clusters from all the vCenter Servers.

  The metric used is `Performance \ Clusters Performance (%)` and is the primary KPI for your entire IaaS. It plots how your IaaS is performing every five minutes, giving you the trend view of overall performance.

- **Multi-Cloud Analysis**

  If the health chart is not green, and you want to find out which clusters are not performing, use the **vSphere Clusters** widget. The table lists all the clusters, starting with the cluster with the lowest performance. By default, the data displayed is from the last 24 hours. The **Worst Performance** column displays the lowest number in the time period. By default, VMware Aria Operations collects data every 5 minutes, so this is the lowest point among 288 datapoints (12 x 24 = 288).

  The **Worst Performance** column displays the lowest performance in the last period, specified under **Time Settings**.

- **Per-Cluster Analysis**

Select a cluster from the **vSphere Clusters** widget to see the trend over time. After you determine the cluster you want to investigate, review the five scoreboards: CPU, memory, disk, network, and others.

- **VM Shares**

  A common root cause for uneven performance problems is uneven shares. Each slice in the pie chart must correspond to a class of service. If the entire cluster is serving one class, then you should see a simple circle with no slices.

- **Resource Pool Analysis**

  Resource pool is another common reason behind uneven VM performances. A cluster with too many resource pools makes performance management difficult.

  The **Resource Pools in the cluster** provides a table listing all the resource pools.

- **ESXi Analysis**

  A cluster is a collection of ESXi hosts and the performance can be affected by uneven performance among the member hosts. You can drill down from a cluster to the ESXi hosts. The **ESXi Hosts in the Selected Cluster** widget lists all the ESXi hosts in the cluster, sorted by the worst performance in the last 24 hours. If the table displays values in green, you do not have to analyze further.

  You can change the time period to the period of your interest. The maximum number will be reflected accordingly.

  The table helps you quickly compare the performance of each ESXi. You can also see the performance over time, to see a trend.

  Certain settings such as power management and hyper threading can impact performance. The **ESXi Hosts in the Selected Cluster** widget displays the relevant property of a selected ESXi Host.

- **VM Analysis**

  When you select a cluster or ESXi from the **Running VMs in the selected Cluster or ESXi widget** the VMs that are running are automatically listed. Use this table to verify if the cluster or host performance problems were caused by VM configuration and usage. It is possible that the VM was not on the same host at the time of problem, due to vMotion.

  To drill down into a particular VM, select it and click the double arrow before the widget title.

- **Datastore Analysis**

  Use the **Shared datastores in the Cluster** widget to see a list of shared datastores accessible by hosts in the cluster. You can also drill down to the selected datastore.

## Cluster Utilization Dashboard

The VMware administrator uses the **Cluster Utilization** dashboard with the **Cluster Contention** dashboard for performance management.

## Design Considerations

This dashboard supports the **Cluster Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when VMs experience a contention. By default, VMware Aria Operations has a 5-minutes collection interval. For five minutes, there may be 300 seconds worth of data points. If a spike is experienced for a few seconds, it may not be visible if the remaining of the 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

## How to Use the Dashboard

- CPU(%) and Memory (%).
  - Review the CPU and Memory distribution charts for an overview of the CPU and memory utilization of the clusters.
  - The highest metric in the last one week is used. Average or 95th percentile is not used as this is utilization and not contention. High utilization does not mean bad performance.
  - One week is used instead of one day to give you a longer time horizon and covers the weekend. Adjust the timeline as you deem fit for your operations.
  - Expect memory to be higher than CPU, as it is a form of cache. The Memory Consumed counter is used as it is more appropriate than the Memory Active counter.
  - Low utilization can actually indicate bad performance, as not much of real work gets done. The chart uses the dark gray color for low utilization.
- Clusters Utilization.
  - The cluster utilization table lists all the clusters, sorted by the highest utilization in the last one week. If the table displays the green color, then there is no need to analyze further.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a cluster from the table.
  - All the utilization charts show the key utilization metrics of the selected cluster.
  - For memory, the high utilization counters are explicitly shown, Balloon, Compressed, and Swapped. Notice they exist even though utilization is not even at 90%, indicating high pressure in the past. If you look only at utilization, you might think you are safe.
  - The line charts show both average and highest among ESXi hosts in the cluster. The reason is unbalanced and it is not rare. There are many settings that can contribute to it (for example, DRS settings, VM Reservation, VM – Host Affinity, Resource Pool, Stretched Cluster, and Large VMs).
  - The disk IOPS is split into read and write to gain insight into the behavior. Some workload is read oriented, while others are write oriented.

- The disk throughput is not shown as it sums all the traffic. In reality, each ESXi host has its own limit.
- The vMotion line chart is added, as a high number of vMotion can indicate that the cluster load is volatile, assuming the DRS Automation level is not set to the most sensitive setting.

### Points to Note

- If your operations team have some forms of standardization that utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold.
- Consider adding a third distribution chart. Show the balloon counter in this third chart, as it complements the consumed counter. If there is no ballooning, a high consumed value is in fact better than a lower value.
- The workload metric can exceed a 100% because it is demand / usable capacity * 100. This can happen if you have four hosts in a cluster with each host running at 100% demand and admission control is set to 50%.
- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the Cluster Contention Dashboard.

## VM Rightsizing Dashboard

The **VM Rightsizing** dashboard helps you adjust the VM size for optimal performance at the lowest cost. It covers both undersized and oversized scenarios. The dashboard looks at the long term trend and covers both undersized and oversized scenarios. This dashboard is designed for the Capacity and the Operations teams, as rightsizing a VM helps in the day-to-day performance.

### How to Use the Dashboard

- **Overall Analysis**

  The scoreboard provides a summary of the total undersized and oversized CPU and memory.

  You can select either a data center or a cluster. In most cases, rightsizing analysis should be done at the cluster level as VMs typically do not move inter-cluster. The counters are displayed to provide better context. Focus on reclaiming VM capacity in a cluster that is low on capacity remaining.

  The distribution charts that display rightsizing are automatically displayed. Other than the bar charts, the **Undersized VM** and **Oversized VM** tables list the actual VMs.

- **VM Analysis**

  Select a VM to investigate further. The utilization is automatically displayed. VM usage reflects the amount of capacity consumed. This is based on the aggregate vCPU usage at five minute granularity, which provides a clear understanding of capacity used.

Rightsizing a VM can help improve performance for the VM as well as the cluster. Because of this, metrics that show performance bottlenecks, such as CPU Ready and CPU Run Queue, are provided to help you confirm how rightsizing may result in less contention and better performance.

Memory utilization is collected from the guest OS via VMTools. If guest OS metrics are not available, then the memory configured value is used instead. Rightsizing memory improves performance by reducing memory ballooning and contention. For example, VMs with over-provisioned memory are more likely to experience ballooning.

## Datastore Performance Dashboard

Use the **Datastore Performance** dashboard to view performance problems related to storage such as high latency, high outstanding IO, and low utilization. This dashboard is designed for both the VMware administrator and the Storage administrator, to foster a closer collaboration between the two teams. Local datastores are treated separately.

### Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

### How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

- **Overall Analysis**

  Select a data center from the **Datacenters** table. The three bar charts which are, **VM Performance**, **Read Performance**, and **Write Performance** provide an overall analysis of the datastore performance in a given vCenter Server data center or vSphere World. They work together to provide better insight. Just like other performance charts, the value displayed is the worst value during the time period. After you select a data center, if the Worst VM Disk Latency displays `No data to display`, it means that you have no observed latency issues for VM disk performance.

  The **VM Performance** chart displays the kind of latency and how many VMs experience that kind of latency. The **VM Performance** chart is your primary chart as it measures latency at the VM level. The **Read Performance** and **Write Performance** charts measure latency at the datastore level, which means they are the normalized average of all VMs in that datastore. Expect the **VM Performance** chart to be higher than the **Read Performance** and **Write Performance** charts. Read and write latency are displayed separately for better insight.

- **Datastore Analysis**

The **Datastores Performance** table automatically lists all the shared datastores in the data center or vSphere World. Both the worst (peak) performance and the 95th percentile are displayed. If the latter is close to the peak and it is also high, then it is a sustained problem. If the latter is low, then the problem is for a short duration. The table is color coded. Select a datastore that you want to troubleshoot. The relevant metrics and configuration are displayed.

- **VM Analysis**

  The list of VMs running in the selected datastore is displayed, with the relevant contention and utilization counters. Select the VM that you want to troubleshoot. The contention and utilization of the VM are automatically displayed. The number is at the VM level. If you suspect one of the virtual disks has high latency, use the counter Peak Virtual Disk Read Latency (ms) and Peak Virtual Disk Write Latency (ms).

- **Relationship**

  From the **Related Clusters and Hosts to selected Datastores** widget, select either an ESXi host, a vSphere cluster, or a vSAN cluster. The relevant contention and utilization counters are displayed.

## ESXi Contention Dashboard

The **ESXi Contention** dashboard is the primary dashboard for managing ESXi host performance. The VMware administrator or architect can use it to monitor and troubleshoot any performance issue. If you determine that there is a performance issue, use the **ESXI Utilization** dashboard to see if the cause for the contention is high utilization.

### Design Consideration

The **ESXi Contention** dashboard complements the Cluster Contention Dashboard, and shares the same design consideration.

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed to be used daily, hence the views are set to show data in the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

### How to Use the Dashboard

- ESXi CPU Performance and ESXI Memory Performance.
  - Review the two distribution charts for an overview of all the ESXi host's utilization and memory performance.
  - Both charts are using the percentage of VM facing performance counter and not the worst performance among VM counter because you are looking at the ESXi performance and not at the single VM performance. See how it handles all the VMs.

  - The bar chart is color coded. Keep the percentage of the VM population not being served under 10%.
- ESXi Hosts Performance.
  - The ESXi hosts performance table lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is displays the green color, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance greater than 24 hours are likely to be irrelevant.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select an ESXi host from the table.
  - All the health charts show the KPI of the selected cluster.
  - For performance, it is important to show both depth and breadth of a performance problem. A problem that impacts one or two VMs require a different troubleshooting than a problem that impacts all VMs in the cluster.
  - Worst CPU overlap among VMs in the host is included as it indicates a lot of interruptions. A running VM might get interrupted because the VMkernel needs the physical core to run something else. High and frequent numbers of interruptions are not healthy and can impact the VM performance.
  - Expect the network error to be 1% and dropped packet to be 0 most of the times, if not always. If it is not zero, analyze it to see if there are any patterns across all ESXi hosts, and bring it up with your network team.

### Points to Note

- Consider adding a third distribution chart and display the CPU co-stop counter in this third chart, as it complements the CPU ready counter. If your environment has relatively slow network and storage IO, you can add IO wait too.
- Unlike the **Cluster Performance** dashboard, there is no average ESXi hosts performance (%) at the vSphere World level. The reason is most ESXi hosts are part of a cluster and monitoring should be done at the cluster level.
- Certain settings such as power management and hyper threading can impact the performance. Consider adding a property widget to show relevant properties of a selected ESXi host.

## ESXi Utilization Dashboard

The VMware administrator uses the **ESXi Utilization** dashboard with the **ESXi Contention** dashboard to manage performance.

## Design Considerations

The **ESXi Utilization** dashboard supports the **ESXi Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, VMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

The dashboard complements the Cluster Utilization Dashboard dashboard, by providing the extra details. Hence it has a similar layout.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

## How to Use the Dashboard

- ESXi Hosts Utilization.
  - It lists all the ESXi hosts, sorted by the highest utilization in the last one week. If the table is all displaying the green color, then there is no need to analyze further.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select an ESXi host from the table.
  - All the utilization charts display the key utilization metrics of the selected cluster.
  - For memory, the high utilization counters are explicitly shown, for example balloon, compressed, or swapped. You might notice they exist even though utilization is not even at 90%, indicating that there was a high pressure in the past. If you look at only utilization, you might think you are safe.
  - The disk IOPS and the disk throughput are split into read and write to gain an insight into the behavior. Some workload is read oriented, while others are write oriented.
  - The network throughput is split into sent (transmit) and received to gain insight into the behavior. The total usage can be misleading because it sums up the send and receive traffic. In reality the network pipe is one for each direction (due to the full duplex nature of Ethernet), and not shared.

## Points to Note

If your operations team have some forms of standardization that the utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold. For more information, see the points to note in the ESXi Contention Dashboard.

# Network Performance Dashboard

Use the **Network Performance** dashboard to view performance problems related to network such as high latency, frequent retransmit, and many dropped packets. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

## Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The dashboard activates you to drill down from the distributed switch to the ESXi host and port groups in the switch, and then to the VM.

## How to Use the Dashboard

- Distributed Switches.
  - The distributed switches table lists all the switches, sorted by the highest packet dropped. The table splits the incoming traffic and the outgoing traffic for better analysis.
  - As the focus is on performance and not capacity, the throughput counters are not shown.
- Select a switch from the distributed switches table.
  - The health chart shows the dropped packet trend over time.
  - It does not narrow down the list of port groups automatically, as the list of port groups are always showing all the port groups in your environment.
  - If necessary, expand the two collapsed widgets. They show the network throughput and broadcast packets. Utilization is also shown so that you can correlate and understand whether the dropped packets are due to higher utilization.
- Port Groups and ESXi Hosts in the selected switch.
  - They get listed when you select a switch from the distributed switches table.
  - Just like the distributed switch, you can also see their relevant counts.
- If your environment has unused network switches, you can filter them out from this list, as this dashboard focuses only on performance.

## Points to Note

- vSphere network is by nature distributed. Each ESXi contributes to the physical NIC. This represents the physical capacity. Distributed switch and its port groups span across these independent network cards. This makes it harder to define and measure its performance. An unbalance can happen among ESXi hosts or physical NIC. In a sense, it is like distributed storage ( vSAN). Capacity management does not apply to a port group, since its upper limit (also known as the physical capacity) can vary by even a minute.

- Latency within a data center should be below 1 millisecond. Use VMware Aria Network Insight to study the latency or the retransmitting problems, caused by moving into the lateral traffic.
- Add a physical network using the appropriate management pack.

Most packets are unicast, between a pair of sender and receiver. If your environment has many VMs sending broadcast packets to everyone and multicast packets to many targets, add a Top-N widget to find out which VMs are sending these packets.

## vSAN Contention Dashboard

The **vSAN Contention** dashboard is the primary dashboard for managing vSAN performance. The VMware administrator or architect can use it to monitor and troubleshoot the vSAN cluster performance. If you determine that there is a performance issue, use the **vSAN Utilization** dashboard to see if the cause for the contention is high utilization.

### Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The **vSAN Contention** dashboard complements the ESXi Contention Dashboard, and shares the same design consideration. It focuses on the storage and vSAN specific metrics, and does not repeat what is already covered. It does not list any non vSAN cluster.

### How to Use the Dashboard

- vSAN Peak VM Latency, vSAN Peak CPU Ready, vSAN Peak Dropped Packet.
  - Review the three distribution charts for an overview of all the vSAN clusters performance.
  - The vSAN peak VM latency chart shows the distribution of disk latency experienced by all the VMs in the cluster. You should expect most of the VMs to experience latency that matches your expectation. For example, in an all flash systems, the VMs should not have >20 ms disk latency. If your vSAN environment is all flash, you must adjust the distribution bucket to a more stringent set.
  - The vSAN peak CPU ready chart shows if any of the vSAN kernel modules has to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not wait for CPU time. vSAN gets higher priority than VM World as it lives in the kernel space.
  - The vSAN peak dropped packet chart shows if any of the vSAN clusters are dropping packet in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.
- vSAN Clusters.
  - It lists all the vSAN clusters, sorted by the least performing.

  - It lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is showing all green, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance issues greater than 24 hours are likely to be irrelevant.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a vSAN cluster from the vSAN clusters table.
  - All the health charts show the KPI of the selected cluster.
  - If you are using SMART, the two heat maps at the bottom of the dashboard provide early warning.

## Points to Note

- A large vSAN cluster can have many components. Each of these components can have multiple performance metrics. The total number of KPI can reach hundreds of metrics. For example, take a 10 node cluster. It can have 530 counters to check. VMware Aria Operations aggregates them by introducing a set of KPIs. This analysis reduces the number to a more manageable number. The following table shows the KPIs and their formula.

| Name | What it is |
|---|---|
| Max Capacity Disk Latency (ms) | Highest latency among all capacity disks take the worst, not average, as the latency in a single capacity disk is already an average of all its VMs. If there are 50 VMs on the disk and 30 are issuing IO on it, then its average is among 30. |
| Min Disk Group Write Buffer Free (%) | Lowest free capacity among all the disk group write buffers. If this number is low, one of your buffers is not enough. While you want to maximize your cache, a low number is an early warning for capacity management. |
| Max Disk Group Read Cache/Write Buffer Latency (ms) | Each disk has a Read Cache Read Latency, Read Cache Write Latency (for writing into cache), Write Buffer Write Latency, and Write Buffer Read Latency (for de-staging purpose). This takes the highest among all these four numbers and the highest among all disk groups. It is the max of the max because each of the four datapoints is an average of all the VMs on it. |
| Sum Disk Group Errors | Sum of the bus reset + sum of commands canceled among all the disk groups. You must use sum and not get the max as each member should return zero. |
| Count Disk Group Congestion Above 60 | The number of disk groups congestion greater than 60. 60 is hardcoded in the vSAN Management Pack as it is a good starting point. As any congestion above 60 serves an early warning, count how many of such occurrences happen. |

| Name | What it is |
| --- | --- |
| Max Disk Group Congestion | The highest congestion among all disk groups. A high number indicates that at least one disk group is not performing. |
| Min Disk Group Capacity Free (%) | The lowest free capacity among all disk groups. A low space triggers rebalance. |
| Min Disk Group Read Cache Hit Rate (%) | The lowest hit rate among the disk group read cache. Ensure that this number is high as it indicates that the read is served by cache. |
| Sum vSAN PortGroup Packets Dropped (%) | Sum of all vSAN VMkernel port RX dropped packet + TX dropped packet. You should expect no dropped packet in your vSAN network. |

# vSAN File Services

The VMware administrator uses the **vSAN File Services** dashboard to monitor the file services running in their vSAN environment.

## Design Considerations

This dashboard is designed to complement the vSAN file services management provided by the vCenter Server. The vCenter Server is more of an administrative tool, while VMware Aria Operations is more of an operations tool. Each tool performs their specific functions and does not duplicate information.

## How to Use the Dashboard

- File Shares by Used Space and Latency.
  - Review the file shares by used space and latency heat map.
  - It shows all the file shares in your environment.
  - The greater the use (consumption), the greater the box, so you can easily see the most consumed ones.
  - The file shares are colored by latency. You must watch out for boxes with red color.
- vSAN Clusters with File Services activated.
  - It lists all the vSAN clusters with file services activated, giving a convenient view to see which clusters have these settings turned on.
- Select a vSAN cluster from the vSAN clusters with file services activated table.
  - The file servers in the selected vSAN cluster are shown. When you select a file server, it filters the file shares list to show the file shares in the selected file server.
  - The file shares in the selected vSAN cluster are shown. Selecting a file share displays all the relevant KPI on the file share.

### Points to Note

vSAN File Servers and vSAN File Shares are two new objects in VMware Aria Operations Management Pack for vSAN.

## vSAN Performance Dashboard

Use the **vSAN Performance** dashboard along with the **Cluster Capacity** dashboard.

### How to Use the Dashboard

The **vSAN Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**

  The **vSAN Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

  The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

  Select a cluster to investigate further. The VM latency distribution is automatically displayed.

- **Contention**

  You can view various disk-related contention counters of the cluster from the **Contention** widget.

- **Utilization**

  Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.

- **Disk Groups**

  You can drill down to the disk group level. All the counters are the worst value among the disk groups.

- **Read Cache**

  All the values displayed are the worst values among the read cache of the disk group.

- **Other KPIs**

  A problem in performance can also be caused by non-storage. vSAN resync is a type of utilization metric, but its presence can impact performance.

- **Disk Group Analysis**

  You can drill down to individual disk groups from the **Disk Groups** widget. Ensure that the disk groups are fairly balanced. Select the disk group that you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Cache Disks**

  You can drill down to cache disks from the **Cache Disks** widget. Ensure that the configuration is consistent. Select the cache disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Capacity Disks**

  You can drill down to capacity disks from the **Capacity Disks** widget. Ensure that the configuration is consistent. Select the capacity disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

## vSAN Utilization Dashboard

The VMware administrator uses the **vSAN Utilization** dashboard with the **vSAN Contention** dashboard to manage performance.

### Design Consideration

The **vSAN Utilization** dashboard supports the **vSAN Contention** dashboard. Use it to identify vSAN clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, VMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

### How to Use the Dashboard

- Clusters Utilization.
  - It lists all the vSAN clusters, sorted by the least performing.
- Select a vSAN cluster from the clusters utilization table.
  - All the health charts show the KPI of selected cluster.
- Disk Groups
  - It lists all the vSAN clusters, sorted by the least performing.
- Select a Disk Group from the disk groups table.
  - All the health charts show the KPI of selected cluster.

### Points to Note

- The **vSAN Utilization** dashboard complements the **vSAN Contention**. For more information, see the points to note in the vSAN Contention Dashboard.

## vSAN ESA Performance Dashboard

The vSAN ESA (Express Storage Architecture) Performance dashboard provides an overview of the performance issues related to your vSAN clusters. The dashboard displays the read and write latency issues that affect the vSAN clusters present in your environment. You can use this dashboard to identify and troubleshoot performance issues related to Contention and Utilization across VMs that are part of the vSAN cluster.

### How to Use the Dashboard

The **vSAN ESA Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**

  The **vSAN ESA Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

  The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

  Select a cluster to investigate further. The VM latency distribution is automatically displayed.

- **Contention**

  You can view various disk-related contention counters of the cluster from the **Contention** widget.

- **Utilization**

  Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.

- **Configuration**

  You can view the configuration details such as storage type, number of VMs, ESXi, Storage Pools, Physical Disks of the selected vSAN cluster.

- **Storage Pools**

You can view the configuration, contention, and utilization details of the Storage Pool for the selected vSAN cluster.

- **Physical Disks**

  You can view the configuration, contention, and utilization details of the Physical Disks for the selected vSAN cluster.

- **Related Objects to Selected vSAN Cluster**

  You can view the details of the related objects for the selected vSAN cluster.

## vSAN OSA Performance Dashboard

The vSAN OSA (Original Storage Architecture) Performance dashboard provides an overview of the performance issues related to your vSAN clusters. The dashboard displays the read and write latency issues that affect the vSAN clusters present in your environment. You can use this dashboard to identify and troubleshoot performance issues related to Contention and Utilization across VMs that are part of the vSAN cluster.

### How to Use the Dashboard

The **vSAN Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**

  The **vSAN Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

  The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

  Select a cluster to investigate further. The VM latency distribution is automatically displayed.

- **Contention**

  You can view various disk-related contention counters of the cluster from the **Contention** widget.

- **Disk Group Analysis**

  You can drill down to individual disk groups from the **Disk Groups** widget. Ensure that the disk groups are fairly balanced. Select the disk group that you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Utilization**

Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.

- **Configuration**

  You can view the configuration details such as storage type, number of VMs, ESXi, Storage Pools, Physical Disks of the selected vSAN cluster.

- **Compute and Network**

  You can view the compute and network metrics for the selected vSAN cluster.

- **Cache Disks**

  You can drill down to cache disks from the **Cache Disks** widget. Ensure that the configuration is consistent. Select the cache disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Capacity Disks**

  You can drill down to capacity disks from the **Capacity Disks** widget. Ensure that the configuration is consistent. Select the capacity disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Related Objects to Selected vSAN Cluster**

  You can view the details of the related objects for the selected vSAN cluster.

# Sustainability Dashboards

Using the Sustainability dashboards in VMware Aria Operations, you can view the reduction in carbon footprint achieved using virtualization, you can view clusters and data centers and compare them based on CO2 emissions or power consumption, you can identify idle VMs, and also identify old hardware, such as compute and storage.

The following Sustainability dashboards are added to the predefined VMware Aria Operations dashboards:

- Carbon Efficiency with Virtualization
- Carbon Transparency
- Environmental Impact of Idle VMs
- Green Supply

## Carbon Efficiency with Virtualization Dashboard

Use this dashboard to view the reduction in carbon footprint achieved using virtualization and also the reduced number of servers, reduced power consumption, and reduced carbon emission achieved using virtualization.

## How to Use the Dashboard

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- The **Greener Planet Contribution** widget provides details of power saved for the vSphere World object in percent and compares the power consumption of a vCenter Server environment before and after virtualization assuming that 100W is the power consumption for a low-range server before virtualization.
- Use the **CO2 Emissions Saved** and **Electricity Cost Saving** widgets to view the sum of carbon emissions and electricity saved for the last 24 hours.
- Using the **Power Consumption** widget, you can view power consumed before and after virtualization. Power consumption before virtualization is based on the assumption that a low-range server consumes 100W. Hence, power consumption before virtualization is calculated using the count of VMs, which is the number of physical servers if not virtualized, at 100W per server.
- Using the **CO2 Emission** widget, you can view CO2 emissions before and after virtualization. CO2 emission per kWh is assumed as 0.709 kg (Based on reference values from Greenhouse Gas Equivalencies Calculator). The value is calculated in Kilogram (kg).
- The **CO2 Emission Chart** and the **Power Consumption Chart** provides CO2 emissions (in kg) for the last 24 hours and power consumption (in kWh) for the last 24 hours respectively.

## Custom Values

- **Predefined Values**
  - Power consumption of a small server (1 socket, 10 cores, 32 GB RAM) = 0.1 kW (Assumption used in calculating power consumption before virtualization).
  - CO2 emission per kWh = 0.709 kg (Reference values from Greenhouse Gas Equivalencies Calculator).
  - Electricity Cost per kWh = $0.108 (Reference from VMware TCO Reference Calculator).
- You can add a custom property for CO2 emission and electricity cost which can be applied per cluster compute resource.

  Custom values of CO2 emission per kWh can be added to each cluster compute resource by creating a custom property with the name **CO2 Emission** and type as **Numeric** and the relevant custom values. You can add custom values for Cost of Power to each cluster compute resource by creating a custom property with name **Electricity Rate** and type as **Numeric** and the relevant custom values.

  **Note** If you do not define custom properties, out of the box values for CO2 emission and electricity costs are used.

### Points to Note

Your actual savings may be much more. The following are not included:

- Physical buildings and land. With virtualization, you consume less carbon foot print. This results in less physical rack.
- Network equipment. Fewer number of physical servers results in lesser network ports. Because firewall, load balancers, IDS, and IPS can be virtual machines, you have less equipment.
- Other components like UPS, lighting, cooling, and labor.

## Carbon Transparency Dashboard

Use this dashboard to view clusters and data centers and compare them based on CO2 emissions or power consumption. You can then identify the most green cluster to provision workloads. You can compare the power consumption of each compute component in the data center, showcase all the compute components with the lowest power consumption, compare physical data centers based on power consumption, and compare hardware models to find optimal power consumption.

### How to Use the Dashboard

As larger clusters consume more power than smaller clusters, the total power consumption cannot be used to determine a cluster to provision the next workload. Hence Power Efficiency is calculated based on the power consumption per GHz of CPU usage.

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- Use the **Top-10 Green Clusters by Power Efficiency** table to view clusters that are power efficient. Power efficiency is calculated as power usage per GHz of CPU Usage. Power efficiency alone cannot be used as the criteria to provision a VM, capacity must also be available in the cluster. View the `Capacity Remaining%` and `Time Remaining` information for each cluster to make an informed decision.
- Use the **Top-10 Green Clusters by Power Consumption** table to view clusters that consume less power. Use this table when you compare clusters based on the total power consumption. Larger clusters will consume more power than smaller clusters, but may still not be the most efficient cluster based on power efficiency. Power consumption alone cannot be used as the criteria to provision a VM, capacity must also be available in the cluster. View the `Capacity Remaining%` and `Time Remaining` information for each cluster to make an informed decision.
- The goal of the **Carbon Transparency** dashboard is to help you identify the greenest cluster to provision the next VM. However, the most green cluster may not have enough capacity to provision a VM, and hence you also need to check capacity before you decide on the greenest cluster to provision. Use relevant metrics for the selected cluster which can help

you identify a target cluster. The metrics are: Power Consumption of the Cluster (kWh), CPU Usage% in the cluster, and Memory usage% in the cluster. Use the associated widgets to view the metrics. The metrics are available for the last 30 days to help you identify the target green cluster to provision the next VM.

- Use the **Heatmap of vSphere Clusters based on Power Consumption** to view all the clusters based on Power Consumption in kWh. The size of each cluster in the heatmap is determined by the number of VMs in the cluster.
- When you add geo tags to clusters or physical data centers, those objects can be mapped in the **Geo** widget. Ensure that you add geo tags to the relevant objects.

## Environmental Impact of Idle VMs Dashboard

Use this dashboard to identify idle VMs in the data center, power off/delete idle VMs to reduce power consumption and improve data center efficiency, and identify remediation plans to offset the damage caused by those idle VMs which cannot be removed.

### How to Use the Dashboard

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- The **Cluster's Idle VM** table provides a list of clusters with reclaimable memory ,vCPUs, and disk space from idle VMs. You can also view the total resources that can be reclaimed from idle VMs. Click a cluster row to view the relevant sustainability metrics.
- Use the **Relevant Metrics of Selected Cluster** widget to view metrics of the selected cluster such as, power wasted by idle VMs (Wh), CO2 emissions by idle VMs (kg), trees to offset CO2 emission of idle VMs.
- The **Wasted Power from Idle VMs** heatmap allows you to view vSphere clusters configured to display wasted power from idle VMs. Power in watt-hour (Wh) is used to build the heatmap. A threshold value of 800 Wh changes the heatmap to red.
- The **CO2 Emission from Idle VMs** heatmap allows you to view vSphere clusters configured to display CO2 Emission from idle VMs. CO2 emission in kilogram (kg) is used to build the heatmap. A threshold value of 750 kg changes the heatmap to red.
- The **Trees Required to Compensate** heatmap allows you to view vSphere clusters configured to display the number of trees to be planted to compensate for the damage caused by idle VMs. A threshold value of 70 trees changes the heatmap to red

### Custom Values

- **Predefined Values**
  - Tree offset for CO2 Emission = 16.511 kg (36.4 pounds of carbon per tree. Refer to Greenhouse Gases Equivalencies Calculator) which is equivalent to 36.4/2.2046 kg of carbon per tree).

  - CO2 emission per kWh = 0.709 kg ( Reference values from Greenhouse Gas Equivalencies Calculator).

- Custom values for CO2 emission can be applied per cluster compute resource by adding a custom property.

  Custom values of CO2 emission per kWh can be added to each cluster compute resource by creating a custom property with the name **CO2 Emission** and type as **Numeric** and the relevant custom values. Custom values of tree offset to compensate CO2 emission can be added to each cluster compute resource by creating a custom property with the name **Trees to Offer** and type as **Numeric** and the relevant custom values.

**Note** If no custom properties are defined, out of the box values for CO2 emission are used.

## Green Supply Dashboard

Use this dashboard to identify old hardware, such as compute and storage, in the data center and replace them with new generation hardware components that are power efficient. You can also use this dashboard to reduce the overheads and buffers and identify smaller clusters that have higher overheads. The aim is to run with fewer overheads and buffer, without compromising on performance.

### How to Use the Dashboard

- Smaller clusters have a relatively higher overhead. A cluster with two nodes has 50% overhead, while a cluster with 10 nodes has only 10% overhead. Clusters with lesser capacity require more hosts and hence consume more electricity. The **Small Clusters** table lists clusters that meet one of the following criteria:

  - <=4 nodes
  - <=120 CPU cores and < 1 TB memory

  Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.

- Advancements in technology help ESXi hosts to deliver higher efficiency. ESXi hosts can deliver more CPU and memory capacity, often with low power requirements. The **Ageing Compute Hardware** table lists ESXi hosts that meet one of the following criteria:

  - ESXi version 6.0 or older
  - <=40 CPU cores and < 256 GB of memory

  Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.

- Just like compute hardware, newer storage hardware is more power efficient than older storage hardware. The **Ageing Storage Hardware** table lists datastores that meet the following criteria:

  - VMFS version 5 or older.

- Not a local datastore.

Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.

If you have goals that are different from those defined, you can modify the criteria of the widgets by updating the filters.

# Dashboard Library

## Executive Summary Dashboards

The requirements of the CIO, Head of Global Infrastructure, and IT Senior Management vary from the requirements of the technical teams. The **Executive Summary** dashboards provide an overall information on capacity and inventory in business terms.

These dashboards allow you to display problems related to budget and resource, and provide visibility to the senior management into the live environment. By doing this, you can prove the need for additional hardware. If there is wastage that has to be reclaimed, you can display where and how large the wastage is using these dashboards. VMware Aria Operations provides two example dashboards to get you started. As each executive might have a unique requirement or preference, the dashboards can be customized accordingly. The five principles displayed in the following figure are used to design the **Executive Summary**

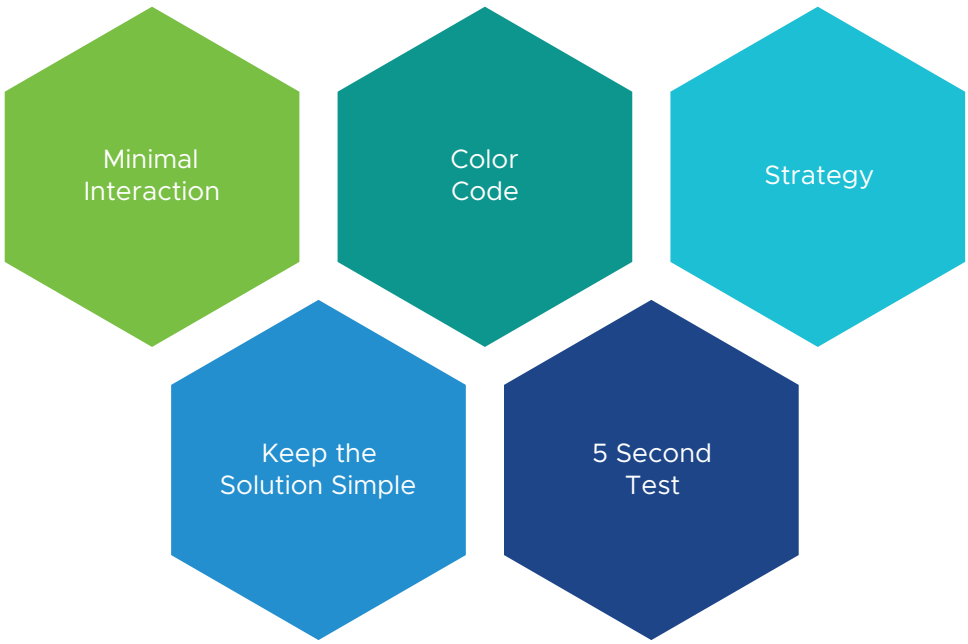


dashboards.

- Keeping the interaction, such as clicking, zooming, and sorting to a minimal.
- Use of color codes to have a user interface that is easy to understand.
- Each dashboard answers a specific question and the information is presented in business terms.

- Keep the solution simple and have a portal that is easy to access.
- Ensure that the dashboards are understood within five seconds.

## Capacity Summary Dashboard

The **Capacity Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Inventory Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

### Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

### How to Use the Dashboard

The **Capacity Summary** dashboard has two sections:

- The top section of the dashboard provides a summary at the vSphere World level.
  - The **VM Growth** widget displays the weekly average of the VM growth and provides holistic visibility of overall growth across all data centers for both running and powered off workloads. If an increase in the VM count is not accompanied by a corresponding increase in utilization, these newly provisioned VMs are likely not yet used.
  - The **Overcommit Ratio** widget highlights the efficiency gained by vSphere virtualization running multiple workloads on a shared infrastructure. Overcommitment has to be further reviewed along with elevated resource contention to understand the impact of performance on VMs competing for resources. In general, Overcommit is required to be financially more economical than the public cloud. As a reference, AWS typically overcommits CPU 2:1 by counting the hyper-threading and does not overcommit memory.

    **Note** VMware Aria Operations uses physical CPU Cores not Logical Cores (Hyper-threading) for all CPU-based capacity calculations.

- The bottom section of the dashboard activates drill down into individual compute or storage capacity.
  - Capacity is split into Compute (vSphere Clusters) and Storage (Datastores) views. The heat map displays capacity by size and color by time remaining. By selecting either Clusters or Datastores, you can further drill down understand the remaining capacity and time (in days).

### Points to Note

- Capacity remaining is not displayed at the vSphere World level as it can be misleading, especially in a global, or large infrastructure. Clusters also tend to serve a different purpose and they are not interchangeable.

- If you are using both on-prem and external cloud, for example, VMware on AWS, consider splitting the dashboard into two columns.

## Inventory Summary Dashboard

The **Inventory Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Capacity Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

### Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

### How to Use the Dashboard

- The **Summary** widget provides a quick view of the key inventory number.
  - The scoreboard is interactive. This widget drives the eight pie charts that are placed at the bottom of the dashboard. Since all the information is at the vSphere World level, clicking any of them will display details of the total inventory.
- Select any data center from the **Datacenters** widget.
  - This widget drives Clusters and Datastores so that you can quickly view what you have in a given data center and related capacity.
  - For a small environment, the vSphere World is displayed so you can view all the VMs in the environment.
  - To sort by any of the columns in the table, click the column title.
- The eight charts in the dashboard provide details of the inventory. They are driven by **Datacenters**, **Compute**, **Storage**, and **Summary** widgets.

### Points to Note

- Understand the relationship hierarchy in vSphere. For example, Compute (cluster) is not a parent of Storage (datastore), so logically it is not possible to display datastores in a cluster. Data center consists of compute (cluster), network (distributed switch), and storage (datastore).
- Datastores do not drive the pie chart. This is a known limitation in the View widget.
- If your senior management wants to view the largest VM in a given environment, add a Top-N widget to list the top 10 largest consumers so that CPU, memory, and the disk details are highlighted.

# Network Operation Center

A dashboard projected on the large screen serves a different business purpose than a dashboard on your laptop or desktop. It is placed strategically because it displays a time sensitive

information. Dashboards complement alerts and cannot replace it. The five principles displayed in the following figure are used to design the predefined **Network Operation Center** dashboards.

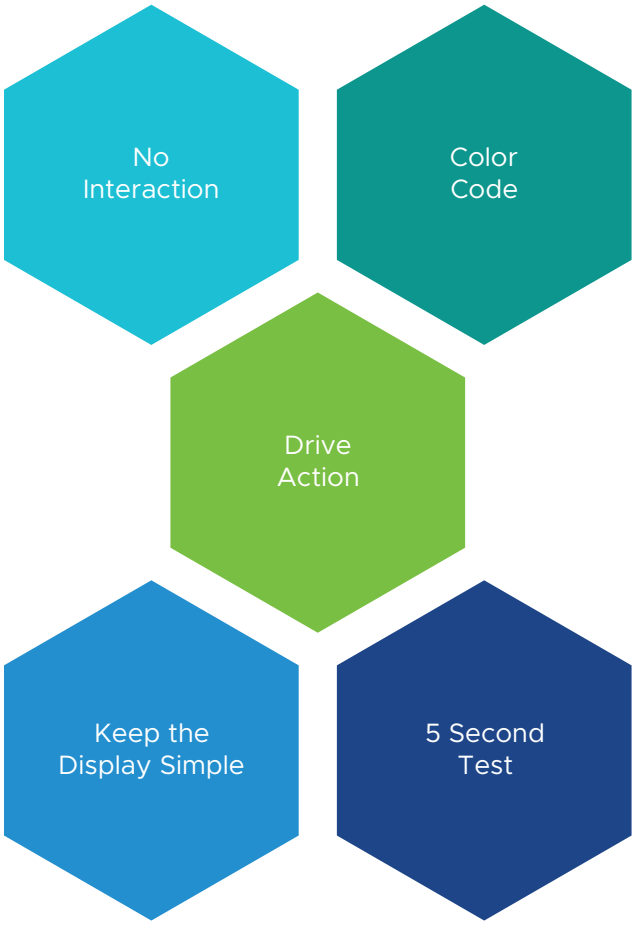


- Keeping the interaction, such as clicking, zooming, and sorting to a minimal. Avoid having buttons, use of mouse or keyboard to view data.
- Use of color codes to have a user interface that is easy to understand.
- Displaying content that drives action. Display of live information as the focus is on immediate remediation. Problems that need immediate actions are displayed, for example, stop provisioning of new VM or take action on VMs that abuse the shared infrastructure.
- Display of problems that do not require immediate attention are avoided, for example, increase supply of infrastructure, such as adding hardware.
- Keep the display simple and have a portal that is easy to access.
- Dashboards are designed to display minimal and critical information only.
- Displays of numbers in percentage, with 0% being poor and 100% being perfect. To display utilisation, you can use the following markers:
  - 50% indicates good and balanced utilization. However, the ideal value is 75%
  - 0% indicates wastage

  - 100% indicates high utilization
- Ensure that the dashboards are understood within five seconds.

## Live! Cluster Performance Dashboard

The **Live! Cluster Performance** dashboard provides live information on whether the requests of the VMs are met by their underlying compute clusters. This dashboard focuses on CPU, Memory, and the performance of the clusters. Use this dashboard to view if there is any problem in meeting the demands of the VMs and if there is any unbalance within a cluster. The **Live! Cluster Performance** dashboard is the primary dashboard and it complements the **Live! Cluster Performance** dashboard which is the secondary dashboard. This secondary dashboard displays if the performance problem is caused by high utilization. The primary dashboard answers the question 'Is our IaaS performing?', while the secondary dashboard answers the question 'Is our IaaS working hard?'.

### Design Consideration

The **Live! Cluster Performance** dashboard displays three heat maps. The heat maps complement each other and must be used together. The location of each cluster and ESXi hosts within those clusters is identical in all heat maps. The fixed positioning allows you to compare if the problem is caused by memory contention, CPU ready, or CPU co-stop.

The sizes of each cluster and ESXi hosts are constant. Variable sizing creates a distraction and can result in small boxes, making it difficult to read.

The focus of the performance is on the population and not on a single VM. This is not a single VM troubleshooting dashboard but a dashboard focusing on infra problem. As the infra counter is mathematically an aggregation of VM counters, you must have a right roll-up strategy. As the goal is to provide an early warning, do not use the average as a roll-up technique. Use the percentage of the population exceeding a threshold. The threshold is set to be stringent to receive an early warning.

### How to Use the Dashboard

Review the heat maps, **Memory Contention**, **CPU Ready**, and **CPU Co-Stop** and see if there is any color other than green.

- Green indicates that almost 100% of the VMs have received the CPU and memory that was requested. The threshold is set such that if the 10% of the VM population does not receive the requested resources, then the heat map turns red.
- Red indicates an early warning. Stringent thresholds are used to activate proactive attention and remediation operations. The heat map can turn red because of the high standard that is applied even when there is no complaint from the VM owner yet.
- The light gray indicates that there is no VM running on the host and the metric is not computing.

View if there is any unbalance.

- There are two types of unbalance, cluster unbalance, and resource type unbalance.

- The ESXi hosts are grouped by the cluster, so that the unbalance within a cluster can be easily viewed. Cluster unbalance is a real possibility and it is best monitored and not just assumed.
- If the three heat maps are different, then there is a resource unbalance. For example, if the memory contention is mostly red, but the two CPU heat maps are green, it means you have an unbalance between memory and CPU.
- If a single ESXi host displays different color across the three heat maps, it indicates that there is an unbalance between the CPU and memory resources in the host.

For NOC Operator, drill-down by selecting one of the VMs on the heat map.

- The **Trends of Selected ESXi Host** widget will automatically display the performance counters. To hide any metric, click the name in the legend.

As part of the deployment, configure auto-rotate among the NOC dashboards. If you want to view one dashboard, then you can remove the VMware Aria Operations menu by using the URL sharing feature. This makes the overall user interface presentable and allows you to focus on the dashboard.

### Points to Note

- You can add Disk Latency if you have the screen real estate. Use the counter 'Percentage of Consumers facing Disk Latency (%)'. It is a part of a datastore object, not a cluster, as a VM in a cluster can have disks across multiple datastores. Organize this storage performance by data center and not by the cluster.

## Live! Cluster Utilization Dashboard

The **Live! Cluster Utilization** dashboard complements the **Cluster Performance** dashboard. Use this dashboard to view the clusters that are working excessively and are close to their physical limit. This dashboard displays ESXi hosts that have CPU or memory saturation that can lead to performance issues for the VMs running on the host.

### Design Considerations

This dashboard is designed to complement the **Live! Cluster Performance** dashboard and it shares design considerations.

### How to Use the Dashboard

As this dashboard has an identical design with the **Live! Cluster Performance** dashboard, it has the same usage procedure. Unlike the heat maps in the **Live! Cluster Performance** dashboard, the three heat maps in this dashboard have a different scale, reflecting the different nature of the counters.

Logically, memory is a form of storage. It acts as a cache to disk as it is much faster. A high utilization is better, as it indicates that more data is being cached. The ideal situation is when ESXi host Consumed metric is red but ESXi host Ballooned metric is green. When Ballooned is red and Consumed is gray, it means that there was high pressure in the past but it is not there anymore. The reason the ballooned stays red is because the ballooned pages were never requested back.

The ballooned memory counter was selected over the swapped or compressed memory counters as it is a better leading indicator. Since all three can co-exist at the same time, they are displayed in the line chart. Ballooned is displayed in absolute amount and not as a percentage, because the higher the size, higher are the chances for it to impact a VM. If you feel using percentage is easier for your operations, create a super metric to translate the value.

The heat map displays Wastage by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

Analyze if the ESXi host is contributing. A light gray box indicates that the host is a part of the cluster but there is no utilization. It is possible for the host to be in the maintenance mode or is powered off.

### Points to Note

- ESXi host chooses to swap over compression if the compression ratio is less than 4x.
- If the ESXi host's physical NIC is saturated in your environment, then you can add a Network Throughput heat map.

## Live! Heavy Hitters Dashboard

The **Live! Heavy Hitters** dashboard helps you analyze the misuse of the shared infrastructure. This dashboard displays details of VMs misusing shared infrastructure and if that has caused performance problems to the other VMs. The shared infrastructure includes risks. The cause for excessive load might be attacks, for example, the denial of service, process runaway, or a mass activation of agents. The most demanding VM is the largest. If a handful of VMs is dominating the shared infrastructure, their collective size is displayed on the dashboard.

### Design Considerations

See the Performance Dashboards page for common design considerations among all the dashboards for performance management.

In a shared environment, it is possible to have a victim-villain problem. In the heat map, the villain VM is the one with the largest box size, while the victim VM is the one with the red box. If a handful of VMs is dominating the shared infrastructure, their collective size will be highly visible on the dashboard.

### How to Use the Dashboard

- The heat maps, Disk IOPS, Disk Throughput, Network Throughput, and CPU Demand displays the four different loads that can be excessive. The heat maps display the relative value and not the absolute value. A VM does not generate a high load in the absolute term just because it has a large configuration.
- Each heat map has its color threshold, reflecting the nature of the contention metrics used in each of them.
- For NOC Operator, drill-down by selecting one of the VMs on the heat map. All the four line charts are automatically displayed, enabling you to get a complete picture of the selected VM.

### Points to Note

- Memory is not displayed as it is a form of storage. The memory counters are space utilization and not speed. Think of disk space instead of IOPS. It can cause a capacity problem on the shared ESXi host, but not performance problems to other VMs.
- In a large environment, it might be difficult to view a small victim VM. Consider having multiple dashboards and use them interchangeably.

## Software Defined Wide Area Network Dashboard

The Software-Defined Wide Area Network (SD-WAN) dashboard allows you to configure and monitor the services related to VeloCloud and SD-WAN using VMware Aria Operations. Using the SD-WAN dashboard, you can also collect the metrics for VeloCloud Orchestrator and VeloCloud Gateway.

By default the SD-WAN dashboards are deactivated, if you want to know how to activate them, see Manage Dashboards. The following services are discovered using the VeloCloud Orchestrator:

- Java Application
- VeloCloud Orchestrator
- Nginx
- ClickHouse
- MySQL
- Redis
- Network Time Protocol

The following services are discovered using VeloCloud Gateway:

- Network Time Protocol
- VeloCloud Gateway

## Troubleshoot SD-WAN Dashboard

You can use the widgets in Troubleshoot SD-WAN dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN.

You can use the dashboard widget in several ways:

- **Troubleshoot Virtual Machine (VM)**: Use this widget to navigate to a specific VM and troubleshoot the issues.
- **Troubleshoot Orchestrator**: Use this widget to navigate to a specific orchestrator and troubleshoot the issues.
- **Troubleshoot Gateway**: Use this widget to navigate to a specific gateway and troubleshoot the issues.
- **Troubleshoot Application**: Use this widget to navigate to a specific application and troubleshoot the issues.
- **Relationship**: Use this widget to view the services and operating system associated with the VeloCloud Orchestrator.
- **Top Alerts**: Use this widget to view the top alerts associated with the SD-WAN.

## Troubleshoot SD-WAN Gateway Dashboard

You can use the widgets in Troubleshoot SD-WAN Gateway dashboard to monitor and troubleshoot all the services and applications associated with the SD-WAN gateway.

You can use the dashboard widget in several ways:

- **Active Alerts on the Gateway**: Use this widget to view the active alerts for the gateway.
- **Health of Gateway Applications**: Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System**: Use this widget to examine the operating system status.
- **Gateway Summary Status**: Use this widget to view the summary information for the gateway.
- **Gateway Process Status**: Use this widget to view the process information for the gateway.
- **Gateway Resource Metrics**: Use this widget to view the resource metrics associated with the gateway.
- **Parent Host**: Use this widget to view the parent host information.
- **Parent Cluster**: Use this widget to view the parent cluster information.

## Troubleshoot SD-WAN Orchestrator Dashboard

You can use the widgets in Troubleshoot SD-WAN Orchestrator dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN Orchestrator.

You can use the dashboard widget in several ways:

- **Active Alerts on the Orchestrator**: Use this widget to view the active alerts for the Orchestrator.
- **Health of Orchestrator Applications**: Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System**: Use this widget to examine the operating system status.
- **Examing MySQL**: Use this widget to examine the MySQL application.
- **Orchestrator Service Status**: Use this widget to view the service status of the Orchestrator.
- **Redis Status**: Use this widget to view the status of the Redis application.
- **API Check Status**: Use this widget to check the API status.
- **Nginx Status**: Use this widget to check the Nginx status.
- **Parent Host**: Use this widget to view the parent host information.
- **Parent Cluster**: Use this widget to view the parent cluster information.

# VMware Aria Automation Dashboards

The VMware Aria Automation dashboards allow you to track performance, health, utilization, and availability attributes of deployments made to SDDC clouds and public cloud end points. You can also monitor the virtual machines and track the utilization and performance of the VMs in your SDDCs and public cloud accounts..

## VMware Aria Automation Predefined Dashboards

The following VMware Aria Automation dashboards are added to the predefined VMware Aria Operations dashboards:

- Cloud Automation Environment Overview
- Cloud Automation Project Cost Overview
- Cloud Automation Resource Consumption Overview
- Cloud Automation Deployment Overview
- Cloud Automation Top-N Dashboard

## Cloud Automation Environment Overview

You can use the widgets in the Cloud Automation Environment Overview dashboard to view the environment details for the vCenter Cloud Zone objects. You can use the Cloud Automation Environment Overview dashboard to view the projects, deployments associated with the vCenter Cloud accounts.

You can use the dashboard widgets in several ways.

- **Environment Summary**: Use this widget to view the SDDC cloud zones, public cloud zones, projects, deployments, blueprints, SDDC VMs and public cloud VM details for the cloud accounts present in your environment.
- **SDDC Cloud Zones**: Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the SDDC cloud zone objects present in your environment.
- **Public Cloud Zones**: Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the public cloud zone objects present in your environment.
- **Project List**: Use this widget to view the total blueprints, cloud zones, deployments, virtual machines, health, risk, efficiency details in your environment.
- **Top Alerts**: Use this widget to view the top alerts in your environment.
- **VM List**: Use this widget to view all the VM details in your environment.
- **Blueprint List**: Use this widget to view the blueprint objects in your environment.
- **Deployment List**: Use this widget to view the blueprint objects deployed in your environment.
- **SDDC Virtual Machines**: Use this widget to view the SDDC VM's resource details.
- **Public Cloud Resources**: Use this widget to view the public cloud resource details.

## Cloud Automation SDDC Project Price Overview

You can use the cloud automation SDDC project price overview dashboard to view the project price details associated with each of VMware Cloud on AWS SDDC and public cloud accounts. With the project price overview dashboard, you can view price details for individual projects and find out deployments with the highest cost. .

You can use the dashboard widgets in several ways.

- **Project Cost**: Use this widget to view the project wise cost for compute, storage, and additional resources associated with your cloud environment.
- **Daily Price Over Time**: Use this widget to view the price of individual projects on a day to day basis.
- **Relationship**: Use this widget to view the relationship between the objects and the projects present in your cloud environment.
- **Deployment Price by Selected Project**: Use this widget to view the deployment cost for the selected project in your cloud environment.
- **Deployment With Highest Cost**: Use this widget to view the highest cost associated with projects that are present in your cloud environment.

## Cloud Automation SDDC Resource Consumption Overview

You can use the widgets in the Cloud Automation SDDC Resource Consumption Overview dashboard to view the resources consumed by VMware Aria Automation on Cloud Accounts.

You can use the Cloud Automation Resource Consumption Overview dashboard widgets in several ways.

- **Cloud Account**: Use this widget to view all the attributes related to the cloud account.
- **SDDC Cloud Zone**: Use this widget to view all the attributes related to the SDDC cloud zones.
- **Project**: Use this widget to view all the project details associated with your cloud account.
- **Cluster List**: Use this widget to view all the details associated with the clusters in your account.
- **Cluster Utilization**: Use this widget to view the cluster utilization details for the cloud accounts.
- **Deployment Heat Map by Project**: Use this widget to view the heat map for each deployed project in your cloud environment.
- **SDDC Cloud Zone Capacity**: Use this widget to view the memory and storage capacity that is allocated, reserved, and free for each cloud zone object.
- **SDDC Cloud Zone Memory Trend**: Use this widget to view and analyze a seven-day trend for the memory allocated, reserved, and free for the cloud zone.
- **SDDC Cloud Zone Storage Trend**: Use this widget to view and analyze a seven-day trend for the storage allocated, reserved, and free for the cloud zone.

## Cloud Automation Deployment Overview

You can use the cloud automation deployment overview dashboard to view the deployment details associated with your cloud environment. You can view details about cloud accounts, cloud zones, projects, and deployments. The dashboard also provides details about the deployment resources and the relationship between deployments and their objects.

You can use the dashboard widgets in several ways.

- **Cloud Account**: Use this widget to know the cloud account details of the selected account.
- **Cloud Zone**: Use this widget to know the cloud zone details like adapter, collection state, policy and object type of the selected account.
- **Project**: Use this widget to know the project details of the selected project.
- **Deployments**: Use this widget to know the deployment details of the cloud accounts in your environment.
- **Deployment Resources**: Use this widget to know the deployment resource details across the cloud accounts in your environment.
- **Object Relationship (Advanced)**: Use this widget to view the relationship between the objects and deployments present in your cloud environment.
- **Deployment Heat Map by Project**: Use this widget to view the heat map for each deployed project in your cloud environment.

- **Deployment Heat Map by Blueprint**: Use this widget to view the heat map for each deployed blueprint in your cloud environment.

## Cloud Automation Top-N Dashboard

You can use the widgets in the Cloud Automation Top-N dashboard to view the projects with most critical alerts, to view the blueprint with most deployments, and to view the deployments with the highest cost.

You can use the dashboard widgets in several ways.

- **Project with Most Critical Alerts**: Use this widget to view the projects which has most critical alerts.
- **Top Alerts**: Use this widget to view the top alerts for the projects in your cloud account.
- **Blueprints with Most Deployment**: Use this widget to view the blueprint which has maximum deployments for the cloud account.
- **Relationship**: Use this widget to analyze the relationship between blueprints and deployments, and deployment and cost.

# VMware Aria Operations Dashboards

With VMware Aria Operations dashboards you can monitor and troubleshoot objects in your cloud infrastructure.

The following vVMware Aria Operations dashboards are added to the predefined VMware Aria Operations dashboards:

- VMware Aria Operations Cloud Billing
- VMware Aria Operations Cloud Universal Billing

## VMware Aria Operations Cloud Billing

The VMware Aria Operations Cloud Billing dashboard provides you object billing details of Operating System Instance (OSI) used in your cloud environment.

### How to Use the Dashboard

- OSIs and Billable Objects widget provides the total count of OSIs and billable objects. You have to update these widgets depending on your subscription limits.
- OSIs Across Object Types widget provides distribution of OSIs across different object types.
- Billable Object Types List widget provides a list of all the object types which are managed by VMware Aria Operations and consume license unit(s).
- OSIs Consumption Across Object Types displays a heat map and maps the magnitude of OSIs consumption for different object types with the relevant heat map colors. The widget also interacts with OSIs Consumption Over Time and displays how OSIs count have been changing for a given object type over a period of time.

- List of Objects widget displays the object details like name, adapter type, object type, policy, creation time, collection state, and collection status. Use the filter option to filter different objects from the list of objects.

### How to Edit OSIs and Billable Objects Widget To Set Correct Color Codes

1 Click the Edit icon at the top right corner of the widget.

2 Go to **Output Data** section.

3 Double click the row which has column heading Yellow, Orange, or Red.

4 Set the **Color Method** to custom.

5 Enter appropriate values as per your subscription limits.

## VMware Aria Operations Cloud Universal Billing

The VMware Aria Operations Cloud Universal Billing dashboard provides the object billing details based on CPU usage.

### How to Use the Dashboard

- CPUs and Billable Objects widgets provides the total count of CPUs and Billable Objects. You have to update these widgets based upon your subscription limits.
- CPUs Across Object Types widget provides distribution of CPUs across different object types.
- Billable Object Types List widget provides a list of all the object types which are actively managed by VMware Aria Operations and consume license unit(s).
- CPUs Consumption Across Object Types displays a heatmap and maps the magnitude of CPUs consumption for different object types with heatmap colors. The widget also interacts with CPUs Consumption Over Time and displays how CPUs count has been changing for a given object type over a period of time.

### How to Edit CPUs and Billable Objects Widget To Set Correct Color Codes

1 Click the Edit icon at the top right corner of the widget.

2 Go to **Output Data** section.

3 Double click the row which has column heading Yellow, Orange, or Red.

4 Set the **Color Method** to custom.

5 Enter appropriate values as per your subscription limits.

# Service Discovery Dashboards

Using the service discovery dashboards, you can determine the inter-dependencies of virtual machines and the dependencies of each service in the respective virtual machines.

The following service discovery dashboards are added to the predefined VMware Aria Operations dashboards:

- Service Distribution
- Service Relationships
- Service Visibility
- Virtual Machine Relationships

## Service Distribution Dashboard

You can use the dashboard to view the distribution of different services in the selected data center, cluster, or a host system. You can also view known and unknown services including the category and distribution percentage across a vSphere resource.

You can use the dashboard widgets in several ways:

- **Inventory Item**: Use this widget to view a hierarchical representation of objects in the form of badges.
- **Known Services Distribution**: Use this widget to view different services discovered from a selected object.
- **Service Categories**: Use this widget to view the service categories that are discovered by selecting an object from the resource widget.
- **User Defined Services Distribution**: Use this widget to view a list of user-defined services.

## Service Relationships Dashboard

You can use the dashboard to view properties of the service such as the install path, the ports used, and the version. You can also view the relationship between the services that run on other VMs.

You can use the dashboard widgets in several ways:

- **List of Services Discovered**: Use this widget to view the services that have been discovered.
- **Connections from the Selected Services**: Use this widget to view the relationship between the services and the other services running on the VMs.
- **Properties of the Selected Service**: Use this widget to view the properties of the selected services.

## Service Visibility Dashboard

You can use the dashboard to view a list of VMs without service visibility and VMs with user-defined services after you select a vSphere object.

You can use the dashboard widgets in several ways:

- **Inventory Tree**: Use this widget to view a hierarchical representation of objects in the form of badges.

- **Virtual Machines without Service Visibility**: Use this widget to view information about services where discovery has failed.
- **Virtual Machines with User-Defined Services**: Use this widget to view a list of VMs where the user has defined such services.

## Virtual Machine Relationships Dashboard

You can use the dashboard to view a list of VMs with service discovery details such as, status, method, incoming/outgoing connections, and protection groups. When you select a VM, the dashboard displays a list of discovered services on the VM, the relationships of the VMs with other VMs based on the relationships of the discovered service.

You can use the dashboard widgets in several ways:

- **List of virtual machines**: Use this widget to view all the VMs discovered by the vCenter Server.
- **Node relationship of the selected VM**: Use this widget to view the relationship between the objects.
- **List of Services running in the selected VM**: Use this widget to view all the properties of the selected VM.
- **Connections of Virtual Machines**: Use this widget to view the relationship between one or more VMs.

# Inventory Dashboards

The three vSphere Inventory dashboards and workload management inventory dashboards cater to the compute, network, and storage aspects of your SDDC. Using these dashboards, you can navigate through the environment and view your inventory and their key metrics at a glance. The Network and Storage dashboards can be shared with the network and storage teams respectively, giving them the necessary visibility, and increasing the collaboration between teams.

## vSphere inventory dashboards

The vSphere inventory dashboards are built specifically for each role, but they share a common design. They have a similar layout and are used in the same manner. This makes learning easier, especially in smaller environments where the same team manages the full environment.

These dashboards help you answer several key questions:

- What is the topology of your vSphere compute inventory?
- What is the topology of your vSphere storage inventory?
- What is the topology of your vSphere network inventory?

## Workload Management Inventory Dashboard

This is a unified dashboard for the new workload management objects. It shows the relationships and KPIs for the workload management objects. For example, you can see the topology view from the Tanzu Kubernetes clusters to the physical infrastructure.

## vSphere Compute Inventory Dashboard

You can use the vSphere Compute Inventory Dashboard to browse through the topology of your vSphere compute inventory which includes information related to vSphere world, vCenter Server, data center, clusters, hosts, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the clusters, ESXi hosts, and virtual machines associated with the object.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to an object in the environment.
- **Metrics**: View the metrics related to the object.
- **Clusters**: View the cluster functionality.
- **ESXi Hosts**: View the data related to the hosts.
- **Virtual Machines**: View VMs that belong to the object.

## vSphere Network Inventory Dashboard

The vSphere Network Inventory Dashboard allows you to browse through the topology of your vSphere network inventory which includes information related to vSphere world, vCenter Server, data center, distributed vSwitches, distributed port groups, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the distributed vSwitches, distributed port groups, virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Distributed vSwitches**: View details related to the distributed vSwitches.
- **Distributed Port Groups**: View data relevant to distributed port groups.
- **Virtual Machines**: View VMs that belong to the object.

## vSphere Storage Inventory Dashboard

The vSphere Storage Inventory dashboard allows you to browse through the topology of your vSphere storage inventory which includes information related to vSphere world, vCenter Server, data center, datastore clusters, datastores, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the datastore clusters, datastores, and virtual machines associated with it.

You can use the dashboard widgets in several ways.

- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Datastore Clusters**: View the datastore cluster functionality.
- **Datastores**: View the datastore functionality.
- **Virtual Machines**: View VMs that belong to the object.

## Workload Management Inventory Dashboard

The Workload Management Inventory dashboards curates the Kubernetes inventory across all the Workload Management activated vSphere environments and displays it here. This includes an end to end topology map showcasing the health of all the objectes along with upstream and downstream dependencies. Upon clicking any object in the relationship tree, the related inventorty of Supervisor Clusters, Namespaces, Pods, Developer Managed VMs and Tanzu Kubernetes clusters can be viewed and exported from this dashboard.

You can select an object type to view the properties and key metrics related to it.

You can use the dashboard widgets in several ways.

- **Environment Summary**: Provides a summary of the supervisor cluster and the child objects.
- **Relationships**: An interactive canvas where you can view the relationship between the different objects in the workload management inventory.
- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Supervisor Clusters**: View the supervisor cluster functionality.
- **Tanzu Kubernetes cluster**: View the Tanzu Kubernetes cluster functionality.
- **Virtual Machines**: View VMs that belong to the object.
- **vSphere Pods**: View information about vSphere Pods.

# Microsoft Azure Dashboards

Use dashboards to monitor and troubleshoot Microsoft Azure issues in VMware Aria Operations.

To access the dashboards, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, navigate to **All > Microsoft Azure**.

The following dashboards are available:

| Dashboard Name | Purpose |
| --- | --- |
| Availability | View the availability of each Microsoft Azure service. Available services are green. Unavailable services are red and will be removed. |
| Inventory | View the adapter instance count in each resource group. Select a resource group to see a sparkline chart and the metrics for all the resources in the group.<br>Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the inventory for the database.<br>**Note** Metrics that are not collected or created are grayed out. |
| Optimization | View whether you are effectively using Microsoft Azure services. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart. |
| Virtual Machine | Select a virtual machine to view its scoreboard, property list, object relationship with resource group, and CPU usage and forecasting. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart. |
| SQL Database | Select an SQL server in the SQL Server widget and then select an SQL database corresponding to the server in the SQL Database widget to view the scoreboard, object relationship, and CPU usage for the database. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart. |
| Load Balancer | Select a load balancer to view its scoreboard, object relationship, and data path availability. This dashboard collects the CPU usage data in the form of metrics for the last 24 hours and displays forecasting information for the next 24 hours in a rolling view chart. |

## Azure Inventory Dashboard

Using the **Azure Inventory** dashboard you can navigate through the Azure inventory and view service details for each Azure subscription.

The **Azure Inventory** dashboard is designed to provide an overview of your Azure inventory. By default, the total number of subscriptions and services displayed are at the Azure World level. You can view the number and types of services and also the growth or decline of a service within a specific period for individual subscriptions. You can also drill down and view the service summary per region for a specific subscription. Using this information, you can understand growth trends, and manage and utilize your inventory cost-effectively.

The **Azure Inventory** dashboard helps you answer several key questions:

- How many subscriptions are being monitored?
- What services are being used?
- What is the growth trend for services?

### How to Use the Dashboard

- **Summary** widget: Displays the total number of Azure subscriptions and resource groups being monitored at the Azure World level.
- **Subscriptions** widget: Select a subscription from the table to refresh the service count and view the services specific to that subscription.
- **Regions** widget: Select a region to filter the service count and view service information per region.
- **Resource Groups** widget: Select a resource group to filter and view service information per resource group.

  **Note** You can filter services either by Regions or by Resource Groups at a time.

- **Inventory by Services** widget: Double click a service to view the growth trend and details regarding the service.

## Azure Virtual Machine Dashboard

Using the **Azure Virtual Machine** dashboard you can manage the network and CPU utilization of your Virtual Machines.

You can use the **Azure Virtual Machine** dashboard to monitor the health of all your Virtual Machines in the Azure inventory and filter them based on subscriptions, regions, and resource groups. The dashboard displays the active and the total number of Virtual Machines. You can also view the size of the instances, their CPU utilization, and network bandwidth at the Azure World level.

The **Azure Virtual Machine** dashboard displays the top 10 Virtual Machines with the highest and the lowest CPU, and network bandwidth at the Azure World level. You can filter this information and view it based on subscriptions, resource groups, and regions per subscription.

The **Azure Virtual Machine** dashboard helps you answer several key questions:

- How many Virtual Machines are there?
- How many Virtual Machines are running or shut down?
- What alerts are being generated for the Virtual Machines?
- What are the key metrics for a specific Virtual Machine?

### How to Use the Dashboard?

The **Azure Virtual Machine** dashboard displays the subscription names, available regions, Virtual Machines by size, and the total number of Virtual Machines at the Azure World level. It also displays the top ten Virtual Machines with the highest and lowest CPU utilization and CPU distribution, the top ten Virtual Machines with the highest and lowest network bandwidth and network distribution. You can also view the Virtual Machines according to their power state and duration for which were running and view the alerts for the Virtual Machines. You can also view the list of Virtual Machines with key metrics at the Azure World level.

- **Subscriptions** widget: Displays the total number of Azure subscriptions being monitored at the Azure World level. Select a subscription from the table to refresh the dashboard and view the Virtual Machine information for the specific subscription.
- **Regions** widget: Select a region to filter the displayed information and view the data for Virtual Machines per region.
- **Resource Groups** widget: Select a resource group to filter the displayed information and view the data for Virtual Machines per resource group.

  **Note** You can filter Virtual Machines either by Regions or by Resource Groups at a time.

- **Virtual Machines List** widget: Displays the list of Virtual Machines with key metrics at the Azure World level. The list of Virtual Machines gets updated according to the subscription and region/resource group you select. Select a Virtual Machine to view its key details, CPU utilization, and network bandwidth for a particular Virtual Machine.
- **Alert List** widget: Displays alerts based on the project and region you select. Click the alert name to view the details of the alert.

## Azure Kubernetes Service Overview Dashboard

The **Azure Kubernetes Service Overview** dashboard is built to provide an overview of your Azure Kubernetes Service environment. You can view the number of clusters configured, along with the configured pods, and pods which are in a Ready state. By default, the numbers are displayed at the Azure World level and you can filter by subscription, region, and resource group.

### How To Use The Dashboard?

- **Subscriptions** widget: Displays all the subscriptions available. Select a subscription from the list.
- **Regions** widget: Displays the regions associated with the subscription you selected.
- **Resource Groups** widget: Displays the resource groups associated with the subscription you selected.
- **Clusters** widget: Displays the total number of Azure Kubernetes Service clusters configured for the selected subscription, region, and resource group.
- **Clusters Configured** widget: Displays a list of configured clusters and their details such as, the name, the FQDN, and the provisioning state.

- **Cluster Details** widget: Displays the details of the selected configured cluster such as the CPU cores available, the memory available, whether RBAC enabled, and the version.
- **Number of Pods** widget: Displays the number of pods for the selected configured cluster.
- **Number of Pods in Ready State** widget: Displays the number of pods for the selected configured cluster that are in the Ready state.

## Azure Virtual Network Overview Dashboard

The **Azure Virtual Network Overview** dashboard is built to provide an overview of virtual network resources configured in your Azure environment. You can view the number of resources configured, and drill down into an individual virtual network to check for round trip times and failed pings. These numbers are displayed at the Azure World level by default, and you can filter by subscription, region and resource group. .

### How To Use The Dashboard?

- **Subscriptions** widget: Displays all the subscriptions available. Select a subscription from the list.
- **Regions** widget: Displays the regions associated with the subscription you selected.
- **Resource Groups** widget: Displays the resource groups associated with the subscription you selected.
- **Virtual Networks** widget: Displays the total number of virtual networks configured for the selected subscription, region, and resource group.
- **Virtual Networks Configured** widget: Displays a list of configured virtual networks and their details such as, the name and whether VM protection and DDoS protection is enabled.
- **Round Trip Time for Pings** widget: Displays the round trip time for pings for the selected virtual network. Round trip time is a key factor for measuring network latency and page load times.
- **Failed Pings** widget: Displays the total number of failed pings for the selected virtual network. When you ping a VM when it is powered off, it results in a failed ping.
- **Virtual Network Topology** widget: Provides a topological view of the configured virtual network.

## Azure Storage Account Overview Dashboard

The **Azure Storage Account Overview** dashboard is built to provide an overview of storage accounts configured in your Azure environment. You can view the number of storage accounts configured, along with utilization statistics and availability. These numbers are displayed at the Azure World level by default and you can filter by subscription, region, and resource group. You can drill down into individual storage accounts and analyze the availability, performance, and usage of the associated Blob, file, table, and queue storage.

### How To Use The Dashboard?

- **Subscriptions** widget: Displays all the subscriptions available. Select a subscription from the list.
- **Regions** widget: Displays the regions associated with the subscription you selected.
- **Resource Groups** widget: Displays the resource groups associated with the subscription you selected.
- **Storage Accounts** widget: Displays the total number of storage accounts configured for the selected subscription, region, and resource group.
- **Top 10** widgets: These widgets display the top ten least available storage accounts, the top ten high capacity storage accounts, the top ten low capacity storage accounts, and the top ten most used storage accounts.
- **Storage Accounts Configured** widget: Displays a list of configured storage accounts and their details such as, the name, availability, and total transactions.
- **Storage Account Details** widget: Displays details of the selected configured storage account such as, used capacity, ingress, egress, transactions, success server latency, and success E2E latency.

  **Capacity Section**

- **Blob, File, Queue, and Table** widgets: The widgets in this section display capacity details for Blob, file, queue, and table. The widgets are updated according to the configured storage account you select.

  **Transaction Section**

- **Availability** widget: Displays the availability of Blobs, files, queues, and tables. The widget is updated according to the configured storage account you select.
- **Transactions** widgets: The widgets display Blob, file, queue, and table transactions. The widget is updated according to the configured storage account you select.
- **Performance** widget: Displays performance data such as E2E latency and success server latency for Blob, file, queue, and table. The widget is updated according to the configured storage account you select.
- **Ingress and Egress** widget: Displays Blob, file, queue, and table ingress and egress information. The widget is updated according to the configured storage account you select.

## AWS Dashboards

Dashboards provide the user interface you use to monitor and troubleshoot Amazon Web Services problems in VMware Aria Operations.

You can access the dashboards by selecting **Dashboards**, and then selecting **AWS**.

## Table 15-3. AWS Dashboards

| Dashboard Name | Purpose |
| --- | --- |
| AWS Alerts | The Alerts dashboard reports system-generated performance information for Amazon Web Services. In VMware Aria Operations 5.8 and later, the dashboard also displays alerts received from Amazon Web Services Cloudwatch. |
| AWS ASG Utilization | Use the Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric.<br>ASG metrics are not collected by default. You must activate them when creating the group. This applies only to the metrics belonging directly to the auto scale group, for example GroupDesiredCapacity. It does not apply to the aggregate instance metrics for the ASG, for example Instance Aggregate CPU Utilization. |
| AWS Disk Space | Use the Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default.<br>For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html |
| AWS Instance Heatmap | Use the Instance Heatmap to monitor CPU/Disk/Network metric elements and identify instances that perform poorly. |
| AWS Instance Utilization | Use to identify which EC2 instances have high use across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by making adjustments to EC2 instances. |

### Table 15-3. AWS Dashboards (continued)

| Dashboard Name | Purpose |
| --- | --- |
| AWS Troubleshooting | This dashboard is most helpful when someone calls in with a problem and you know which device they are using. You can search for that type of device or the specific device, if you know the name.<br>When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas. You can use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another. |
| AWS Volume Performance | Use the Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations. |
| AWS Availability | Use this dashboard to view the availability of each AWS service. |
| AWS Inventory | Use this dashboard to view the count of each AWS service instance in each region. |
| AWS Optimization | Use this dashboard to view if you are effectively using AWS services. |

### Table 15-4. AWS - All Other Dashboards

| Dashboard Name | Purpose |
| --- | --- |
| AWS Services<br>■ CloudFormation Stacks<br>■ Compute: EC2<br>■ Compute: Elastic Containers<br>■ Compute: Lambda Functions<br>■ Database: Dynamo<br>■ Database: Elasticache<br>■ Database: RDS<br>■ Database: Redshift<br>■ Desktop: Workspaces<br>■ Network: Load Balancers<br>■ Network: VPS<br>■ Simple Queue Services<br>■ Storage | Select **AWS Services** and then select a dashboard to view a specific service-related information. |

## AWS Instance Utilization Dashboard

Use the AWS Instance Utilization dashboard to identify which EC2 instances have a high usage across the metrics for CPU, Disk IO, Network Transmissions, Received/Sent, and Memory. Use that information to determine whether you can optimize the system by adjusting the EC2 instances.

For example, you might determine that you need to resize the EC2 instance to make it larger or smaller.

You most often use this dashboard to troubleshoot issues with the listed metrics based on a support request from a user.

You can also identify which EC2 instances have been running for the longest and shortest amount of time. Then, you can use that information to determine whether EC2 instances can be decommissioned, or discover instances that have been added and need to be tracked in inventory.

Memory metrics require that you implement an add-on for each EC2 instance. These add-ons cost extra, and are not included by default.

## AWS Auto Scaling Group Dashboard

Use the AWS Auto Scaling Group (ASG) dashboard to identify which ASG groups have a high utilization across the metrics CPU, Disk IO, Network Transmissions, Received/Sent, and Number of Instances in the ASG. Use that information to determine whether any action is needed to adjust the ASG parameters. For example, you might need to raise or lower the scaling threshold for the CPU metric.

## AWS Troubleshooting Dashboard

When a user calls in with a problem and you know the name of the device they are using, can search for that type of device or the specific device and use the AWS Troubleshooting dashboard to get an overview of the system functionality.

When you select the device, the relationship tree displays the item, its parents, and children. You can observe the Health, Workload, Anomalies, and Faults to get an overview of how the system is functioning in those areas.

Use information in the Interesting Metrics widget to help identify the root cause of issues. The Health, Anomalies, and Events Mash-up widget allows you to compare changes in the system to see how they might affect one another.

There is a suggested flow to using the widgets in this dashboard.

1 Start with only the AWS Object widget open, and find the item you want to inspect.

2 Select the item, then expand the AWS Relationship widget to view the item's status.

3 Select one or all the related objects, then view the Ordered Symptoms, Interesting Metrics, and Mash-up.

4 Optionally, drag widgets into a new configuration if it makes it easier for you to compare information that is meaningful to you.

5 Examine the list of ordered symptoms and determine which of these events, in the given order might cause the problem to occur.

## AWS Instance Heatmap Dashboard

Use the AWS Instance Heatmap dashboard to monitor CPU/Disk/Network metric elements and identify instances that perform poorly.

You can use the Troubleshooting dashboard to find more detail, and research the root cause of issues. Then you can view the specific object instance to identify faulty processes and take a corrective action.

## AWS Volume Performance Dashboard

Use the AWS Volume Performance dashboard to identify Elastic Block Store (EBS) volumes that are experiencing high disk read time, high disk write time, a high volume of disk read operations, or a high volume of disk write operations. When you identify the EC2 instance that generates the load, use the Troubleshooting dashboard to investigate further.

## AWS Disk Space Dashboard

Use the AWS Disk Space dashboard to monitor EBS volumes to see whether they are running out of disk space and take appropriate action to anticipate future storage needs. Amazon Web Services does not report disk space by default.

For more information on accessing additional metrics, including disk space, and corresponding pricing, go to the Amazon Web Services documentation page at http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/mon-scripts.html.

## AWS Alerts Dashboard

The AWS Alerts dashboard reports system-generated performance information for Amazon Web Services. In VMware Aria Operations 6.6 and later, the dashboard also displays alerts received from Amazon Web Services Cloud watch.

## AWS Inventory Dashboard

Using the **AWS Inventory** dashboard you can navigate through the AWS inventory and view service details for each AWS account.

The **AWS Inventory** dashboard is built to provide an overview of your AWS inventory. By default, the total number of accounts and services displayed is at the AWS World level. You can view the number and types of services and also the growth or decline of a service within a specific period for individual accounts. You can also drill down and view the service summary per region for a specific account. Using this information you can understand growth trends, and manage and utilize your inventory cost-effectively.

The Inventory dashboard helps you answer several key questions:

- How many accounts are being monitored?
- How many services are running?
- What services are being used?
- What is the growth trend of the services?

### How to Use the Dashboard

The **AWS Inventory** dashboard displays the total number of accounts, the account names, available regions, and the total number of services running at the AWS World level.

- **Total Number of Accounts Being Monitored** widget: Displays the total number of AWS accounts being monitored at the AWS World level.
- **Accounts** widget: Select an account from the table to refresh the service count and view the services specific to that account.
- **Regions** widget: Select a region to filter the service count and view service information per region.
- **Services** widget: Services are arranged in the alphabetical order, double click a service to view the growth trend and details regarding the service.

## Compute EC2 Dashboard

Using the **Compute EC2** dashboard you can cater to the network, memory, and CPU utilization of your EC2 instances at the AWS World level.

You can use the **Compute EC2** dashboard to monitor the health of all your EC2 instances in the AWS inventory and filter them based on accounts and regions. The dashboard displays the total number of EC2 instances and the number of active instances. You can also view the size of the instances and their CPU, network, and memory utilization and distribution at the AWS World level.

The **Compute EC2** dashboard displays the top 10 EC2 instances with the highest and the lowest CPU, network, and memory utilization and distribution at the AWS World level. You can filter this information and view it based on accounts and regions per account.

The **Compute EC2** dashboard helps you answer several key questions:

- How many EC2 instances are there?
- How many EC2 instances are running or shut down?
- What alerts are generated for the EC2 instances?
- What are the key metrics for specific VMs in an EC2 instance?

### How to Use the Dashboard?

The **Compute EC2** dashboard displays the account names, available regions, EC2 instances by size, and the total number of EC2 instances at the AWS World level. It also displays the top ten EC2 instances with the highest and lowest CPU utilization and CPU distribution, the top ten EC2 instances with the highest and lowest network bandwidth and network distribution, and the top ten EC2 instances with the highest and lowest memory usage and distribution. To view memory-related data for any EC2 instances, the AWS CloudWatch agent must be installed on the EC2 instances.

You can also view the EC2 instances according to their power state and running time and view the alerts for the EC2 instances. You can also view the list of EC2 instances with their key metrics at the AWS world level.

- **Accounts** widget: Displays the total number of AWS accounts being monitored at the AWS World level. Select an account from the table to refresh the dashboard and view the EC2 instance information for the specific account.
- **Regions** widget: Select a region to filter the displayed information and view the data for EC2 instances per region.
- **EC2 Instances with Key Metrics** widget: Displays the list of EC2 instances with key metrics at the AWS World level. The list of EC2 instances gets updated according to the account and region you select. Select an EC2 instance to view its key metrics, CPU utilization, network bandwidth, and memory usage for a particular EC2 instance.

## AWS Network Services Overview Dashboard

The **AWS Network Services Overview** dashboard is built to provide an overview of your AWS VPCs and key network services linked to it.

You can view the number of VPCs, Subnets, NAT Gateways, and VPN connections.

The **AWS Network Services Overview** dashboard helps you answer several key questions.

- What is the network service count for a specific account or region?
- How many VPCs are used for a specific account or region?
- What are the NAT gateways for a specific VPC?
- How many VPNs and Subnets are connected to a specific VPC?

### How to Use the Dashboard?

In the **AWS Network Services Overview** dashboard, the numbers of VPCs, Subnets, NAT Gateways, and VPN connections are displayed at the AWS World level by default. You can filter them by account and by region. You can also drill down to a specific VPC and check the NAT Gateways and VPN Connections linked to the selected VPC.

- **Accounts** widget: Lists the AWS accounts. Select an account from the table to refresh the dashboard and view the VPC information for the specific account.

- **Regions** widget: Lists the AWS regions. Select a region to filter the displayed information and view the data for VPCs per region.
- **Network Service Count** widget: Displays the total number of VPCs, number of VPC NAT Gateways, and the number of VPN connections.
- **VPC List**: Displays the names and regions of the VPCs deployed. The first VPC in the list is selected by default. The list of VPCs gets updated according to the account and region you select. Select a VPC instance to view the associated VPC NAT gateways, VPN connection, and connected services.
- **VPC NAT Gateways**: Displays the VPC NAT Gateway information for the first VPC in the VPC list by default. The VPC NAT gateways list gets updated according to the VPC you select.
- **VPN Connection**: Displays the number of VPN connections for the first VPC in the VPC list by default. The VPN connections information gets updated according to the VPC you select.
- **Connected Services**: Displays the number of connected services for the first VPC in the VPC list by default. It displays the number of VPNs and the number of subnets connected to the VPC. The information gets updated according to the VPC you select.

## AWS Storage Services Overview Dashboard

Using the **AWS Storage Services Overview** dashboard you view the total storage service count (EBS volumes and S3 buckets).

The **AWS Storage Services Overview** dashboard provides an overview of your EBS volumes and S3 buckets. You can view the number of EBS and S3 buckets used in your environment.

The **AWS Storage Services Overview** dashboard helps you answer several key questions.

- What is the usage of EBS volumes?
- Which are the top read and top written EBS volumes?
- Which S3 buckets have maximum usage?

### How to Use the Dashboard?

The **AWS Storage Services Overview** dashboard is divided into two sections based on the two storage types, **EBS Volumes** and **S3 Buckets**. By default, the numbers are displayed at the AWS World level and you can filter the information by account and by region. You can check the usage of the EBS volumes and drill down to individual volumes. For S3 buckets, you can view the buckets with maximum usage and errors. You can also select an individual bucket to check the traffic and errors associated with it.

- **Accounts**: Lists the AWS accounts. Select an account from the table to refresh the storage information and view the EBS volumes and S3 buckets specific to that account.
- **Regions**: Lists the AWS regions. Select a region to filter the storage information and view the EBS volumes and S3 buckets specific to that region.

- **EBS Volumes**: The EBS volumes section displays the top read EBS volumes, top written EBS volumes, EBs volumes with maximum idle time, and the EBS volumes list based on the load that they have. The lists get updated according to the account or region you select.

  The **EBS Volumes List** displays the name, region, capacity, maximum queue, and maximum write time of the EBS Volumes storage. The first EBS volume in the list is selected by default. The list of EBS volumes gets updated according to the account and region you select. Select an EBS volume to view the associated read time and write time.

- **S3 Buckets**: The S3 buckets section displays the S3 buckets with maximum total latency, S3 buckets with maximum first byte latency, S3 buckets with most 4xx errors, S3 buckets with most 5xx errors, and the S3 bucket lists. The lists get updated according to the account or region you select.

  The **S3 Buckets List** displays the name, region, object, and size of the S3 buckets storage. The first S3 bucket in the list is selected by default. The list of S3 buckets gets updated according to the account and region you select. Select a S3 bucket to view the associated requests and the 4xx and 5xx errors related to it.

# Dashboards in VMware Cloud on AWS

The **VMware Cloud on AWS** dashboards allow you to track the capacity, cost, and inventory overviews of the SDDCs. You can also track the virtual machines monitoring and the utilization and performance of these SDDCs.

## VMC Capacity Dashboard

Use the **VMC Capacity** dashboard to view the capacity overview of each VMware Cloud on AWS SDDC. You can view the capacity of Clusters, Hosts, VMs, Datastores, and Disk groups.

Table 15-5. Widgets in VMC Capacity Dashboard

| Widget | Description |
| --- | --- |
| VMC SDDC by Capacity Remaining % | Displays the SDDCs as cards that show the remaining capacity percentage. |
| VMC SDDC by Time Remaining % | Displays the SDDCs as cards that show the remaining time percentage. |
| VMC SDDC by Virtual Machine Remaining (based on avg VM profile) | Displays the SDDCs as cards that show the remaining number of virtual machines. |

When you select one of the SDDC cards, the details of that SDDC are automatically populated in the widgets after the VMC SDDC by Virtual Machine Remaining (based on avg VM profile) widget.

**Note** The key kpis are color-coded to help in identifying capacity bottlenecks.

## VMC Cost Overview Dashboard

Use the **VMC Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

### Table 15-6. Widgets in VMware Cloud on AWS Dashboard

| Widget | Description |
|---|---|
| Organization Cost Overview | Displays a list of organizations with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD). |
| Outstanding Expense Trend | Displays the outstanding expense trend of the organization selected in the Organization Cost Overview widget. |
| Total Expense Trend (Monthly) | Displays the total monthly expenses trend of the organization selected in the Organization Cost Overview widget. |
| Commit Expense Trend (Monthly) | Displays the committed monthly expense trend of the organization selected in the Organization Cost Overview widget. |
| On-Demand Expense Trend (Monthly) | Displays the on-demand monthly expenses trend of the organization selected in the Organization Cost Overview widget. |
| Purchase History | Displays the bill line items/purchases from the available bills. |
| Currency Information | Represents the metrics currency unit set in this management pack account. |

**Note** The YTD metric is an aggregation from the beginning of the calender year, until the last available bills.

## VMC Inventory Dashboard

Use the **VMC Inventory** dashboard to view the inventory overview of all the SDDCs configured in VMware Cloud on AWS.

### Widgets in VMC Inventory Dashboard

VMC SDDCs: displays the SDDCs as cards that show the number of virtual machines running in the SDDC. The SDDC card also shows a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that SDDC, the SDDC card indicates this by changing colors.

When you select one of the SDDC cards, the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs with key configuration details of that SDDC are populated in the widgets after the VMC SDDCs widget.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

# VMC Management VM Monitoring Dashboard

Use the **VMC Management VM Monitoring** dashboard to monitor the utilization and performance of the key management VMs running in your SDDC. This dashboard ensures that the management components (such as vCenter and NSX) are not facing any resource bottlenecks from the CPU, memory, network, and storage perspectives.

Table 15-7. Widgets in VMC Management VM Monitoring Dashboard

| Widget | Description |
| --- | --- |
| CPU Usage & Performance | Displays the list of all the management components in each SDDC with key CPU utilization and performance KPIs. Select a management VM to see the usage and performance trends of all the CPU cores. |
| Memory Usage & Performance | Displays the list of all the management components in each SDDC with key Memory utilization and performance KPIs. Select a management VM to see the memory usage and performance trends. |
| Network Usage & Performance | Displays the list of all the management components in each SDDC with key Network utilization and performance KPIs. Select a management VM to see the memory usage and performance trends. |
| | |
| Storage Usage & Performance | Displays the list of all the management components in each SDDC with key storage utilization and performance KPIs. Select a management VM to see the network usage and performance trends. |

# VMC Utilization and Performance Dashboard

Use the **VMC Utilization and Performance** dashboard to view the utilization and performance overview of each SDDC based on heavy hitter VMs and impacted VMs over the last 30 days. This dashboard helps you in finding the VMs in your environment that are negatively impacting the capacity or performance from a CPU, memory, storage, or network perspective.

## Widgets in VMC Utilization and Performance Dashboard

List of VMC SDDCs: displays the list of all the SDDCs with aggregate CPU, memory, and storage utilization with 95th percentile and maximum values over the last 30 days.

When you select one of the SDDC from the List of VMC SDDCs widget, you can see the list of top VMs that are consuming compute, network & storage resources in that SDDC. The widgets after that show the compute (CPU & memory) utilization and performance analysis, network, storage, and utilization and performance analysis.

Each section in the dashboard is based on the last 30 days data with 95th percentile transformation which is configurable to Max, Average, Current, Standard Deviation, or other mathematical transformations.

# VMC Configuration Maximums Dashboard

Use the **VMC Configuration Maximums** dashboard to view the VMC limits and your consumption against those limits. This dashboard displays alerts for configuration maximum, and details of organization, SDDC, vSAN, and cluster maximums.

### Table 15-8. Widgets in VMC Configuration Maximums Dashboard

| Widget | Description |
|---|---|
| Select an Environment | Select an environment for which you want to view the alerts and other details. Once you select an environment, the details of that environment are automatically populated in the widgets below. |
| VMC Configuration Maximums Alerts | Displays the list of alerts for the selected environment. |
| Number of SDDCs | Displays the number of SDDCs for the organization maximums, the provisioned, and the soft limit used. |
| Number of Hosts | Displays the number of hosts for the organization maximums, the provisioned, and the soft limit used. |
| Public IP Addresses (Elastic IPs) | Displays the public IP addresses for the organization maximums, the provisioned, and the soft limit used. |
| Maximum Clusters | Displays the maximum clusters for the SDDC maximums, the provisioned, and the hard and soft limit used. |
| Maximum Hosts | Displays the maximum hosts for the SDDC maximums, the provisioned, and the limit used. |
| Maximums VMs | Displays the maximum VMs for the SDDC maximums, the provisioned, and the limit used. |
| Linked VPCs | Displays the linked VPCs for the SDDC maximums, the provisioned, and the limit used. |
| Clusters with No SLA | Displays the maximum number of clusters and the number of provisioned clusters with no SLA per SDDC. An empty list means no clusters have been identified with no SLA. |
| Clusters with Limited SLA | Displays the maximum number of clusters and the number of provisioned clusters with limited SLA per SDDC. An empty list means no clusters have been identified with limited SLA. |
| Max hosts per Cluster (including stretched clusters) | Displays the maximum hosts per cluster including the stretched clusters, the provisioned, and the limit used. |
| Datastore Utilization | Displays the datastore utilization for vSAN maximums, the used space, utilization limit, and the remediation needed. |
| VMs per Host Limit Used | Displays the maximum number VMs that can be deployed per host, VMs that are provisioned per host, and the percentage of limit used. |
| VMs per Host Limit Used of Selected Host | Displays the VMs that are used per host limit for a selected host. |

# Google Cloud VMware Engine Dashboards

The Google Cloud VMware Engine dashboards allow you to track the configuration maximums, cost, and inventory overviews of the Google Cloud VMware Engine private clouds.

## GCVE Configuration Maximums Dashboard

Use the **GCVE Configuration Maximums** dashboard to view the Google Cloud VMware Engine limits and your consumption against those limits. This dashboard displays alerts for configuration maximum, and details of the adapter instance, private cloud, vSAN, and host and cluster maximums.

### Table 15-9. Widgets in the GCVE Configuration Maximums Dashboard

| Widget | Description |
|---|---|
| Select an Environment | Select an environment for which you want to view the alerts and other details. Once you select an environment, the details of that environment are automatically populated in the widgets below. |
| GCVE Configuration Maximums Alerts | Displays the list of alerts for the selected environment. |
| **Private Cloud Maximums** | |
| Maximum Clusters | Displays the maximum clusters for Google Cloud VMware Engine private cloud maximums, the provisioned, and the hard and limit used. |
| Maximum Hosts | Displays the maximum hosts for Google Cloud VMware Engine private cloud maximums, the provisioned, and the hard limit used. |
| **Cluster Maximums** | |
| Max hosts per Cluster (including stretched clusters) | Displays the maximum hosts per cluster including the stretched clusters, the provisioned, and the limit used. |
| **vSAN Maximums** | |
| Datastore Utilization | Displays the datastore utilization for vSAN maximums, the used space, utilization limit, and the remediation needed. |

## GCVE Cost Overview Dashboard

Use the **GCVE Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

### Table 15-10. Widgets in the GCVE Cost Overview Dashboard

| Widget | Description |
|---|---|
| Organization Cost Overview | Displays a list of bills with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD). |
| Outstanding Expense Trend | Displays the outstanding expense trend of the bill selected in the **Organization Cost Overview** widget. |

Table 15-10. Widgets in the GCVE Cost Overview Dashboard (continued)

| Widget | Description |
| --- | --- |
| Total Expense Trend (Monthly) | Displays the total monthly expenses trend of the bill selected in the **Organization Cost Overview** widget. |
| Commit Expense Trend (Monthly) | Displays the committed monthly expense trend of the bill selected in the **Organization Cost Overview** widget. |
| On-Demand Expense Trend (Monthly) | Displays the on-demand monthly expenses trend of the bill selected in the **Organization Cost Overview** widget. |
| Purchase History | Displays the bill line items/purchases from the available bills. |
| Currency Information | Represents the metrics currency unit set in this management pack account. |

**Note** The YTD metric is an aggregation from the beginning of the calender year, until the last available bills.

## GCVE Inventory Dashboard

Use the **GCVE Inventory** dashboard to view the inventory overview of all the private clouds configured in Google Cloud VMware Engine.

### Widgets in the GCVE Inventory Dashboard

**GCVE Private Clouds**: Displays a card per Google Cloud VMware Engine private cloud with the number of virtual machines running in each private cloud. The card also displays a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that private cloud, the card indicates this by changing colors.

When you select one of the private cloud cards, the widgets below are populated and you can see the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs in that private cloud with key configuration details.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

# Dashboards in NSX Management Pack

The **NSX Main** dashboard provides an overview of the network objects. It displays the topology of a selected object, how it connects to the elements in the network, and a view of related alerts.

### Table 15-11. Widgets in NSX Main Dashboard

| Widget | Description |
| --- | --- |
| NSX Instances | Displays the list of environments that are being monitored. When you select an environment in this widget, the other widgets in the **NSX Main** dashboard display data for the selected adapter. |
| Environment Overview | Displays a top-level view of the selected environment and following key components.<br>■ NSX Manager<br>■ Controller Node<br>■ Logical Router<br>■ Logical Switch<br>■ Load Balancer Virtual Server<br>■ Transport Zone |
| Top Alerts | Displays all the open alerts for the selected object in the **Environment Overview** widget. |
| Topology Graph | Displays the topology of the selected object in the **Environment Overview** widget. |

## NSX Configmax Metrics

The **NSX Configmax Metrics** dashboard in VMware Aria Operations provides an overview of all the configuration maximum metrics in all the NSX instances.

### Table 15-12. Widgets in NSX Configmax Metrics Dashboard

| Widget | Description |
| --- | --- |
| Select an adapter instance | Displays the list of all the NSX and NSX on VMC instances. When you select an instance in this widget, the other widgets in the **NSX Configmax Metrics** dashboard display data for the selected instance. |
| Relationship view | Displays the objects hierarchy for the instance selected in the **Select an adapter instance** widget. Only the objects with configuration maximum metrics are shown in the relationship view. |
| Select object from relationship view for the configmax metric | Displays all the configmax metrics for the selected object in the **Relationship View** widget. |
| Trend View | Displays all the MGW, CGW, and Distributed firewall section rule trends of the instance selected in the **Select an adapter instance** widget.<br>**Note** The Trend View widget loads the trends only for the firewall sections object on VMware Cloud on AWS instances. |

## NSX Configmax Metrics

The NSX Configmax Metrics provides information about the new metrics added to the NSX Configmax Metrics dashboard.

Table 15-13. NSX Configmax Metrics

| Metric Type | Metric | Description |
|---|---|---|
| Group | Configuration Maximums\|Count\|Tag Count | This metric displays the number of tags for the selected group. |
| Logical Router | Configuration Maximums\|ARP Entries Count | This metric displays the number of ARP entries of the logical router. |
| | Configuration Maximums\|Router Port Count | This metric displays the number of ports available in the logical router. |
| Management Cluster | Configuration Maximums\|Prepared vC Cluster Count | This metric displays the number of prepared vCenter clusters in the management cluster. |
| | Configuration Maximums\|Compute Manager Count | This metric displays the number of compute manager present in the management cluster. |
| Edge Cluster | Configuration Maximums\|Edge Node Count | This metric displays the number of edge nodes present in the edge cluster. |

## NSX Inventory Dashboard

The **NSX Inventory** dashboard in VMware Aria Operations provides a holistic view of NSX environment topology. It offers an graphical representation of the interconnections between various NSX components.

This dashboard helps administrators with essential insights into the following components:

- NSX Management Nodes: Displays information related to the status and configuration of NSX management nodes.
- T0 and T1 Routers: Displays details of the t0 and t1 routers including their configurations and associations.
- NSX Load Balancers (LBs) and LB Pools: Displays information about NSX Load Balancers, virtual servers, pool members, and health checks.
- Edge Clusters: Displays data on Edge Clusters including cluster membership and resource consumption.
- Transport Zones: Displays insights into transport zones, offering visibility into network segmentation and connectivity options.
- Edge Nodes and Host Nodes: Monitors Edge Nodes and Host Nodes that are crucial components of the NSX infrastructure.
- Logical Switches: Displays a comprehensive list of logical switches with associated properties and configurations.
- NSX Groups: Displays details related to NSX Security Groups and Service Groups, simplifying group membership and policy management.

## NSX System Dashboard

The **NSX System** dashboard in VMware Aria Operations provides essential insights into the core components of the NSX infrastructure specific to inventory and health.

This dashboard offers visibility into the following aspects:

- Management Clusters: Displays information about NSX Management Clusters including cluster health, member nodes, and their roles.
- Advanced Object Relationship: Allows you to explore intricate connections between different NSX objects, enhancing understanding of component interactions.
- Summary: Displays a concise snapshot of NSX system, presenting key performance metrics and alert notifications.
- Alerts: Monitors and manages alerts generated within the NSX environment, offering insights into critical issues that require attention and resolution.

## NSX Edge Dashboard

The **NSX Edge** dashboard in VMware Aria Operations provides detailed insights into the performance and configuration of Logical Routers, which is a foundational component of the NSX Edge environment.

This dashboard equips administrators with necessary tools to effectively manage, optimize, and troubleshoot the performance of Logical Routers within the NSX Edge environment. This dashboard offers the following features:

- Logical Routers: Displays information related to Logical Routers, including configurations, routing tables, and interface details.
- Advanced Object Relationship for Logical Router: Displays insights into complex relationships specific to logical routers.
- Logical Router RX and TX: Monitors traffic received (RX) and traffic transmitted (TX) by logical routers to optimize network performance and troubleshoot issues.
- VMs Using Logical Router: Identifies virtual machines (VMs) utilizing specific logical routers, offering visibility into network traffic patterns and dependencies.

## NSX Switch Dashboard

The **NSX Switch** dashboard in VMware Aria Operations provides information about logical switches. Its includes an inventory of logical switches with their associated properties and configurations.

Additionally, the dashboard offers the following functionalities:

- Inventory of Logical Switches: Displays detailed information on logical switches for efficient tracking and management of network segmentation.

- List of NSX Logical Switches: Displays a comprehensive list of NSX logical switches, simplifying the tracking of network segmentation.
- Advanced Object Relationship of Logical Switch: Provides insights into the intricate relationships and dependencies between logical switches and other NSX components.
- Logical Switch Inbound and Outbound Throughput: Monitors inbound and outbound traffic throughput for logical switches to optimize network performance and identify potential bottlenecks.

## NSX Transport Node Performance Dashboard

The **NSX Transport Node Performance** dashboard in VMware Aria Operations provides an extensive view of performance metrics related to NSX Transport Nodes. It provides essential insights into the operation and efficiency of these nodes, aiding network administrators in optimizing resource allocation and ensuring seamless data packet forwarding within the NSX environment.

You can monitor the following key metrics:

- DPDK CPU Core Average Usage
- Non-DPDK CPU Core Average Usage
- Memory used for Transport Node
- Transport Node Interface Statistics

## NSX Edge Performance Dashboard

The **NSX Edge Performance** dashboard in VMware Aria Operations provides performance information on NSX Edge devices. It offers visibility into the operation of these devices, assisting users in maintaining optimal routing and network services while swiftly addressing potential issues.

You can monitor the following key metrics:

- List of NSX Logical Routes
- RX MBPs (Megabits per second)
- TX MBPS (Megabits per second)
- RX Packets Per Second
- TX Packets Per Second
- RX Packets Dropped Per Second
- TX Packets Dropped Per Second

## NSX Switch Performance Dashboard

The **NSX Switch Performance** dashboard in VMware Aria Operations focuses on the performance of logical switches of the NSX environment. It simplifies the tracking and

management of network segmentation while offering insights into the dependencies between logical switches and other NSX components.

You can monitor the following key metrics:

- List of NSX Logical Switches
- Advanced Object Relationships
- Performance Statistics (Inbound and Outbound Throughput)

## NSX Load Balancer Performance Dashboard

The **NSX Load Balance Performance** dashboard in VMware Aria Operations offers detailed insights into the performance of NSX Load Balancer services. It includes critical data related to the operation of Load Balancers, providing administrators with the necessary information to optimize configurations and maintain high availability and performance for applications and services.

You can monitor the following key metrics:

- CPU Usage of the Load Balancer
- Memory Usage of the Load Balancer
- Session Statistics

# Cloud Federation Adapter Dashboards

Cloud Federation Adapter caters to certain out-of- the-box use cases, which are delivered through dashboards. After the configuration of the Cloud Federation Adapter management pack, you can access these out-of- the-box dashboards.

### Prerequisites

Ensure that you have enabled the following Management Packs for the Cloud Federation Adapter dashboards to display data.

- VMware Aria Automation
- CloudHealth
- VMware Cloud on AWS
- AWS
- Azure

### Procedure

1. To access the dashboards, in the left pane of the VMware Aria Operations, click **Dashboards**.
2. Click **Cloud Federation Analytics** to view all dashboards.

3 From the dashboard list, select **Getting Started - Cloud Federation** dashboard in the list of all the dashboards and associated use cases.

## Getting Started - Cloud Federation Dashboard

The Getting Started - Cloud Federation dashboard displays the catalog of Cloud Federation Analytics dashboards and the Cloud Federation Analytics deployment map.

You can step through the multiple categories under Cloud Federation Analytics page to cater to specific use cases. The deployment map on this page provides a topology of your Cloud Federation Analytics Deployment with colors depicting the current health and collection status.

### Infrastructure Operations

The Getting Started page provides access to the following OOTB categories. The infrastructure category caters to Senior Executives and NOC users by providing them an operations summary and alerts. The dashboards under this category are.

- Operations Summary
- Alerts Overview

### Operations Summary Dashboard

The Operations Summary dashboard provides a count of your infrastructure inventory objects across all the VMware Aria Operations instances being monitored by the Cloud Federation Adapter. It provides you the overall alert volume in these environment, along with a summarized view of health and capacity of the vSphere clusters in your environment.

### Alerts Overview

The Alerts Overview dashboard provides the alerts page of VMware Aria Operations. On this page, you can not only view the alerts triggered on your VMware Aria Operations Cloud Federation Analytics instance, but also view the alerts triggered in your child VMware Aria Operations instances.

The overview page includes all the alerts triggered on the VMware Aria Operations instance and is not dependent on the objects you are monitoring from the child VMware Aria Operations instance.

This page can help your NOC teams to have a single pane for all the alerts across your environments. To view the details of a specific alert, click on that alert to get directed to the child VMware Aria Operations instance where the alert was originally triggered for troubleshooting purposes.

## Infrastructure Capacity

The Infrastructure Capacity category caters to Senior Executives and NOC users by providing them an overview of overall capacity and reclamation opportunities.

The following dashboards are under the Infrastructure Capacity:

- Capacity Summary
- Capacity Optimization

## Capacity Summary Dashboard

The Capacity Summary Dashboard provides you with a summary of the total physical capacity available across all your environments being monitored by multiple instances of VMware Aria Operations. The dashboard also provides you a list view of all your clusters across your environment with details around inventory, capacity and utilization of those clusters. You can export this list into a CSV file for reporting purposes.

Figure 15-1. Capacity Summary Dashboard

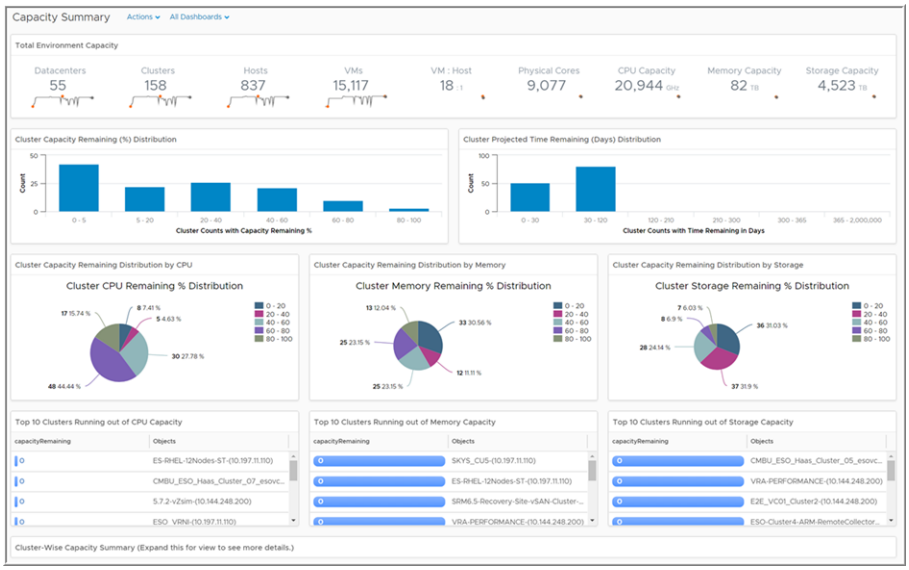


## Capacity Optimization Dashboard

The Capacity Optimization dashboard provides you a quick view of resource optimization opportunities within your virtual infrastructure. This dashboard is focused on improving the efficiency of your multiple instances of VMware Aria Operations environment by reducing the wastage of resources.

Figure 15-2. Capacity Optimization Dashboard

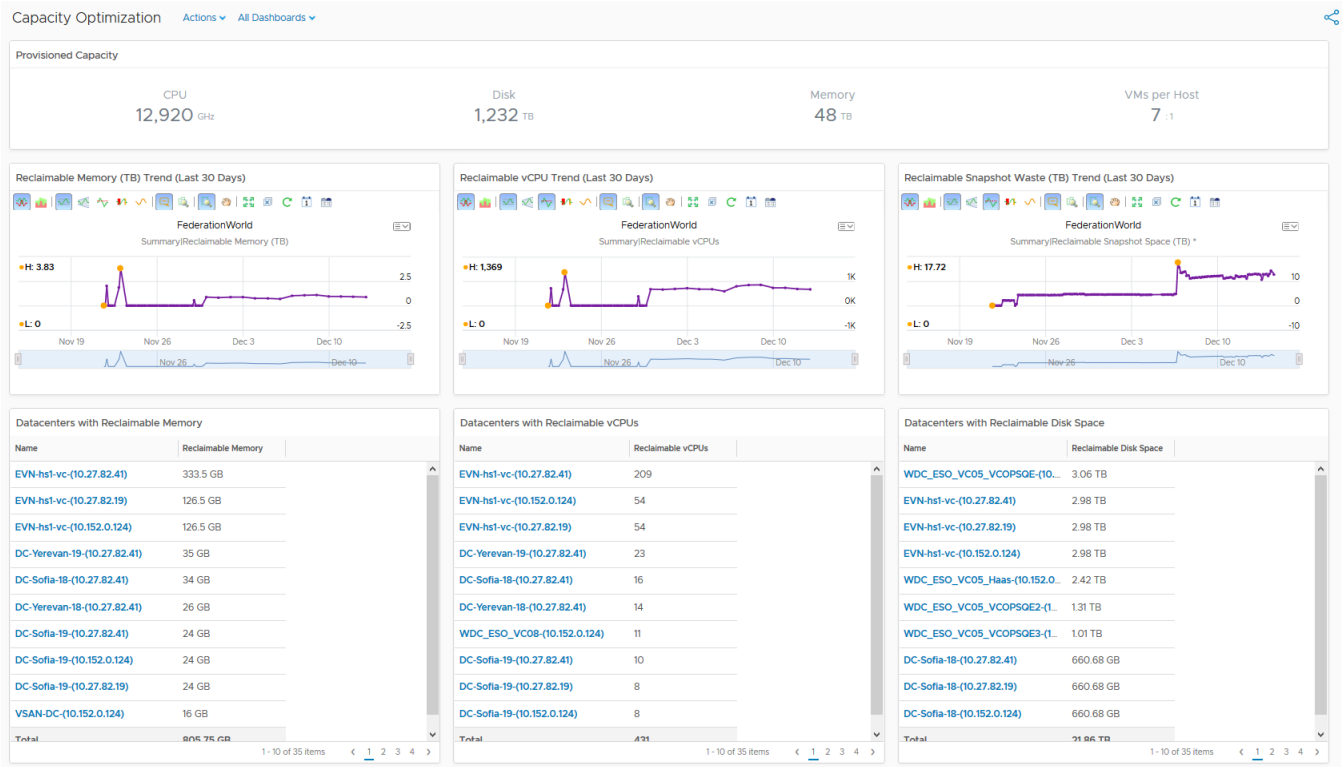


# Infrastructure Configuration

Infrastructure Configuration category caters to Senior Executives and virtual infrastructure admin by providing them the summary of ESXi Host and cluster configurations.

The dashboards under Infrastructure Configuration category are:

- vSphere Cluster Configuration
- vSphere Host Configuration

## vSphere Cluster Configuration Dashboard

The vSphere Cluster Configuration dashboard provides a quick overview of your vSphere cluster configurations from various instances of VMware Aria Operations. The dashboard highlights the areas that are important in delivering performance and availability to your virtual machines.

Figure 15-3. vSphere Cluster Configuration Dashboard

The dashboard also highlights if there are vSphere clusters which are not configured for DRS, High Availability (HA), or admission control to avoid any resource bottlenecks or availability issues when a host fails. The heat map in this dashboard helps you to identify if you have hosts where vMotion was not enabled as this may not allow the VMs to move from or to that host. This can cause potential performance issues for the VMs on that host if the host gets too busy.

## vSphere Host Configuration Dashboard

The vSphere Host Configuration dashboard provides an overview of your ESXi host configurations, and displays inconsistencies so that you can take corrective action.

Figure 15-4. vSphere Host Configuration Dashboard

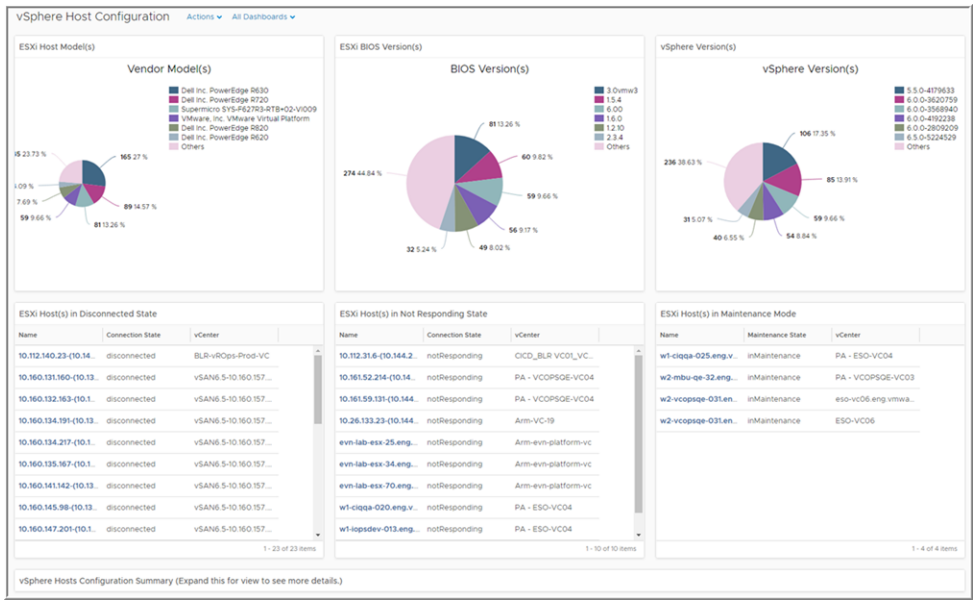

The dashboard also measures the ESXi hosts against the vSphere best practices and indicates deviations that can impact the performance or availability of your virtual infrastructure. Although you can view this type of data in other dashboards, in this dashboard you can export the ESXi configuration view and share it with other administrators.

## SDDC Health and Configuration

This category caters to Senior Executives and virtual infrastructure administrator by providing a summary of health and configuration of the SDDC Management Stack.

Health & Configuration dashboard is available under the SDDC Health and Configuration category.

### Health and Configuration Dashboard

This dashboard provides health of all the SDDC components. The Cloud Federation Analytics adapter monitors and displays the health of vCenter Server, VMware Aria Operations health, VMware Aria Automation and other stack components.

**Note** Verify to install and configure the SDDC health adapter installed on the child VMware Aria Operations instances to display the health and configuration information in this dashboard.

Figure 15-5. Health and Configuration Dashboard

This dashboard provides information of summary of the number of SDDC components, applications, deployments and its health status. The **SDDC instances and Issue count** widget in the dashboard displays detailed information of which object requires an immediate attention and from which instance the issue has occurred. When you click on the particular VMware Aria Operations child instance, you can drill down to analyze the root cause of an issue.

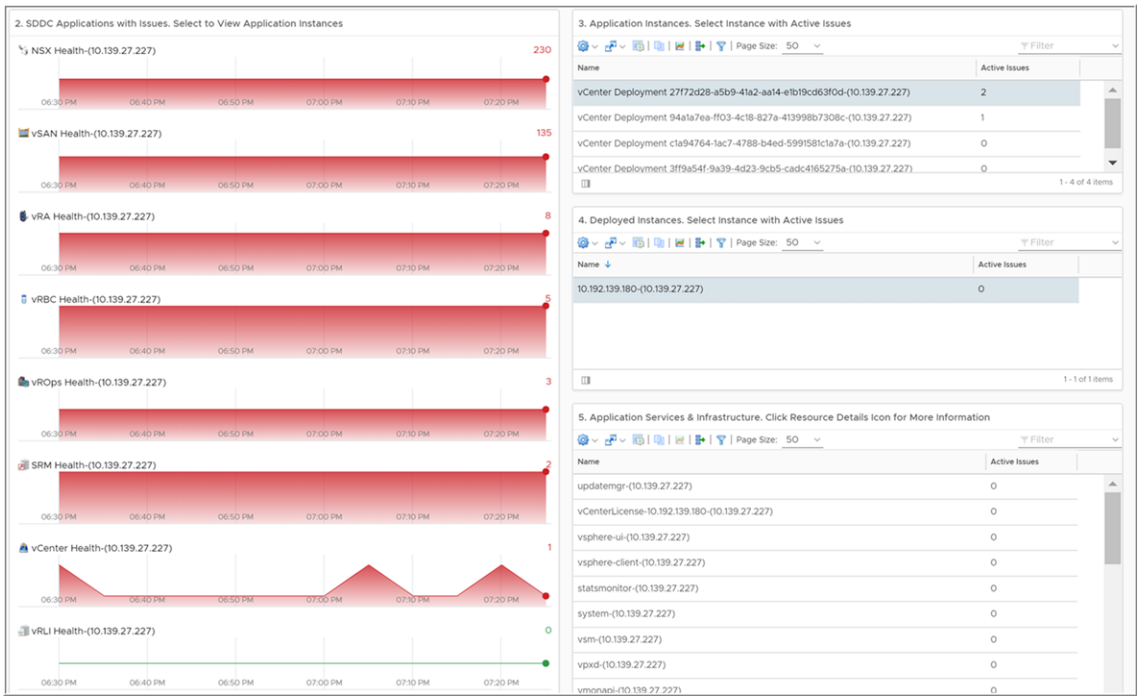


# Dashboards in VMware Infrastructure Health

Use the VMware Infrastructure Health dashboards to view data related to its certificates, historic trend, health, VM, and versions of all products under the VCF management domain.

### Table 15-14. VMware Infrastructure Health Dashboards

| Dashboard | Description |
| --- | --- |
| VMware Infrastructure Health Certificate Overview Dashboard | The VMware Infrastructure Health dashboard provides certificate information related to vCenter Server, NSX Server, NSX Manager Node, ESXi, SDDC Product, and Adapter certificate alerts. |
| VMware Infrastructure Health Historic Trend | Use VMware Infrastructure Health health trend dashboard to view the health trend for each component in the cloud management applications.<br>The dashboard provides an overall trend of the health components of the cloud management applications, that are monitored for the last seven days. |

Table 15-14. VMware Infrastructure Health Dashboards (continued)

| Dashboard | Description |
| --- | --- |
| VMware Infrastructure Health Health Overview | You can use VMware Infrastructure Health dashboard to view and analyze the application-specific problems in the accounts in the VMware Cloud Management Plane. This Dashboard provides the health information for each of the accounts. You can select the component in the cloud management applications from the widgets available in the dashboard. The widgets help in rendering the infrastructure health of that component with its service health and associated configuration alerts, if any. |
| VMware Infra VCF BOM Version | The VMware Infra VCF BOM Version dashboard displays all the versions of all products under the VCF management domain. |

# Dashboards in Google Cloud Platform

With the dashboards in the Management Pack for Google Cloud Platform, you can monitor the following Google Cloud Platform services:

- Google Compute Engine
- Container/Kubernetes Engine
- Cloud Storage
- Cloud VPN
- Big Query

## Access the Dashboards

1 To access the dashboards, from the left menu, click **Visualize > Dashboards**.

2 From the dashboard list, select the required Google Cloud Platform dashboard.

## GCP Availability Dashboard

You can view the heat maps of the Google Cloud Platform services monitored by this Management Pack.

## GCP Big Data: Big Query Dashboard

You can view the information about the monitored Cloud Big Query type resources, with the curated tables and charts.

## GCP Compute: Container Dashboard

You can view the information about the monitored Cloud Kubernetes Engine type resources, with the curated tables and charts.

## GCP Compute: CE Dashboard

You can view the information about the monitored Compute Engine type resources, with the curated tables and charts.

## GCP Inventory Dashboard

You can view the information about the number of monitored resources of a Google Cloud Platform services type.

## GCP Network: VPN Dashboard

You can view the information about the monitored Cloud VPN type resources, with the curated tables and charts.

## GCP Storage: Buckets Dashboard

You can view the information about the monitored Cloud Storage type resources, with the curated tables and charts.

## GCP Inventory Dashboard

Using the **GCP Inventory** dashboard you can navigate through the GCP inventory and view service details for each GCP project.

The GCP Inventory dashboard is designed to provide an overview of your GCP inventory. By default, the total number of projects and services displayed are at the GCP World level. You can view the number and types of services and also the growth or decline of a service within a specific period for individual projects. You can also drill down and view the service summary per region for a specific project. Using this information, you can understand growth trends, and manage and utilize your inventory cost-effectively.

The Inventory dashboard helps you answer several key questions:

- How many projects are being monitored?
- What services are being used?
- What is the growth trend for specific services, projects, or regions?

### How to Use the Dashboard

- **GCP Inventory** widget: The inventory dashboard displays the total number of projects, the project names, available regions, and the total number of services running at the GCP World level.
- **Number of Projects** widget: Displays the total number of GCP projects being monitored at the GCP World level.
- **Projects** widget: Select a project from the table to refresh the service count and view the services specific to that project.

- **Regions** widget: Select a region to filter the service count and view service information per region.
- **Inventory by Services** widget: Double click a service to view the growth trend and details regarding the service.

## GCP Compute: CE Dashboard

Using the **GCP Compute: CE** dashboard you can manage the network, memory, and CPU utilization of your CE Instances.

The **GCP Compute: CE** dashboard helps you monitor the health of all your CE Instances in the GCP inventory and filter them based on projects and regions. This dashboard displays the active and the total number of CE Instances. You can also view the size of the instances and their CPU, network, and memory utilization and distribution at the GCP World level.

The **GCP Compute: CE** dashboard displays the top 10 CE Instances with the highest and the lowest CPU, network, and memory utilization and distribution at the GCP World level. You can filter this information and view it based on projects and regions per project.

The **GCP Compute: CE** dashboard helps you answer several key questions:

- How many CE Instances are there?
- How many CE Instances are running or shut down?
- What alerts are being generated for the CE Instances?
- What are the key metrics for a specific CE Instance?

### How to Use the Dashboard?

The **GCP Compute: CE** dashboard displays the project names, available regions, CE Instances by size, and the total number of CE Instances at the GCP World level. It also displays the top ten CE Instances with the highest and lowest CPU utilization and CPU distribution, the top ten CE Instances with the highest and lowest network bandwidth and network distribution, and the top ten CE Instances with the highest and lowest memory usage and distribution. You can also view the CE Instances according to their power state and duration for which were running and view the alerts for the CE Instances. You can also view the list of CE Instances with key metrics at the GCP world level.

- **Projects** widget: Displays the total number of GCP projects being monitored at the GCP World level. Select a project from the table to refresh the dashboard and view the CE Instance information for the specific project.
- **Regions** widget: Select a region to filter the displayed information and view the data for CE Instances per region.
- **CE Instances with Key Metrics** widget: Displays the list of CE Instances with key metrics at the GCP World level. The list of CE Instances gets updated according to the project and region you select. Select a CE Instance to view its key details, CPU utilization, network bandwidth, and memory usage for a particular CE Instance.

- **Alert List** widget: Displays alerts based on the project and region you select. Click the alert name to view the details of the alert.

## GCP Kubernetes Service Overview Dashboard

Using the **GCP Kubernetes Service Overview** dashboard you can get an overview of your GKE environment.

The **GCP Kubernetes Service Overview** dashboard helps you answer several key questions.

- How many clusters are configured in your environment and what is their average resource utilization?
- How many node pools are associated with a selected cluster?
- How many pods and containers are running on the nodes?

### How to Use the Dashboard?

In the **GCP Kubernetes Service Overview** dashboard, the data is displayed at the GCP World level by default. You can filter Kubernetes clusters by project and by region. In addition, you can also check the node pools associated with a selected cluster, the average resource utilization of the cluster, number of pods, and the containers running on the nodes.

- **Projects** widget: Displays the total number of GCP projects being monitored at the GCP World level. Select a project from the table to refresh the dashboard and view information for a specific project.
- **Regions** widget: Select a region to filter the displayed information and view the data for a specific region.
- **Kubernetes Clusters** widget: Displays the total number of Kubernetes clusters for the selected project and region.
- **Kubernetes Clusters** widget: Displays the list of Kubernetes clusters based on the selected project/region and their details such as name, status, reserved cores, memory limit, and total capacity. The first Kubernetes cluster in the list is selected by default. Select a Kubernetes cluster to view the associated node pools, pods, and containers.
- **Node Pool Count** widget: Displays the total number of node pools for the selected cluster.
- **Disk Utilization Average** widget: Displays the average disk utilization for the selected cluster.
- **Memory Utilization Average** widget: Displays the average memory utilization for the selected cluster.
- **Kubernetes Node Pool** widget: Displays the list of Kubernetes node pools for the selected cluster and their details such as name, version, image type, instance groups, and disk size.
- **Pods and Container Count** widget: Displays the total number of Kubernetes pods and containers for the selected Kubernetes node pool.

- **Kubernetes Pods** widget: Displays the list of Kubernetes pods for the selected Kubernetes node pool and their details such as name, termination grace period, QOS class, and volume capacity.
- **Kubernetes Containers** widget: Displays the list of Kubernetes containers for the selected Kubernetes node pool and their details such as name, memory limit, limit cores, and volume mounts.

## GCP Network Overview Dashboard

Using the **GCP Network Overview** dashboard you can get an overview of your GCP VPCs and key network services linked to it.

The **GCP Network Overview** dashboard helps you answer several key questions.

- What is the network service count for a specific project?
- How many VPC Networks are configured in your environment?
- What are the VPN gateways for a specific VPC Network?
- How many VPN Tunnels and Subnets are connected to a specific VPC Network?

### How to Use the Dashboard?

In the **GCP Network Overview** dashboard, the data is displayed at the GCP World level by default. You can filter them by project. You can view the number of VPCs, Subnets, VPN Gateways, and VPN tunnels. You can also choose a specific VPC and drill down into the network resources linked to the selected VPC.

- **Projects** widget: Displays the total number of GCP projects being monitored at the GCP World level. Select a project from the table to refresh the dashboard and view information for a specific project.
- **Network Services** widget: Displays the total number of VPC Networks, number of Subnetworks, number of VPN Gateways, and the number of VPN tunnels.
- **VPC Networks** widget: Displays the list of VPC Networks. The first VPC in the list is selected by default. The list of VPC Networks gets updated according to the selected project. Select a VPC Network to view the associated VPN gateways, subnets, and VPN tunnels.
- **Services** widget: Displays the total number of VPN gateways, subnets, and VPN tunnels for the selected VPC network.
- **VPN Gateways** widget: Displays the list of VPN gateways for the selected VPC Network and details such as name, region, description, VPN interfaces, and connections.
- **VPN Tunnels** widget: Displays the list of VPN tunnels for the selected VPC Network and VPN Gateway. It displays details such as name, IKE version, region, status, VPN gateway interface, and peer external gateway interface. The first VPN tunnel in the list is selected by default. Select a VPN tunnel to view the associated details in the **Incoming Packets Dropped**, **Outgoing Packets Dropped**, **Received Bytes**, and **Sent Bytes** widgets.

- **Subnets** widget: Displays the list of subnetworks for the selected VPC Network and details such as name, purpose, IP CIDR range, and gateway address.

## GCP Storage Bucket Overview Dashboard

Using the **GCP Storage Bucket Overview** dashboard you can get an overview of the cloud storage resources in your GCP environment.

The **GCP Storage Bucket Overview** dashboard helps you answer several key questions.

- How many Storage Buckets are configured in your environment?
- What are the storage class and location types of the buckets?
- What are the utilization statistics of storage buckets?
- What is the egress/ingress traffic for the top 10 buckets?
- What are the top 10 buckets with the most errors?

### How to Use the Dashboard?

In the **GCP Storage Bucket Overview** dashboard, the data is displayed at the GCP World level by default.You can filter them by project and by region. You can also check utilization statistics, egress/ingress traffic, and errors on buckets, or drill down to a specific bucket.

- **Projects** widget: Displays the total number of GCP projects being monitored at the GCP World level. Select a project from the table to refresh the dashboard and view information for a specific project.
- **Regions** widget: Select a region to filter the displayed information and view the data for a specific region.
- **Storage Buckets** widget: Displays the total number of storage buckets for the selected project and region.
- **Buckets By Storage Class** widget: Displays the number of buckets for storage classes such as standard, regional, nearline, coldline, and archive.
- **Buckets By Location Type** widget: Displays the number of buckets for location type such as region, multi-region, and dual-region.
- **Top 10 Buckets** widget: Displays the top 10 buckets with highest/lowest space usage, high/low egress traffic, high/low ingress traffic, highest object count, and the most 4XX and 5XX errors.
- **Bucket Space Usage Distribution** widget: Displays the number of buckets and their average consumption of bytes.
- **Bucket Egress Traffic Distribution** widget: Displays the number of buckets and their average consumption of sent bytes.
- **Bucket Ingress Traffic Distribution** widget: Displays the number of buckets and their average consumption of received bytes.

- **Storage Bucket List** widget: Displays the list of storage buckets for the selected project and region. The first storage bucket in the list is selected by default. Select a storage bucket to view the associated details in the **Object Count Average**, **Request Count Average**, **Total Size Usage**, **Received Bytes Average**, and **Sent Bytes Average** widgets.

# Widgets in VMware Aria Operations

16

Widgets are the panes on your dashboards. You add widgets to a dashboard to create a dashboard. Widgets display information about attributes, resources, applications, or the overall processes in your environment.

You can configure widgets to reflect your specific needs. The available configuration options vary depending on the widget type. You must configure some of the widgets before they display any data. Many widgets can provide or accept data from one or more widgets. You can use this feature to set the data from one widget as filter and display related information on a single dashboard.

Read the following topics next:


- Widget Interactions
- Configuration Files
- Widget Definitions List


## Widget Interactions

Widget interactions are the configured relationships between widgets in a dashboard where one widget provides information to a receiving widget. When you are using a widget in the dashboard, you select data on one widget to limit the data that appears in another widget, allowing you to focus on a smaller subset data.

### How Interactions Work

If you configured interactions between widget at the dashboard level, you can then select one or more objects in the providing widget to filter the data that appears in the receiving widget, allowing you to focus on data related to an object.

To use the interaction option between the widgets in a dashboard, you configure interactions at the dashboard level. If you do not configure any interactions, the data that appears in the widgets is based on how the widget is configured.

When you configure widget interaction, you specify the providing widget for the receiving widget. For some widgets, you can define two providing widgets, each of which can be used to filter data in the receiving widget.

For example, if you configured the Object List widget to be a provider widget for the Top-N widget, you can select one or more objects in the Object List widget and the Top-N displays data only for the selected objects.

For some widgets, you can define more than one providing widget. For example, you can configure the Metric Chart widget to receive data from a metrics provider widget and an objects providing widget. In such case, the Metric Chart widget shows data for any object that you select in the two provider widgets.

# Configuration Files

You can create configuration files to upload SVG content and define topological hierarchies in dashboards through supported widgets. You can also create configuration files that displays adapter kind, resource kind, and the associated metrics that can be displayed in dashboards using supported widgets.

## Configuration Files for Widget Metric Configuration

You can create an XML file that displays the adapter type, resource kind, and the associated metrics that can be displayed in a widget and dashboard.

### How Configuration Files Work

From the **Configuration Files** page, you can create a XML configuration file with the adapter type, object type, and metric details to be displayed in a dashboard. The final output is an XML file. The supported widgets are Metric Chart, Property List, Rolling View Chart, Scoreboard, Sparkline Chart, and Topology Graph. To use the metric configuration, you must set the widget Self Provider to **Off** and create a widget interaction with a provider widget.

### Where You Find Widget Metric Configuration File

To manage the configuration files, from the left menu, click **Configure > Configuration Files**. From the right panel, click the **Widget Metric Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

Table 16-1. New Configuration File Page Options

| Option | Description |
| --- | --- |
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |

### Table 16-1. New Configuration File Page Options (continued)

| Option | Description |
| --- | --- |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the adapter type, object type, and metrics.<br>From the **Adapter Kind** drop down above the text box, select the adapter type. From the **Resource Kind** drop down above the text box, select the object type. From the **Metric** drop down above the text box, select a metric.<br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

### Table 16-2. Widget Metric Configuration Toolbar and Data Grid Options

| Option | Description |
| --- | --- |
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files. To import:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select the file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the configuration file, and click **Done**. |

Table 16-2. Widget Metric Configuration Toolbar and Data Grid Options (continued)

| Option | Description |
| --- | --- |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a dashboard and supported widget.<br>**Yes**: Indicates that the configuration file is used in a dashboard. Hover and click on the green tick mark to view the name of the dashboard and the widget that uses this configuration file.<br>**No**: Indicates that the configuration file is not used in a dashboard. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, *admin* or *maintenanceAdmin*. |

# Configuration Files for Text Widget Content

You can upload an SVG in Text widgets that are utilized in different dashboards.

## How Configuration Files Work

From the **Configuration Files** page, you can upload an SVG file to be displayed in a Text widget.

## Where You Find Text Widget Configuration File

To manage the configuration files, from the left menu, click **Configure > Configuration Files**. From the right panel, click the **Text Widget Content** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

Table 16-3. New Configuration File Page Options

| Option | Description |
| --- | --- |
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |

### Table 16-3. New Configuration File Page Options (continued)

| Option | Description |
| --- | --- |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | From your system, using an XML editor, open the SVG file to be uploaded. Copy the text from the SVG file into the text box. Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the SVG that was uploaded.<br>**Note** Ensure that the text in the text box starts with `<svg`. |

### Table 16-4. Text Widget Toolbar and Data Grid Options

| Option | Description |
| --- | --- |
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save**. |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files. To import:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select the file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |

Table 16-4. Text Widget Toolbar and Data Grid Options (continued)

| Option | Description |
| --- | --- |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a Text widget and dashboard.<br>**Yes**: Indicates that the configuration file is used in a **Text** widget. Hover and click on the green tick mark to view the name of the dashboard and the Text widget that uses this configuration file.<br>**No**: Indicates that the configuration file is not used in a dashboard. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, *admin* or *maintenanceAdmin*. |

# Management Packs Configuration

You can create an XML file for specific adapter kinds with additional threshold details such as configuration limits, name, type, and value. The configuration file can be used in a widget and dashboard.

## How Configuration Files Work

From the **Management Packs Configuration** page, you can create a XML file with adapter specific details with additional thresholds to be displayed in dashboards and widgets. The final output is an XML file. The final output is an XML file. The supported widgets are Metric Chart, Property List, Rolling View Chart, Scoreboard, Sparkline Chart, and Topology Graph.

## Where You Find Resource Kind Metric Configuration File

To manage the configuration files, from the left menu, click **Configure > Configuration Files**. From the right panel, click the **Solutions Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

### Table 16-5. New Configuration File Page Options

| Option | Description |
| --- | --- |
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the adapter kind and thresholds such as configuration limits, name, type, and value.<br>Select the tags above the text box to add to the configuration file.<br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

### Table 16-6. Solutions Configuration Toolbar and Data Grid Options

| Option | Description |
| --- | --- |
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files. To import:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select the file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the configuration file, and click **Done**. |

Table 16-6. Solutions Configuration Toolbar and Data Grid Options (continued)

| Option | Description |
| --- | --- |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, *admin* or *maintenanceAdmin*. |

# Configuration Files for the Topology Widget

You can create an XML file with topology hierarchies to be displayed in a Topology widget.

## How Configuration Files Work

From the **Configuration Files** page, you can create a XML file with topology hierarchies to be displayed in a Topology widget. The final output is an XML file.

## Where You Find Resource Kind Metric Configuration File

To manage the configuration files, from the left menu, click **Configure > Configuration Files**. From the right panel, click the **Topology Widget Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

Table 16-7. New Configuration File Page Options

| Option | Description |
| --- | --- |
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |

### Table 16-7. New Configuration File Page Options (continued)

| Option | Description |
| --- | --- |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the topology hierarchies.<br>Select the tags above the text box to add to the configuration file.<br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

### Table 16-8. Topology Widget Configuration Toolbar and Data Grid Options

| Option | Description |
| --- | --- |
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them |
| Import | Select this option if you want to import configuration files. To import:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select the file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |

Table 16-8. Topology Widget Configuration Toolbar and Data Grid Options (continued)

| Option | Description |
| --- | --- |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a dashboard and topology widget.<br>**Yes**: Indicates that the configuration file is used in a dashboard. Hover and click on the green tick mark to view the name of the dashboard and the widget that uses this configuration file.<br>**No**: Indicates that the configuration file is not used in a dashboard. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, *admin* or *maintenanceAdmin*. |

# Widget Definitions List

A widget is a pane on a dashboard that contains information about configured attributes, resources, applications, or the overall processes in your environment. Widgets can provide a holistic, end-to-end view of the health of all the objects and applications in your enterprise. If your user account has the necessary access rights, you can add and remove widgets from your dashboards.

Table 16-9. Summary of Widgets

| Widget Name | Description |
| --- | --- |
| Alert List | Shows a list of alerts for the objects that the widget is configured to monitor. If no objects are configure, the list displays all alerts in your environment. |
| Alert Volume | Shows a trend report for the last seven days of alerts generated for the objects it is configured to monitor. |
| Anomalies | Shows a chart of the anomalies count for the past 6 hours. |
| Anomaly Breakdown | Shows the likely root causes for symptoms for a selected resource. |
| Capacity Remaining | Shows a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource. |

### Table 16-9. Summary of Widgets (continued)

| Widget Name | Description |
|---|---|
| Container Details | Shows the health and alert counts for each tier in a single selected container. |
| Container Overview | Shows the overall health and the health of each tier for one or more containers. |
| Current Policy | Shows the highest priority policy applied to a custom group. |
| Data Collection Results | Shows a list of all supported actions specific for a selected object. |
| DRS Cluster Settings | Shows the workload of the available clusters and the associated hosts. |
| Efficiency | Shows the status of the efficiency-related alerts for the objects that it is configured to monitor. Efficiency is based on generated efficiency alerts in your environment. |
| Environment | Lists the number of resources by object or groups them by object type. |
| Environment Overview | Shows the performance status of objects in your virtual environment and their relationships. You can click an object to highlight its related objects and double-click an object to view its Resource Detail page. |
| Environment Status | Shows statistics for the overall monitored environment. |
| Faults | Shows a list of availability and configuration issues for a selected resource. |
| Forensics | Shows how often a metric had a particular value, as a percentage of all values, within a given time period. It can also compare percentages for two time periods. |
| Geo | Shows where your objects are located on a world map, if your configuration assigns values to the Geo Location object tag. |
| Health | Shows the status of the health-related alerts for the objects that it is configured to monitor. Health is based on generated health alerts in your environment. |
| Health Chart | Shows health information for selected resources, or all resources that have a selected tag. |
| Heat Map | Shows a heat map with the performance information for a selected resource. |
| Mashup Chart | Brings together disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs). This widget is typically used for a container. |
| Metric Chart | Shows a chart with the workload of the object over time based on the selected metrics. |
| Metric Picker | Shows a list of available metrics for a selected resource. It works with any widget that can provide resource ID. |
| Object List | Shows a list of all defined resources. |
| Object Relationship | Shows the hierarchy tree for the selected object. |
| Object Relationship (Advanced) | Shows the hierarchy tree for the selected objects. It provides advanced configuration options. |
| Property List | Shows the properties and their values of an object that you select. |
| Recommended Actions | Displays recommendations to solve problems in your vCenter Server instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines. |

Table 16-9. Summary of Widgets (continued)

| Widget Name | Description |
| --- | --- |
| Risk | Shows the status of the risk-related alerts for the objects that it is configured to monitor. Risk is based on generated risk alerts in your environment. |
| Rolling View Chart | Cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget. |
| Scoreboard | Shows values for selected metrics, which are typically KPIs, with color coding for defined value ranges. |
| Scoreboard Health | Shows color-coded health, risk, and efficiency scores for selected resources. |
| Sparkline Chart | Shows graphs that contain metrics for an object . If all the metrics in the Sparkline Chart widget are for an object that another widget provides, the object name appears at the top right of the widget. |
| Tag Picker | Lists all defined resource tags. |
| Text Display | Reads text from a Web page or text file and shows the text in the user interface. |
| Time Remaining | Shows a chart of the Time Remaining values for a specific resource over the past 7 days. |
| Top Alerts | Lists the alerts most likely to negatively affect your environment based on the configured alert type and objects. |
| Top-N | Shows the top or bottom N number metrics or resources in various categories, such as the five applications that have the best or worst health. |
| Topology Graph | Shows multiple levels of resources between nodes. |
| View | Shows a defined view depending on the configured resource. |
| Weather Map | Uses changing colors to show the behavior of a selected metric over time for multiple resources. |
| Workload | Shows workload information for a selected resource. |
| Workload Pattern | Shows a historical view of the hourly workload pattern of an object. |

For more information about the widgets, see the VMware Aria Operations help.

## Alert List Widget

The Alert List widget is a list of alerts for the objects it is configured to monitor. You can create one or more alert lists in VMware Aria Operations for objects that you add to your custom dashboards. The widget provides you with a customized list of alerts on objects in your environment.

## How the Alert List Widget and Configuration Options Work

You can add the Alert List widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. You edit an Alert List widget after you add it to a dashboard. The changes you make to the options create a custom alert list to meet the needs of the dashboard users.

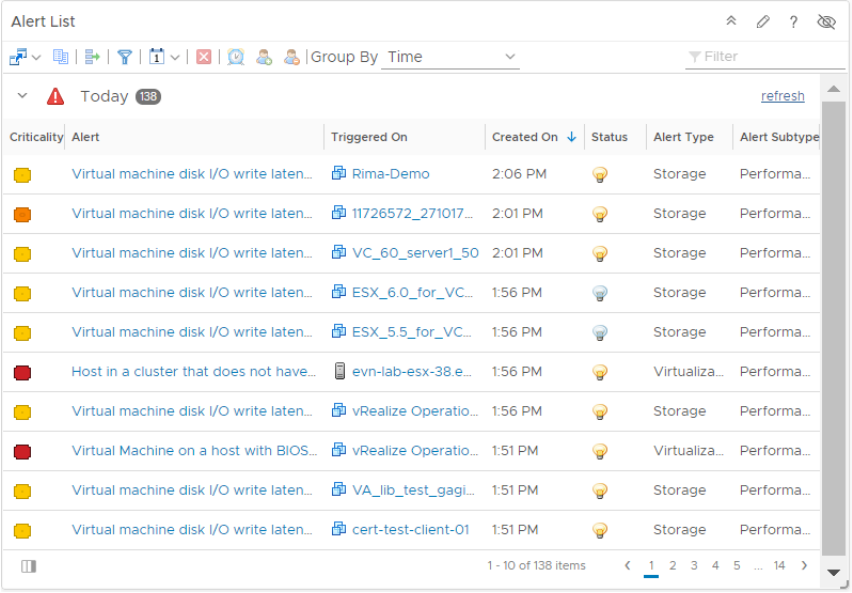


## Where You Find the Alert List Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Alert List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Dashboard Navigation | Actions you can run on the selected alert.<br>For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.<br>Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards. |
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Display Filtering Criteria | Displays the object information on which this widget is based. |
| Select Date Range | Limits the alerts that appear in the list to the selected date range. |
| Cancel Alert | Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.<br>You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |

| Option | Description |
| --- | --- |
| Suspend | Suspend an alert for a specified number of minutes.<br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.<br>The user who suspends the alert becomes the assigned owner.<br>**Note** You can cancel or retrigger the alert, if it is still active when its suspension period has ended, by rerunning the automated actions connected to the alert. In this case, you can suppress cancelation and update on all instances of an alert on an object. To activate this option, open the property file `/usr/lib/vmware-vcops/user/conf/analytics/advanced.properties` and add **`retriggerExpiredSuspendedActiveAlerts = true`** to the property file, and restart the analytics service or the cluster. |
| Take Ownership | As the current user, you make yourself the owner of the alert.<br>You can only take ownership of an alert, you cannot assign ownership. |
| Release Ownership | Alert is released from all ownership. |
| Group By | Group alerts by the options in the drop-down menu. |
| Filter | Locate data in the widget. |

### Table 16-10. Group By Options

| Option | Description |
| --- | --- |
| None | Alerts are not sorted into specific groupings. |
| Time | Group alerts by time triggered. The default. |
| Criticality | Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the Alert List Widget Data Grid table. |
| Definition | Group alerts by definition, that is, group like alerts together. |
| Object Type | Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together. |

## Alert List Widget Data Grid Options

The data grid provides information on which you can sort and search.

Expand the grouped alerts to view the data grid.

| Option | Description |
| --- | --- |
| Criticality | Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon.<br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**. |
| Alert | Description of the alert. |
| Triggered On | Name of the object for which the alert was generated. |
| Created On | Date and time when the alert was generated. |
| Status | Current state of the alert. |
| Alert Type | Alert type is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution.<br>The possible values include:<br>■ Application<br>■ Virtualization/Hypervisor<br>■ Hardware (OSI)<br>■ Storage<br>■ Network |
| Alert Sub-Type | Alert subtype is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution.<br>The possible values include:<br>■ Availability<br>■ Performance<br>■ Capacity<br>■ Compliance<br>■ Configuration |
| Importance | Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm. |

## Alert List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |

| Option | Description |
| --- | --- |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |

| Option | Description |
| --- | --- |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |
| Alert Related | A group of filters limits the alerts that appear in this alert list to those that meet the selected criteria.<br>If the objects on which the alerts are based have an input transformation applied, you define filters for the alerts based on the transformed objects.<br>You can configure the following filters:<br>■ Alert Type. Select the subtype in the type list. This value was assigned when you configured the alert definition.<br>■ Status. Select one or more alert states to include in the list.<br>■ Control State. Select one or more control states to include in the list.<br>■ Criticality. Select one or more levels of criticality.<br>■ Impact. Select one or more alert badges to include in the list. |

## Alert Volume Widget

The Alert Volume widget is a trend report for the last seven days of alerts generated for the objects it is configured to monitor in VMware Aria Operations. You can create one or more alert volume widgets for objects that you add to your dashboards. The alert volume provides you with a customized trend report on objects that helps you identify changes in alert volume, indicating a problem in your environment.

## How the Alert Volume Widget and Configuration Options Work

You can add the Alert Volume widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

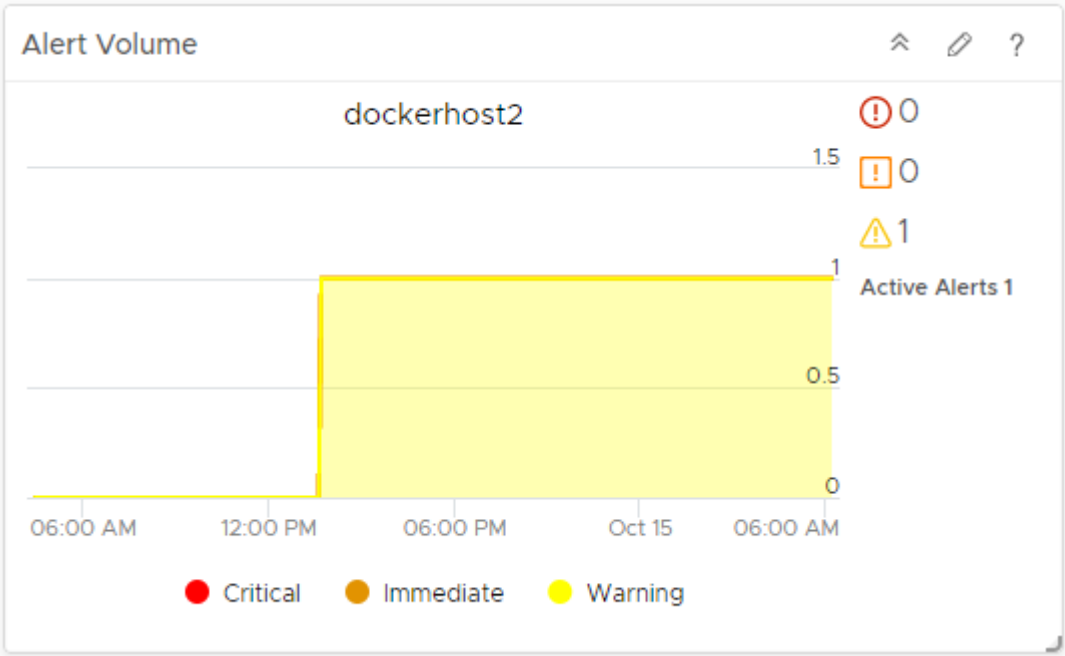


## Where You Find the Alert Volume Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Alert Volume Widget Display Options

The Alert Volume widget displays a trend chart, symptoms by criticality, and active alerts.

| Option | Description |
| --- | --- |
| Trend chart | Volume of critical, immediate, and warning symptoms for the configured objects. |
| Symptoms by criticality | Number of symptoms for each criticality level. |
| Active Alerts | Number of active alerts. Alerts can have more than one triggering symptom. |

## Alert Volume Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

| Option | Description |
| --- | --- |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Anomalies Widget

The Anomalies widget displays the anomalies for a resource for the past 6 hours at time intervals you set.

The Anomalies widget shows or hides time periods when the metric violates a threshold that is configured. The widget color indicates the criticality of the violation.

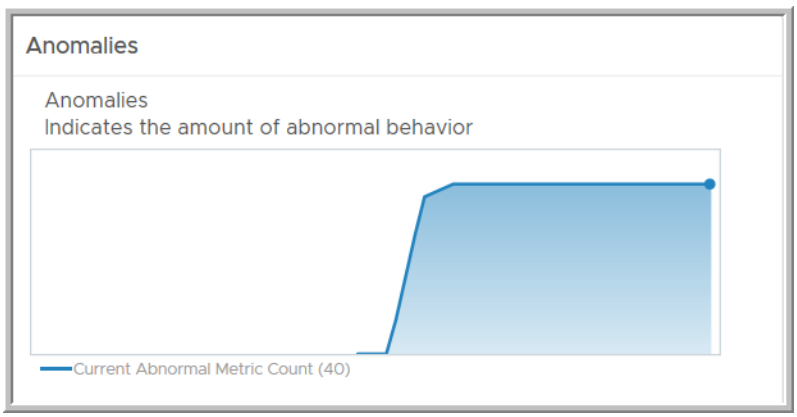


### Where You Find the Anomalies Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize** > **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize** > **Dashboards**. To create your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions** > **Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Anomalies Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | ■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Anomaly Breakdown Widget

The Anomaly Breakdown widget shows the likely root causes for symptoms for a selected resource.

## How the Anomaly Breakdown Widget and Configuration Options Work

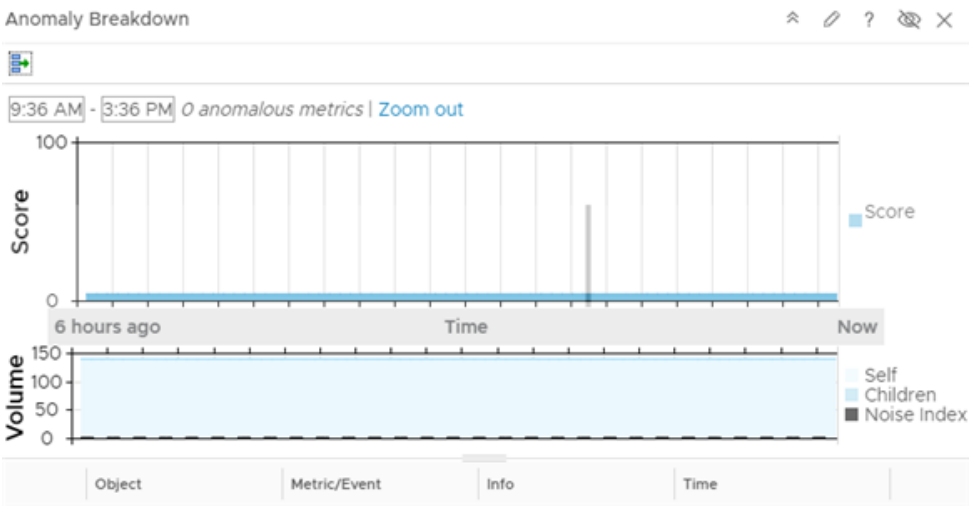


You can add the Anomaly Breakdown widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

## Where You Find the Anomaly Breakdown Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Anomaly Breakdown Widget Display Options

The Anomaly Breakdown widget displays scores, volume, and a list of anomaly metrics.

| Option | Description |
|---|---|
| Score | Anomaly value. |
| Volume | VMware Aria Operations full set metric count for the selected object in the specified time range. |
| Anomaly Metrics List | List of alarms for the selected object in the specified time range. |

## Anomaly Breakdown Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Show Bar Details | If the widget is displaying data for multiple objects, you can select a row and click this button to view the list of alarms for the selected object. |
| Perform Multiple Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |

## Anomaly Breakdown Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

| Option | Description |
|---|---|
| Mode | Display a single object or multiple objects. |
| Show | Select the number of objects to display in multiple objects mode. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |

## Container Details Widget

The Container Details widget displays graphs that show a summary of child objects, metrics, and alerts of an object in the inventory.

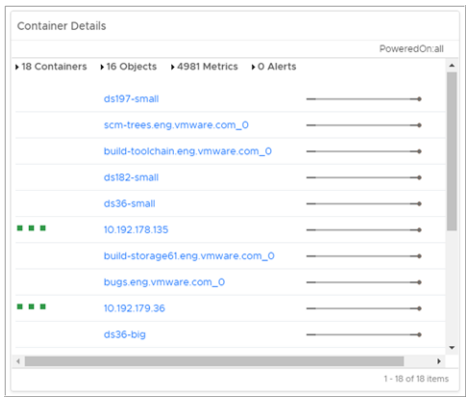


### How the Container Details Widget and Configuration Options Work

The Container Details widget treats objects from the inventory as containers and objects. Containers are objects that contain other objects. The widget lists the containers and shows the number of containers, objects, metrics, and alerts of the observed object. The widget also displays the alerts of each container and an icon links to its child objects. For example, if you select from the inventory a host that contains three objects such as, two virtual machines and one datastore, the Container Details widget displays summary information with three containers, two objects that are the child objects of the two virtual machines, and the number of alerts for

the host and the number of metrics for the child objects of the host. The widget also lists each of the three containers, with the number of alerts for each object. Clicking an object in the graph takes you to the object details page. When you point to the icon next to the object, a tool tip shows the name of the related resource and its health. For example, when you point to the icon next to a virtual machine, the tool tip shows a related datastore and its health. Clicking the icon takes you to the object detail page of the related object, which is the datastore following the example.

You edit a container details widget after you add it to a dashboard. You can configure the widget to take information from another widget in the dashboard and to analyze it. When you select **Off** from the Self Provider option and set source and receiver widgets in the **Widget Interactions** menu during editing of the dashboard, the receiver widget shows information about an object that you select from the source widget. For example, you can configure the Container Details widget to display information about an object that you select from the Object Relationship widget in the same dashboard.

## Where You Find the Container Details Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Container Details Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

| Option | Description |
| --- | --- |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Mode | You can change the size of the graph using the Compact or Large buttons. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Capacity Remaining Widget

The Capacity Remaining widget displays a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.

## Where You Find the Capacity Remaining Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Capacity Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Container Overview Widget

The Container Overview widget gives a graphical presentation of the health, risk, and efficiency of an object or list of objects in the environment.

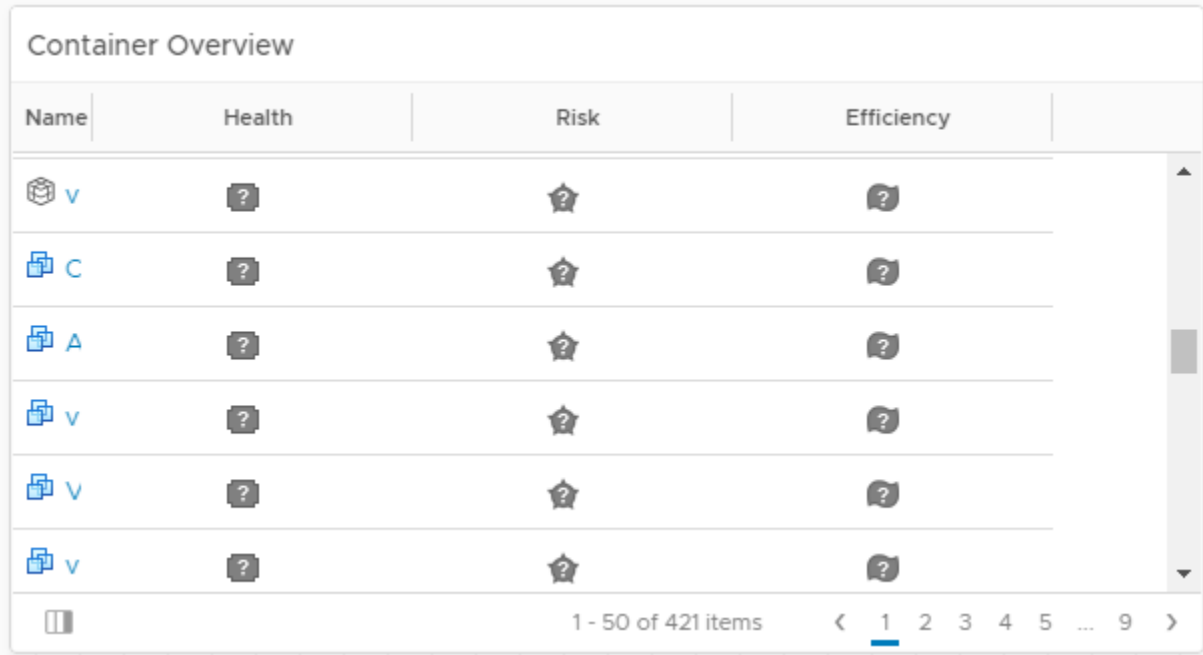


## How the Container Overview Widget and Configuration Options Work

The Container Overview widget displays the current status, the status for a previous time period of the health, risk, and the efficiency of an object or list of objects. You can configure the widget to display information for one or more objects that you are interested in when you select the **Object** mode during configuration of the widget. The widget displays information for all objects from an object type or types when you select the **Object Type** mode during configuration of the widget. You can open the object detailed page of each object in the data grid when you click the object.

You edit a container overview widget after you add it to a dashboard. You can configure the widget to display information about an object or to display information about all objects from an object type by using the **Object** or **Object Type** mode. The configuration options change depending on your selection of mode.

## Where You Find the Container Overview Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize** > **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize** > **Dashboards**. To create your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions** > **Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Container Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about other widgets or dashboards.

| Option | Description |
| --- | --- |
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Filter | You can filter the objects in the data grid. |
| Dashboard Navigation | You can explore information from another dashboard.<br>**Note** This toolbar icon exists when you configure the widget to interact with a widget from another dashboard. Use **Dashboard Navigation** menu during dashboard configuration to configure the widgets to interact.<br>When you select an object from an object data grid and click the toolbar icon, it takes you to a related dashboard. For example, you can configure the widget to send information to a Topology Graph widget that is on another dashboard, for example dashboard 1. When you select a VM from the data grid, click **Perform Multi-Select Interaction** , click **Dashboard Navigation** and select **Navigate > dashboard 1**. It takes you to dashboard 1, where you can observe selected VM and objects related to it. |

## Container Overview Widget Data Grid Options

The data grid provides information on which you can sort and search.

| Option | Description |
| --- | --- |
| Name | Name of the object |
| Health | Shows information about the health parameter.<br>Status displays the badge of the current health status of an object. You can check the status in a tool tip when you point to the badge.<br>Last 24 Hours displays the statistic of health parameter for last 24 hours. |

| Option | Description |
| --- | --- |
| Risk | Shows information about the risk parameter.<br>Status displays the badge of the current risk status of an object. You can check the status in a tool tip when you point to the badge.<br>Last Week displays the statistics of the health parameter for the last week. |
| Efficiency | Shows information about the efficiency parameter.<br>Status displays the badge of the current efficiency status of an object. You can check the status in a tool tip when you point to the badge.<br>Last Week displays statistic of the efficiency parameter for the last week. |

## Container Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Mode | Use **Object** to select an object from the environment to observe.<br>Use **Object Type** to select the type of the objects to observe. |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

| Option | Description |
| --- | --- |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| Object Type | Select an object type in your environment on which you want to base the widget data.<br>1 Click the **Add Object Type** icon to search for and add an object type.<br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.<br>2 Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type. |

## Current Policy Widget

The Current Policy widget displays the active operational policy that is assigned to your object or object group. VMware Aria Operations uses the assigned policy to analyze your objects, control

the data that is collected from those objects, generate alerts when problems occur, and display the results in the dashboards.

## How the Current Policy Widget and Configuration Options Work

You add the Current Policy widget to a dashboard so that you can quickly see which operational policy is applied to an object or object group. To add the widget to a dashboard, you must have access permissions associated with the roles assigned to your user account.

The configuration changes that you make to the widget creates a custom instance of the widget that you use in your dashboard to identify the current policy assigned to an object or object group. When you select an object on the dashboard, the policy applied to the object appears in the Current Policy widget, with an embedded link to the policy details. To display the inherited and local settings for the applied policy, click the link.

## Where You Find the Current Policy Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Current Policy Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |

| Option | Description |
|---|---|
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.<br>For example, to view the policy applied to each object that you select in the Object List widget, select **Off** for Self Provider. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Data Collection Results Widget

The Data Collection Result widget shows a list of all supported actions specific for a selected object. The widget retrieves data specific to a selected object actions and uses the action framework to run data collection actions.

## How the Data Collection Results Widget and Configuration Options Work

You can add the Data Collection Results widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The Data Collection Results widget is a receiver of a resource or metric ID. It can interact with any resource or metric ID that provides widgets such as Object List and Metric Picker. To use the widget, you must have an environment that contains the following items.

- A vCenter Adapter instance
- A VMware Aria Operations for Horizon View Adapter
- A VMware Aria Operations for Horizon View Connection Server

You edit a Data Collection Result widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

## Where You Find the Data Collection Results Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Data Collection Results Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Results | Shows all finished and currently running actions for the selected object. |
| Choose Action | Shows a list with all supported actions specific for the selected object. The selected object is a result of widget interactions. |

## Data Collection Results Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget updates only when you open the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

| Option | Description |
| --- | --- |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Config | Specifies self provider choice and selection of a resource instance. |
| Selected Object | When you select an object, this text box is populated by the object. |
| Start new data collection on interaction change | Indicates whether to start a new data collection action when the object selection changes in the source widget. |
| Objects | List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget. |
| Defaults | Specifies the default data collection action selected for each object type. |
| Object Types | List of object types in your environment that you can search or sort by column so that you can locate the object type on which you are basing the data that appears in the widget. You can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. |
| Default Data Collection Action | This panel is populated by the object type that you select in the object types list.<br>You can select only one default data collection action for an object type. |

# DRS Cluster Settings Widget

The DRS Cluster Settings widget displays the workload of the available clusters and the associated hosts. You can change the Distributed Resource Scheduler (DRS) automation rules for each cluster.

## How the DRS Cluster Settings Widget and Configuration Options Work

You can view CPU workload and memory workload percentages for each of the clusters. You can view CPU workload and memory workload percentages for each host in the cluster by selecting a cluster in the data grid. The details are displayed in the data grid below. You can set the level of DRS automation and the migration threshold by selecting a cluster and clicking **Cluster Actions > Set DRS Automation**.

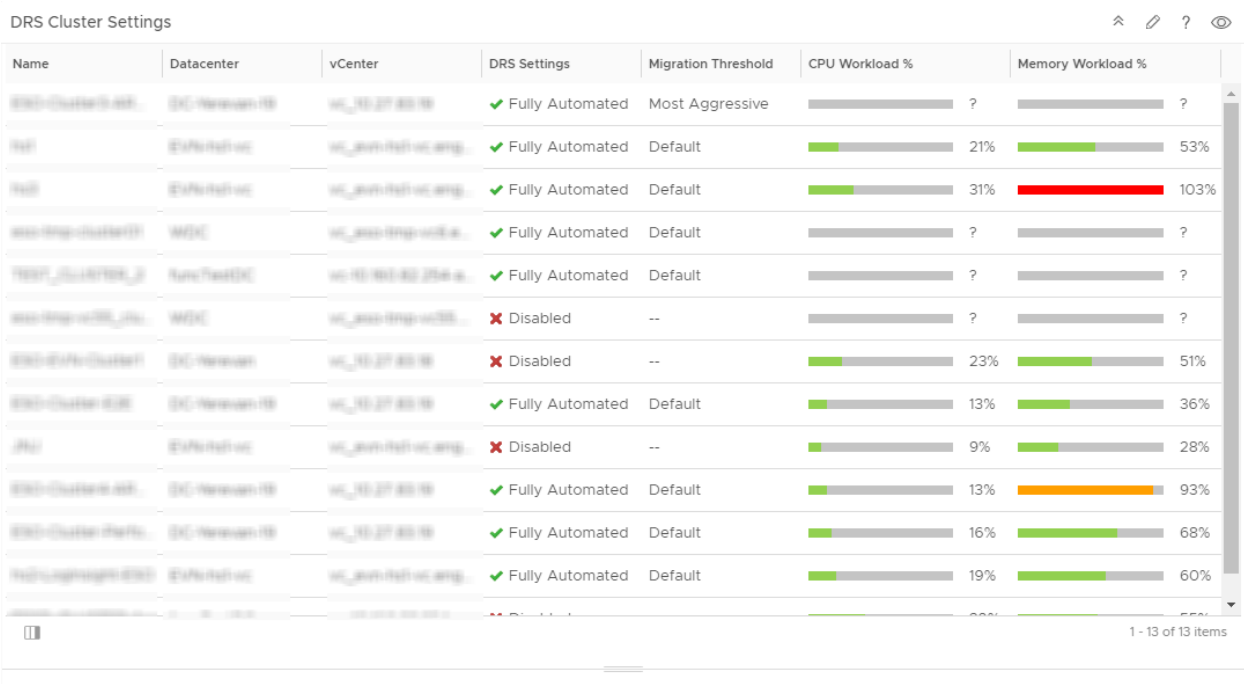


You edit a DRS Cluster Settings widget after you add it to a dashboard. To configure the widget, click the edit icon at the upper-right corner of the widget window. You can add the DRS Cluster Settings widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The DRS Cluster Settings widget appears on the dashboard named vSphere DRS Cluster Settings, which is provided with VMware Aria Operations.

## Where You Find the DRS Cluster Settings Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## DRS Cluster Settings Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Cluster Actions | Limits the list to actions that match the cluster you select. |
| Show | The drop-down menu displays the parent vCenter Server instances where the clusters reside. You can also view the data centers under each parent vCenter Server instance. Select a parent vCenter Server to view the workload of the available clusters in the data grid.<br>The default setting displays the clusters across all vCenters. |
| Filter | Filters the data grid by name, data center, vCenter, DRS settings, and migration threshold. |

## DRS Cluster Settings Widget Data Grid Options

The data grid provides information on which you can sort and search.

| Option | Description |
| --- | --- |
| Name | Displays the names of the clusters in the selected parent vCenter Server instance. |
| Datacenter | Displays the data centers that belong to each cluster. |
| vCenter | Displays the parent vCenter Server instance where the cluster resides. |
| DRS Settings | Displays the level of DRS automation for the cluster.<br>To change the level of DRS automation for the cluster, select **Cluster Actions > Set DRS Automation** from the toolbar. You can change the automation level by selecting an option from the drop-down menu in the Automation Level column. |
| Migration Threshold | Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. |
| CPU Workload % | Displays the percentage of CPU in GHz available on the cluster. |
| Memory Workload % | Displays the percentage of memory in GB available on the cluster. |

## DRS Cluster Settings Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

# Efficiency Widget

The efficiency widget is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in VMware Aria Operations usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.

## How the Efficiency Widget and Configuration Options Work

You can add the efficiency widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning efficiency alerts generated over time, if the monitored object is a group.
- A trend line displays the efficiency status of the monitored object over time if the object does not provide its resources to any other object, or where no other object depends on the monitored object's resources. For example, if the monitored object is a virtual machine or a distributed switch.
- A pie chart displays the reclaimable, stress, and optimal percentages for the virtual machines that are descendants of the monitored object for all other object types. You use the chart to identify objects in your environment from which you can reclaim resources. For example, if the object is a host or datastore.

If the **Badge Mode** is set to **On**, only the badge appears.

Edit an efficiency widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

## Where You Find the Efficiency Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Efficiency Widget Display Options

The Efficiency widget displays an efficiency badge. The widget also displays an efficiency trend when not in badge mode.

| Option | Description |
| --- | --- |
| Efficiency Badge | Status of the objects configured for this instance of the widget.<br>Click the badge to open the **Alerts** tab for the object that provides data to the widget. |
| Efficiency Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears. |

## Efficiency Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br>Select one of the following options:<br>■ On. Only the badge appears in the widget.<br>■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Environment Widget

The Environment widget displays the resources which collects data. You can create one or more lists in VMware Aria Operations for the resources that you add to your custom dashboards.

## How the Environment Widget and Configuration Options Work

The Environment widget lists the number of resources by object or groups them by object type. You can add the Environment widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an Environment widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

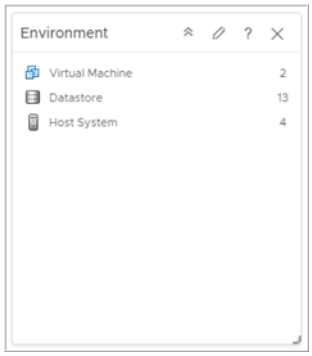


## Where You Find the Environment Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Environment Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |

| Option | Description |
|---|---|
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Environment Overview Widget

The Environment Overview widget displays the health, risk, and efficiency of resources for a given object from the managed inventory.

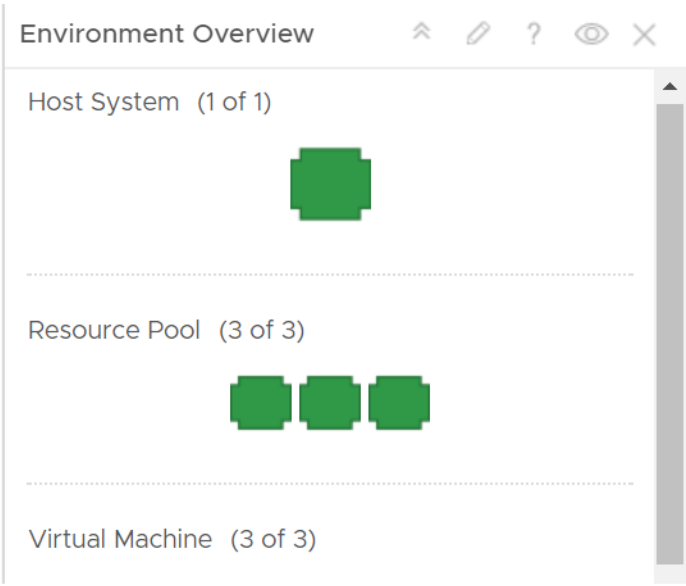


### How the Environment Overview Widget and Configuration Options Work

You can add the Environment Overview widget to one or more custom dashboards.

The widget displays data for objects from one or several types. The data that the widget displays depends on the object type and category that you selected when you configured the widget.

The objects in the widget are ordered by object type.

The parameters for the health, risk, and efficiency of an object appear in a tool tip when you point to the object.

When you double-click an object on the Environment Overview widget, you can view detailed information for the object.

To use the Environment Overview widget, you must add it to the dashboard and configure the data that appears in the widget. You must select at least one badge and an object. Additionally, you can select an object type.

The Environment Overview widget has basic and advanced configuration options. The basic configuration options are activated by default.

To use all features of the Environment Overview widget, you must change the default configuration of the widget. Log in to the VMware Aria Operations machine and set `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. The change is propagated after you use the `service vmware-vcops-web restart` command to restart the UI.

You must use the **Badge** tab to select the badge parameters that the widget shows for each object. You must use the **Config** tab to select an object or object type. To observe a concrete object from the inventory, you can use the **Basic** option. To observe a group of objects or objects from different types, you must use the **Advanced** option.

## Where You Find the Environment Overview Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Environment Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about badges.

| Option | Description |
| --- | --- |
| Badge | You can select a Health, Risk, or Efficiency badge for objects that appear in the widget. The tool tip of a badge shows the standard name of the badge. |
| Status | You can filter objects based on their badge status and their state. |
| Sort | You can sort objects by letter or by number. |

## Environment Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| Selected Object | Object that is the basis for the widget data.<br>To populate the text box, select **Config > Basic** and select an object from the list. |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Badge | Defines a parameter to observe. You can select or deselect Health, Risk, and Efficiency parameters using check boxes. Default configuration of the widget selects all badges.<br>Select at least one badge parameter. |

| Option | Description |
| --- | --- |
| Config | **Basic**<br>List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget. |
| | **Advanced**<br>You can use Object Types to select a type of the objects to observe information about health, risk, and efficiency. Double-click the object type to select it.<br>Use the **Adapter Type** drop-down menu to filter the objects types based on an adapter.<br>You can use the **Use vSphere Default** button to observe the main vSphere object types.<br>To remove an object type from the list, click **Remove Selected** next to **Use vSphere Default**.<br>You can use the **Object Type Categories** menu to select a group or groups of object types to observe.<br>You can use the Object tree to select an object to filter the displayed objects. For example, to observe a datastore of a VM, double-click **Datastore** from the **Object Types** menu to select it. Click the datastore when it is in the list of object types, and find the VM in the object tree and select it. To return to your previous configuration of the widget, click **Datastore** from the list of object types and click **Deselect All** in the object tree window.<br>The metrics tree and badge data grids are available configuration options only if the default configuration of the widget is changed. To use these configuration options, log in to the VMware Aria Operations machine and set `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. |

## Environment Status Widget

The Environment Status widget displays the statistics for the overall monitored environment.

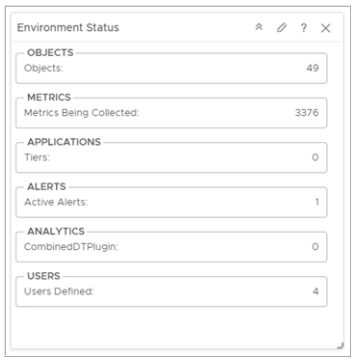

## How the Environment Status Widget and Configuration Options Work

You customize the output of the widget by choosing a category such as Objects, Metrics, Applications, Alerts, Analytics, and Users. You can filter the data by using the tags tree from **Select which tags to filter** in the configuration window.

You edit an environment status widget after you add it to a dashboard. To configure the widget, click the pencil at the right corner of the widget window. You must select at least one type of information from **OBJECTS**, **METRICS**, **APPLICATIONS**, **ALERTS**, **ANALYTICS**, **USERS** categories for the widget to display. By default, the widget displays statistics information about all objects in the inventory. You can use the Select which tags to filter option to filter the information. The widget can interact with other widgets in the dashboard, taking data from them and displaying statistics . For example, you can have a Object List widget , which is the source of the data and an Environment Status widget, which is the destination. If you select objects and perform a multiselection interaction from the Object List widget, the Environment Status widget results are updated based on the selections you made in the Object List.

## Where You Find the Environment Status Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Environment Status Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.<br>The widget is also updated when it is in interaction mode. For example, when an item is selected in the provider widget, the content of the Environment Status widgets is refreshed. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |

| Option | Description |
| --- | --- |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Objects | The widget shows summarized information about the objects in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of resources. For example, if you select **Adapter Types > Container** from Select which tag to filter and click **Objects** and **Objects Collecting** , the widget displays the number of containers and collecting containers. |
| Metrics | The widget shows summarized information about available metrics. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of metrics. |
| Applications | The widget shows summarized information about available applications. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of applications. |
| Alerts | The widget shows summarized information about alerts in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of alerts. |
| Analytics | The widget shows summarized information about the analytics plug-ins. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of analytics. |
| Users | The widget shows the number of users defined in VMware Aria Operations. Select **Administration > Access Control > User Accounts**. |
| **Output Filter** | |

| Option | Description |
| --- | --- |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Faults Widget

The Faults widget displays detailed information about faults experienced by an object

The Faults widget configuration options are used to customize each instance of the widget that you add to your dashboards.

## Where You Find the Faults Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Faults Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Forensics Widget

The Forensics widget shows how often a metric has a particular value as a percentage of all values, within a given time period. It can also compare percentages for two time periods.

## How the Forensics Widget and Configuration Options Work

You can add the Forensics widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit the Forensics widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

## Where you Find the Forensics Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Forensics Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

| Option | Description |
| --- | --- |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Percentile | Indicates how much data is above or below the specific value. For example, it indicates that 90% of the data is more than 4 when a vertical line occurs on the value 4. |
| **Input Data** | |
| | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list. |

## Geo Widget

If your configuration assigns values to the Geo Location object tag, the geo widget shows where your objects are located on a world map. The geo widget is similar to the **Geographical** tab on the Inventory page.

## How the Geo Widget and Configuration Options Work

You can move the map and zoom in or out by using the controls on the map. The icons at each location show the health of each object that has the Geo Location tag value. You can add the geo widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a Geo widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

## Where You Find the Geo Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Geo Widget Toolbar Options

| Option | Description |
| --- | --- |
| Zoom in | Zooms in on the map. |
| Zoom out | Zooms out on the map. |

## Geo Widget Configuration Options

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

| Option | Description |
| --- | --- |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Heatmap Widget

The Heatmap widget contains graphical indicators that display the current value of two selected attributes of objects of tag values that you select. In most cases, you can select only from internally generated attributes that describe the general operation of the objects, such as health or the active anomaly count. When you select a single object, you can select any metric for that object.

## How the Heatmap Widget and Configuration Options Work

You can add the Heatmap widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The Heatmap widget has a General mode and an Instance mode. The General mode shows a colored rectangle for each selected resource. In the Instance mode, each rectangle represents a single instance of the selected metric for an object.

You can click a color or the size metric box in the bottom of the Heatmap widget to filter the display of cells in the widget. You can click and drag the color filter to select a range of colors. The Heatmap widget displays cells that match the range of colors.

When you point to a rectangle for an object, the widget shows the resource name, group-by values, the current values of the two tracked attributes, virtual machine details, the metric name, and the value of the color. Click **Show Sparkline** to view the value.

You edit a Heatmap widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

## Where You Find the Heatmap Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Heatmap Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Dashboard Navigation | Actions you can run on the selected alert.<br>For example, you use the option to open a vCenter Server, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems. |
| Group Zoom | You can roll-up non-significant resources with similar characteristics into groups to obtain only the relevant data among the thousands of resources in the system. The roll-up method improves performance and decreases the memory usage. The roll-up box encompasses the average color and the sum of the sizes of all the resources. You can view all the resources by zooming in the roll-up box. |
| Show/Hide Text | Show or hide the cell name on the heatmap rectangle. |
| Show Details | If you configure the Heatmap widget as a provider to another widget, such as the Metric Chart widget , you can double-click a rectangle to select that object for the widget. If the widget is in Metric mode, double-clicking a rectangle selects the resource associated with the metric and provides that resource to the receiving widget. Optionally, you can select a cell from the heatmap and click the **Show Details** icon to see details about the cell. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget. |
| Reset Zoom | Resets the heatmap display to fit in the available space. |
| Heatmap Configuration Drop-down | Select from a list of predefined heatmaps. |

## Heatmap Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |

| Option | Description |
| --- | --- |
| **Output Data** | |
| Configurations | List of saved heatmap configuration options. You can create a configuration and save it in the list. From the options on the right, you can also delete, clone, and reorder the configurations. |
| Name | Name of the widget. |
| Group by | First-level grouping of the objects in the heatmap. |
| Then by | Second-level grouping of the objects in the heatmap. |
| Relational Grouping | After you select the Group by and Then by objects, select the **Relational Grouping** check box to reorganize the grouping of the objects, and to relate the objects selected in the Group by text box with the objects selected in the Then by text box. |
| Mode | **General mode**<br>The widget shows a colored rectangle for each selected resource. The size of the rectangle indicates the value of one selected attribute. The color of the rectangle indicates the value of another selected attribute.<br>**Instance mode**<br>Each rectangle represents a single instance of the selected metric for a resource. A resource can have multiple instances of the same metric. The rectangles are all the same size. The color of the rectangles varies based on the instance value. You can use instance mode only if you select a single resource kind. |
| Object Type | Object that is the basis for the widget data. |
| Size by | An attribute to set the size of the rectangle for each resource.<br>Resources that have higher values for the Size By attribute have larger areas of the widget display. You can also select fixed-size rectangles. In most cases, the attribute lists include only metrics that VMware Aria Operations generates. If you select a resource kind, the list shows all the attributes that are defined for the resource kind. |
| Color by | An attribute to set the color of the rectangle for each resource. |

| Option | Description |
| --- | --- |
| Solid Coloring | Select this option to use solid colors instead of a color gradient. By default, the widget assigns red color for high value, brown color for intermediate value and green color for low value. Click the color box to set a different color for the values. You can add up to seven color thresholds by clicking color range. |
| Color | Shows the color range for high, intermediate and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes. By default, green indicates a low value and red indicates the high end of the value range. You can change the high and low values to any color and set the color to use for the midpoint of the range. You can also set the values to use for either end of the color range, or let VMware Aria Operations define the colors based on the range of values for the attribute.<br>If you leave the text boxes blank, VMware Aria Operations maps the highest and lowest values for the **Color By** metric to the end colors. If you set a minimum or maximum value, any metric at or beyond that value appears in the end color. |
| **Output Filter** | |

| Option | Description |
| --- | --- |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Health Widget

The Health widget is the status of the health-related alerts for the objects it is configured to monitor in VMware Aria Operations. Health alerts usually require immediate attention. You can create one or more health widgets for different objects that you add to your custom dashboards.

## How the Health Widget and Configuration Options Work

You can add the Health widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

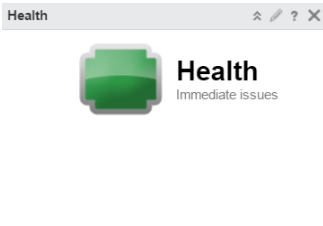


The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A trend line displays the health status of the monitored object if the object does not provide its resources to any other object. For example, if the monitored object is a virtual machine or a distributed switch.
- A weather map displays the health of the ancestor and descendant objects of the monitored object for all other object types. For example, if the monitored object is a host that provides CPU and memory to a virtual machine.

If the **Badge Mode** is set to **On**, only the badge appears.

You edit a Health widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

## Where You Find the Health Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Heath Widget Display Options

The Health widget displays a health badge. The widget also displays a health trend when not in badge mode.

| Option | Description |
|---|---|
| Health Badge | Status of the objects configured for this instance of the widget.<br>Click the badge to open the **Alerts** tab for the object that provides data to the widget.<br>If the **Badge Mode** option is off, a health weather map or trend chart appears for the object. Whether the map or chart appears depends on the object type. The health weather map displays tool tips for up to1000 objects. |
| Health Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears. |

## Heath Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

| Option | Description |
| --- | --- |
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br>Select one of the following options:<br>■ On. Only the badge appears in the widget.<br>■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Health Chart Widget

The Health Chart widget displays Health, Risk, Efficiency, or custom metric charts for selected objects. You use the widget to compare the status of similar objects based on the same value or name.

## How the Health Chart Widget and Configuration Options Work

You can add the Health Chart widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

If the widget is configured to display Health, Risk, or Efficiency, the chart values are based on the generated alerts for the selected alert type for the selected objects.

If the widget is configured to display custom metrics, chart values are based on the metric value for the configured time period.

You edit the Health Chart widget after you add it to the dashboard. The changes you make to the options create a custom widget with the selected charts.

The charts are based either on Health, Risk, or Efficiency alert status, or you can base them on a selected metric. You can include a single object, multiple objects, or all objects of a selected type.

To view the value of the object at a particular time, point your cursor over the chart. A date range and metric value tool tip appear.

A context drop-down menu for each chart can be accessed at the top-right corner after the last metric value.

For each chart, you can view the minimum, maximum, and last metric values. The values are displayed at the top-right corner of each chart. Each of the values is preceded by an appropriate icon of the same color as the state of the metric value.

If there is not enough space to view the metric values, a blue information icon is displayed. Point your cursor over the icon to view the metric value details.

## Where You Find the Health Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Health Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Date Controls | Use the date selector to limit the data that appears in each chart to the time period you are examining.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br>**Dashboard Time** is the default option. |

## Health Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

| Option | Description |
| --- | --- |
| Close | Deletes the chart. |
| Save a snapshot | Creates a PNG file of the current chart. The image is the size that appears on your screen.<br>You can retrieve the file in your browser's download folder. |
| Save a full screen snapshot | Downloads the current graph image as a full-page PNG file, which you can display or save.<br>You can retrieve the file in your browser's download folder. |

| Option | Description |
|---|---|
| Download comma-separated data | Creates a CSV file that includes the data in the current chart.<br>You can retrieve the file in your browser's download folder. |
| Units | Select the units in which the widget displays data. This option is visible when you select a custom source of data in the widget configuration. |

## Health Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Order By | Determines how the object charts appear in the widget.<br>You can order them based on value or name, and in ascending or descending order. |
| Chart Height | Controls the height of all charts. Choose from three possible choices - Small, Medium, Large. Default is Medium. |

| Option | Description |
| --- | --- |
| Pagination number | Number of charts that appears on a page.<br>If you prefer scrolling through the charts, select a higher number. If you prefer to page through the results, select a lower number. |
| Auto Select First Row | Determines whether to start with the first row of data. |
| Metric | Determines the source of the data.<br>■ Health, Risk, or Efficiency. The displayed charts are based on one of these alert badges.<br>■ Custom. The displayed charts are based on the selected metric and use either alert symptom state colors or the selected custom color. You can select a unit for the custom metric from the drop-down menu or choose to allow the widget to automatically pick a unit.<br>If you apply custom colors, enter the value in each box that is the highest or lowest value that should be that color. You can select a unit for the metric. |
| Metric Unit | Select a unit for the custom metric. |
| Show | Select one or more of the following items to display in the widget:<br>■ Select **Object Name** to display the name of the object in the widget.<br>■ Select **Metric Name** to display the name of the metric in the widget. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter** | |

| Option | Description |
| --- | --- |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

## Log Analysis Widget

The Log Analysis widget displays the logs for objects using VMware Aria Operations for Logs. You can configure the objects for which data appears in the widget, or provide objects using the dashboard widget interactions options.

The following screenshot shows the log analysis widget.

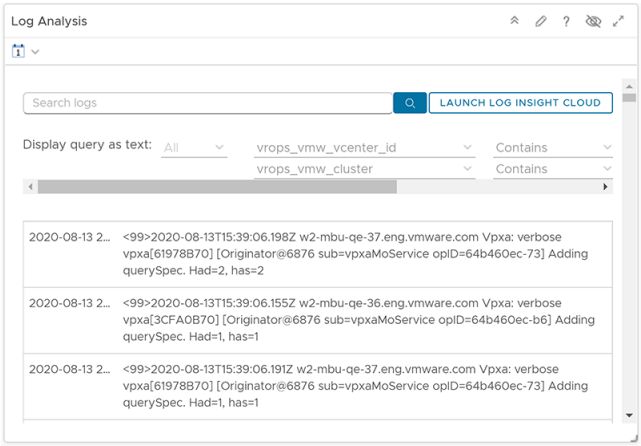


You can view, filter, and search the logs that are displayed. For in-depth analysis of the logs, you can run VMware Aria Operations for Logs.

## Where You Find the Log Analysis Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Log Analysis Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

| Option | Description |
| --- | --- |
| Self Provider | ■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Note** Do not use the log analysis form widget for reporting. The widget requires user inputs to show results. The widget is interactive and performs cross-product queries.

## Mashup Chart Widget

The Mashup Chart widget shows disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs).

### How the Mashup Chart Widget and Configuration Options Work

The Mashup Chart widget contains charts that show different aspects of the behavior of a selected resource. By default, the charts show data for the past six hours.

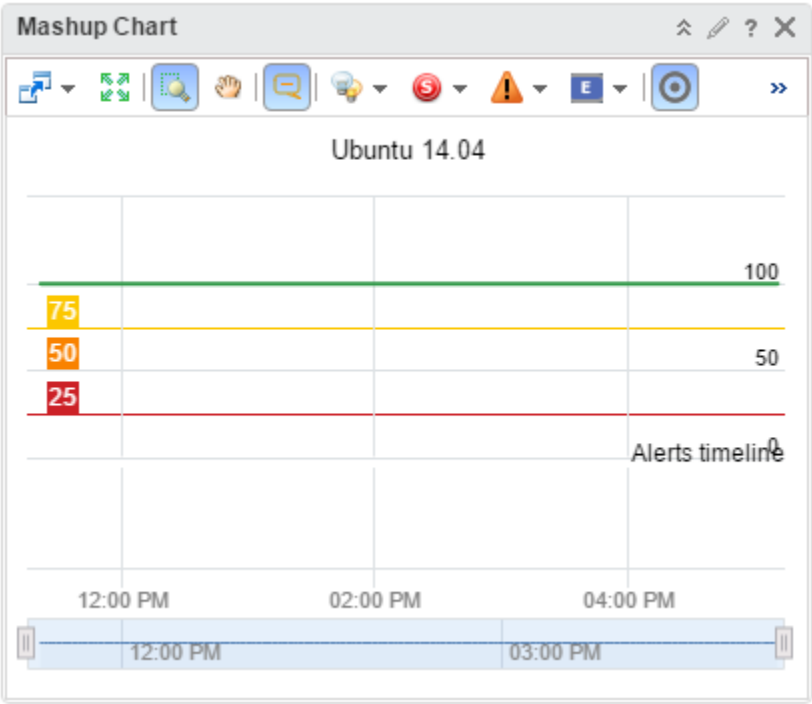

The Mashup Chart widget contains the following charts.

- A Health chart for the object, which can include each alert for the specified time period. Click an alert to see more information, or double-click an alert to open the Alert Summary page.
- Metric graphs for any or all the KPIs for any objects listed as a root cause object. For an application, this chart shows the application and any tiers that contain root causes. You can select the KPI to include by selecting **Chart Controls > KPIs** on the widget toolbar. Any shared area on a graph indicates that the KPI violated its threshold during that time period.

The metric graphs reflect up to five levels of resources, including the selected object and four child levels.

You edit a Mashup Chart widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

## Where You Find the Mashup Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Mashup Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view.

| Option | Description |
| --- | --- |
| Filters | Filter data based on criticality, status, and alert type. |
| Event Filters | Filter based on the type of event such as, change, notification, and fault. |
| Date Controls | Use the date selector to limit the data that appears in each chart to the time period you are examining.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br>**Dashboard Time** is the default option. |
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. |

### Mashup Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Metric Chart Widget

You can use the Metric Chart widget to monitor the workload of your objects over time. The widget displays data based on the metrics that you select.

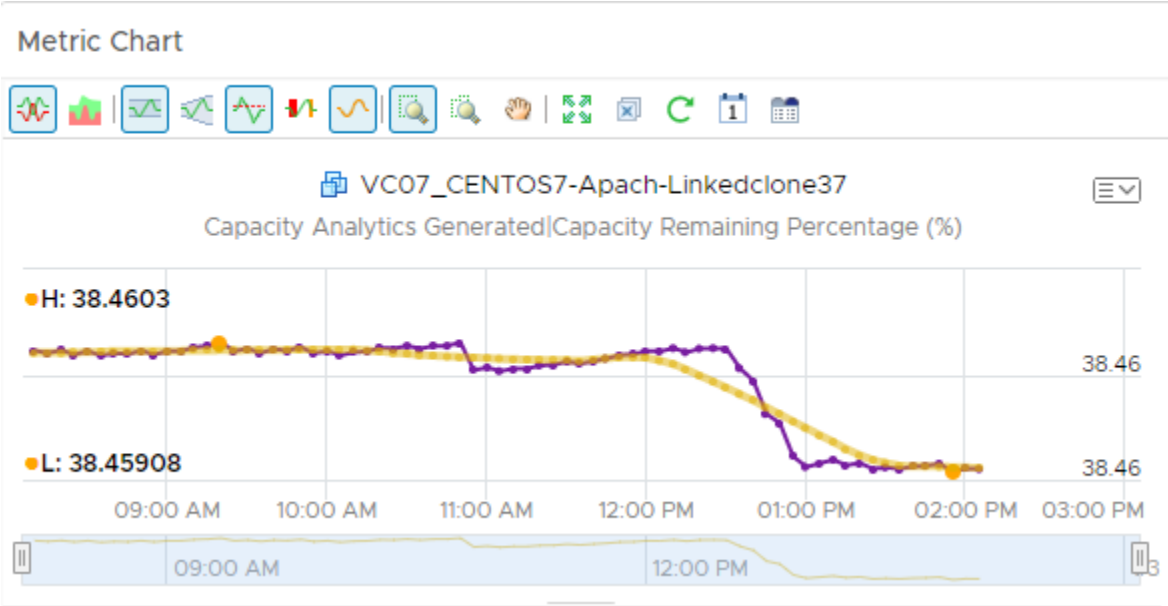


## How the Metric Chart Widget and Configuration Options Work

You can add the Metric Chart widget to one or more custom dashboards and configure it to display the workload for your objects. The data that appears in the widget is based on the configured menu items for each widget instance.

You edit the Metric Chart widget after you add it to a dashboard. The changes you make to the menu items create a custom widget with the selected metrics that display the workload on your objects.

To select metrics, you can select an object from the object list, then select the metrics. Or, you can select a tag from the object tag list to limit the object list, then select an object. You can configure multiple charts for the same object or multiple charts for different objects.

To use the metric configuration, which displays a set of metrics that you defined in an XML file, the dashboard and widget configuration must meet the following criteria:

- The dashboard **Widget Interaction** menu items are configured so that another widget provides objects to the target widget. For example, an Object List widget provides the object interaction to a chart widget.
- The widget **Self Provider** options are set to **Off**.
- The custom XML file in the **Metric Configuration** drop-down menu is in the `/usr/lib/vmware-vcops/tools/opscli` directory and has been imported into the global storage using the import command.

## Where You Find the Metric Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize** > **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Metric Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
|---|---|
| Split Charts | Displays each metric in a separate chart. |
| Stacked Chart | Consolidates all charts into one chart. This chart is useful for seeing how the total or sum of the metric values vary over time. To view the stacked chart, ensure that the split chart option is turned off. |
| Dynamic Thresholds | Shows or hides the calculated dynamic threshold values for a 24-hour period. |
| Show Entire Period Dynamic Thresholds | Shows or hides dynamic thresholds for the entire time period of the graph. |
| Static Thresholds | Shows or hides the threshold values that have been set for a single metric. |
| Anomalies | Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below. |
| Trend Line | Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points. |
| Show Data Values | Activates the data point tooltips if you switched to a zoom or pan option. **Show Data Point Tips** must be activated. |
| Zoom All Charts | Resizes all the charts that are open in the chart pane based on the area captured when you use the range selector.<br>You can switch between this option and **Zoom the View**. |
| Zoom the View | Resizes the current chart when you use the range selector. |
| Pan | When you are in zoom mode, allows you to drag the enlarged section of the chart so that you can view higher or lower, earlier or later values for the metric. |
| Zoom to Fit | Resets the chart to fit in the available space. |
| Remove All | Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts. |
| Refresh Charts | Reloads the charts with current data. |

| Option | Description |
| --- | --- |
| Date Controls | Opens the date selector.<br>Use the date selector to limit the data that appears in each chart to the time period you are examining.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br>**Dashboard Time** is the default option. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Generate Dashboard | Saves the current charts as a dashboard. |
| Save as PDF | Saves the current metric chart as a PDF file.<br>If you have added many metrics for different objects to a metric chart, you can quickly download and share the PDF file with another user or VM owner. You can create ad-hoc reports when analyzing metrics. You can also add notes to the metrics to provide context to the user. |

## Metric Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

| Option | Description |
| --- | --- |
| Close | Deletes the chart. |
| Save a snapshot | Creates a PNG file of the current chart. The image is the size that appears on your screen.<br>You can retrieve the file in your browser's download folder. |
| Download comma-separated data | Creates a CSV file that includes the data in the current chart.<br>You can retrieve the file in your browser's download folder. |
| Save a full screen snapshot | Downloads the current graph image as a full-page PNG file, which you can display or save.<br>You can retrieve the file in your browser's download folder. |
| Units | You can display the data with dots or as a percentage. |
| Thresholds | You can choose to show/hide **Critical**, **Immediate**, and **Warning** thresholds in the current chart. |

| Option | Description |
| --- | --- |
| Scales | You can choose a scale for a stacked chart.<br>■ Select **Linear** to view a chart in which the Y axis scale increases in a linear manner. For example, the Y axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on.<br>■ Select **Logarithmic** to view a chart in which the Y axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values.<br>**Note** If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph.<br>■ Select **Combined** to view overlapping graphs for the metrics. The chart uses individual scales for each graph instead of using a relative scale, and displays a combined view of the graphs.<br>■ Select **Combined by Unit** to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs. |
| Move Down | Moves the chart down one position. |
| Move Up | Moves the chart up one position. |

You can take the following actions on the Metric Chart graph.

| Option | Description |
| --- | --- |
| Y Axis | Shows or hides the Y-axis scale. |
| Chart | Shows or hides the line that connects the data points on the chart. |
| Data Point Tips | Shows or hides the data point tooltips when you hover the mouse over a data point in the chart. |
| Zoom by X | Enlarges the selected area on the X axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simultaneously. |
| Zoom by Y | Enlarges the selected area on the Y axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simultaneously. |
| Zoom by Dynamic Thresholds | Resizes the Y axis of the chart so that the highest and the lowest values on the axis are the highest and the lowest values of the dynamic threshold calculated for this metric. |
| Vertical resize | Resizes the height of a graph in the chart. |
| **Remove** icon next to each metric name in a stacked chart | Removes the graph for the metric from the chart. |

## Metric Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |

| Option | Description |
| --- | --- |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |

| Option | Description |
| --- | --- |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

## Metric Picker Widget

The Metric Picker widget displays a list of available metrics for a selected object.

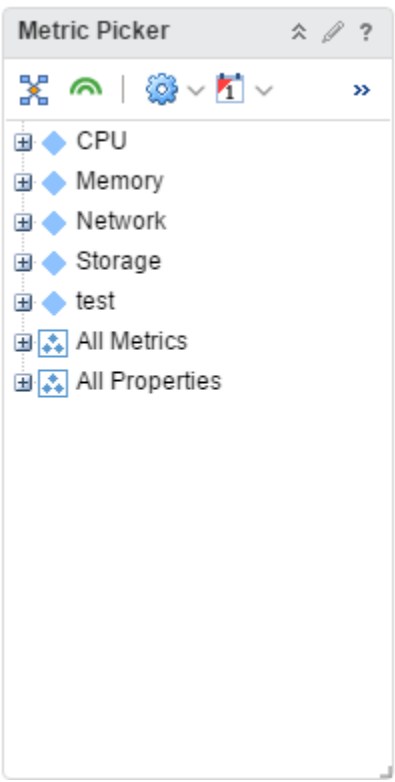

## How the Metric Picker Widget and Configuration Options Work

With the Metric Picker widget, you can check the list of the object's metrics. To select an object to pick its metrics, you use another widget as a source of data, for example, Topology Graph widget. To set a source widget that is on the same dashboard, you use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a dashboard that contains the source widget. You can also search for objects using tags.

You edit a Metric Picker widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

## Where You Find the Metric Picker Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Metric Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
| --- | --- |
| Show common metrics | Filter based on common metrics. |
| Show collecting metrics | Filter based on collecting metrics. |
| Metrics or Properties | Filter based on metrics or property metrics. |
| Time Range | Filter based on selected time range. |
| Search | Search for dashboards, views, and network IP addresses using tags. |

## Metric Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Action |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

# Object List Widget

The Object List widget displays a list of the objects available in the environment.

## How the Object List Widget and Configuration Options Work

The Object List widget displays a data grid with objects in the inventory. The default configuration of the data grid appears in Object List Widget Options section. You can customize it by adding or removing default columns. You can use the **Additional Column** option to add metrics when you configure the widget.

You edit an Object List widget after you add it to a dashboard. Configuration of the widget helps you observe parent and child objects. You can configure the widget to display the child objects of an object selected from another widget, for example, another Object List or Object Relationship widget, in the same dashboard.

Click the legend at the bottom of the widget to filter the objects based on threshold. Point your cursor over any of the boxes to view tooltips.

## Where You Find the Object List Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Object List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Action | Selects from a set of actions specific for each object type. To see available actions, select an object from the list of objects and click the toolbar icon to select an action. For example, when you select a datastore object in the graph, you can select **Delete Unused Snapshots for Datastore**. |
| Dashboard Navigation | Navigates you to the object. For example, when you select a datastore from the list of objects and click **Dashboard Navigation**, you can open the datastore in vSphere Web Client. |
| Reset Grid Sort | Returns the list of resources to its original order. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.<br>Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Display Filtering Criteria | Displays the object information on which this widget is based. |
| Page Size | |
| Filter | Locate data in the widget.<br>You can search for objects or filter the list based on the values of the metrics or properties in the additional columns of the Configuration section. |

## Object List Widget Data Grid Options

The data grid provides a list of inventory objects on which you can sort and search.

| Option | Description |
| --- | --- |
| ID | Unique ID for each object in the inventory, randomly generated and produced by VMware Aria Operations. |
| Name | Name of the object in the inventory. |
| Description | Displays the short description of the object given during creation of the object |
| Adapter Type | Shows the adapter type for each object. |

| Option | Description |
| --- | --- |
| Object Type | Displays the type of the object in the inventory. |
| Policy | Displays policies that are applied to the object. To see policy details and create policy configurations, in the menu click **Administration**, and then in the left pane click **Policies**. |
| Creation Time | Displays the date, time, and time zone of the creation of an object that was created in the inventory. |
| Identifier 1 | Can contain the custom name of the object in the inventory or default unique identifier, depending on the type of inventory object. For example, My_VM_1 for a VM in the inventory, or 64-bit hexadecimal value for VMware Aria Operations Node. |
| Identifier 2 | Can contain the abbreviation of an object type and the unique decimal number or parent instance, depending on the type of the object. For example, vm-457 for a VM and an IP address for VMware Aria Operations Node. |
| Identifier 3 | Can contain a unique number identifying an adapter type. For example, 64-bit hexadecimal value for vCenter Adapter |
| Identifier 4 | Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses. |
| Identifier 5 | Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses. |
| Object Flag | Displays a badge icon for each object. You can see the status when you point to the badge. |
| Collection State | Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the state icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**. |
| Collection Status | Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the status icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**. |
| Relevance | Displays the user interest on objects based on the number of clicks. The relevance is determined using a system-wide ranking algorithm that rates the object with most clicks as most relevant object. |
| Internal ID | Unique number that VMware Aria Operations uses to identify the object internally. For example, the internal ID appears in log files used for troubleshooting. |

## Object List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Select First Row | Determines whether to start with the first row of data. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |

| Option | Description |
| --- | --- |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |
| **Additional Columns** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.<br>1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**. |

# Object Relationship Widget

The Object Relationship widget displays the hierarchy tree for the selected object. You can create one or more hierarchy trees in VMware Aria Operations for the selected objects that you add to your custom dashboards.

## How the Object Relationship Widget and Configuration Options Work

You can add the Object Relationship widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

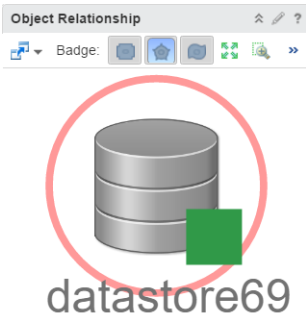


You edit an Object Relationship widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

## Where You Find the Object Relationship Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize** > **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize** > **Dashboards**. To create your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions** > **Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Object Relationship Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To be able to navigate to another dashboard, configure the relevant option when you create or edit the dashboard. |
| Badge | Displays the Health, Risk, or Efficiency alerts on the objects in the relationship map. You can select a badge for objects that appear in the widget. The tool tip of a badge shows the object name, object type, and the name of the selected badge with the value of the badge. You can only select one badge at a time. |
| Zoom to fit | Resets the chart to fit in the available space. |
| Pan | Click this icon and click and drag the hierarchy to show different parts of the hierarchy. |
| Show values on point | Shows or hides the data point tooltips when you hover the mouse over a data point in the chart. |
| Zoom the view | Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |
| Zoom in | Zooms in on the hierarchy. |
| Zoom out | Zooms out on the hierarchy. |
| Reset to Initial Object | If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Show Alerts | Select the resource in the hierarchy and click this icon to show alerts for the resource. Alerts appear in a pop-up window. You can double-click an alert to view its Alert Summary page. |

## Object Relationship Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Zoom to Fixed Node Size | You can configure a fixed zoom level for object icons in the widget display.<br>If your widget display contains many objects and you always need to use manual zooming, this feature is useful because you can use it to set the zoom level only once. |
| Node Size | You can set the fixed zoom level at which the object icons display. Enter the size of the icon in pixels.<br>The widget shows object icons at the pixel size that you configure. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |

| Option | Description |
| --- | --- |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Object Relationship (Advanced) Widget

The Object Relationship (Advanced) widget displays a graph or tree view that depicts the parent-child relationship of the selected object. It provides advanced configuration options. You can create a graph or tree view in VMware Aria Operations for the selected objects that you add to your custom dashboards.

## How the Object Relationship (Advanced) Widget and Configuration Options Work

You can add the **Object Relationship (Advanced)** widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an **Object Relationship (Advanced)** widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

You can double-click any object in the graph or tree view and see the specific parent-child objects for the focus object. When you double-click the object again, you see the original graph or tree view. If you point your cursor over an object icon, you see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object.

## Where You Find the Object Relationship (Advanced) Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Object Relationship (Advanced) Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Options | Description |
| --- | --- |
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To navigate to another dashboard, configure the relevant option when you create or edit the dashboard. |
| Reset to Initial Object | If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |
| View Tree/View graph | Displays a tree or graph view of the relationships. |
| Vertical/Horizontal | Displays a vertical or horizontal view of the graph or tree view. |
| Hide Text/Show Text | Hides or displays the object names. |
| Standard View/Fit View | The **Standard View** option fixes the view to a specific zoom level<br>The **Fit View** option adjusts the graph or tree view to fit the screen. |
| Group Items/Ungroup Items | Groups by objects types. You can view further details by double-clicking on the object. You can also choose to display the graph or tree view without grouping the object types. |
| Path Exploration | Displays the relative relationship path between two selected objects on the graph or tree view. To highlight the path, click the **Path Exploration** icon and then select the two objects from the graph or tree view. |

| Options | Description |
| --- | --- |
| Layers | ■ **Parent/Child**: Displays a graph or tree view of the parent and child relationship for the specific object selected.<br>■ **Custom**: Indicates the relationship between the objects that are part of the custom relationship. These objects have a connection via the selected custom relationship. |
| Quick Filter | Enter the name of an object that you want to see in the graph or tree view. |

## Object Relationship (Advanced) Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Name | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Parents Depth | Select the depth of parent objects to be displayed. |
| Children Depth | Select the depth of child objects to be displayed. |

| Option | Description |
| --- | --- |
| Inventory trees | Select an existing predefined traversal spec for the initial object relationship graph or tree view. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

## Property List Widget

You can use the Property List widget to view the properties of objects and their values.

## How the Property List Widget and Configuration Options Work

To observe the properties of objects in the Property List widget, you can select object property metrics when you configure the widget itself (Self Provider mode activated). Alternatively, you can select objects or object property metrics from another widget (Self Provider mode deactivated). You can also view a default or custom set of properties by selecting a preconfigured XML file in the Metric Configuration drop-down menu of the widget configuration window.

You edit a Property List widget after you add it to a dashboard. You can configure a widget to receive data from another widget by selecting **Off** for Self Provider mode. When the widget is not in Self Provider mode, it displays a set of predefined properties and their values of an object that you select on the source widget. For example, you can select a host on a Topology widget and observe its properties in the Property List widget. To configure the Property List as a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To configure a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

## Where You Find the Property List Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Property List Widget Data Grid Options

The data grid provides information on which you can sort and search.

| Option | Description |
| --- | --- |
| Object Name | Name of the object, whose properties you observe. You can sort the properties by object name. To open the Object Details page, click an object name. |
| Property Name | Name of the property. You can sort the properties by property name. |
| Value | Value of the property. You can sort the properties by value. |

## Property List Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Visual Theme | Select a predefined visual style for each instance of the widget. The options are: Original and Compact. |
| Show Metric Full Name | You can choose to view the full name of the metrics. The options are: On and Off. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |

| Option | Description |
| --- | --- |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | 1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br><br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2 Optionally, you can define measurement units for the metrics and properties in the list. Double-click a metric or properties box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.<br>3 You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**. |

| Option | Description |
| --- | --- |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Recommended Actions Widget

The Recommended Actions widget displays recommendations to solve problems in your vCenter Server instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines.

## How the Recommended Actions Widget and Configuration Options Work

The Recommended Actions widget appears on the Home dashboard, and displays the health status for the objects in your vCenter Server instance. At a glance, you can see how many objects are in a critical state, and how many objects need immediate attention.

From the Recommended Actions widget, you can focus in on problems further by, for example, clicking an object where the alerts triggered, and by clicking an individual alert.

You can edit the Recommended Actions widget on the Home dashboard, or on another dashboard where you add the widget. With the widget configuration options, you can assign a new name to the widget, set the refresh content, and set the refresh interval.

The Recommended Actions widget includes a selection bar, a summary pane, a toolbar for the data grid, and alert information for your objects in a data grid.

## Where You Find the Recommended Actions Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Recommended Actions Widget Selection Bar and Summary Pane

| Option | Description |
|---|---|
| Scope | Allows you to select an instance of vCenter Server, and a data center in that instance. |
| Object tabs | Displays the object types with the number of objects affected in parentheses. You can display the actions for virtual machines, host systems, clusters, vCenter Server instances, and datastores. |
| Badge | Select the Health, Risk, or Efficiency badge to display alerts on your objects. Health alerts require immediate attention. Risk alerts require attention in the immediate future. Efficiency alerts require your input to reclaim wasted space or to improve the performance of your objects. For each badge, you can view critical, immediate, and warning alerts.<br>■ Health Status. With the Health badge selected, displays the number of affected objects and a summary of their health based on the alerts that triggered on the object. Lists the objects that have the worst health, and the number of alerts that triggered on each object.<br>■ Risk Status. With the Risk badge selected, displays the number of affected objects and a summary of their risk based on the alerts that triggered on the object. Lists the objects that have the highest, and the number of alerts that triggered on each object.<br>■ Efficiency Status. With the Efficiency badge selected, displays the number of affected objects. Lists the objects that have the lowest efficiency based on the alerts that triggered on the object, and the number of alerts that triggered on each object. |
| Search filter | Narrows the scope of the objects that appear. Enter a character or a number to search and display an object. When a filter is active, the name of the filter appears below the Search filter text box. |

## Recommended Actions Widget Toolbar Options

The toolbar allows you to address an alert, and to filter the alert list.

| Option | Description |
| --- | --- |
| Cancel Alert | Cancels the selected alert.<br>You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |
| Suspend | Suspends an alert for a specified number of minutes.<br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.<br>The user who suspends the alert becomes the assigned owner. |
| Quick Filter | Narrows the search to one of the available filter types. For example, you can display all alerts that are related to the Compliance Alert Subtype. |

## Recommended Actions Widget Data Grid Options

The data grid displays the alerts that triggered on your objects. To resolve the problems indicated by the alerts, you can link to the alerts and the objects on which the alerts triggered.

For more information, see Accessing Alerts VMware Aria Operations.

| Option | Description |
| --- | --- |
| Criticality | Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon.<br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**. |
| Actionable | When an alert has an associated action, you can run the action on the object to resolve the alert. |
| Suggested Fix | Describes the recommendation to resolve the problem. For example, for Compliance alerts, the recommendation instructs you to use the *vSphere Hardening Guide* to resolve the problem.<br>You can find the *vSphere Hardening Guides* at http://www.vmware.com/security/hardening-guides.html.<br>You can view other available recommendations and their associated actions, if any, to resolve the problem when you click the drop-down menu. |
| Name | Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name.<br>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object. |
| Alert | Name of the alert definition that generated the alert.<br>Click the alert name to view the alert details tabs where you can begin troubleshooting the alert. |
| Alert Type | Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/Hypervisor, Hardware, Storage, and Network. |

| Option | Description |
| --- | --- |
| Alert Subtype | Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration. |
| Time | Date and time that the alert triggered. |
| Alert ID | Unique identification for the alert. This column is hidden by default. |

### Recommended Actions Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

## Risk Widget

The risk widget is the status of the risk-related alerts for the objects it is configured to monitor. Risk alerts in VMware Aria Operations usually indicate that you should investigate problems in the near future. You can create one or more risk widgets for objects that you add to your custom dashboards.

### How the Risk Widget and Configuration Options Work

You can add the risk widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the Badge Mode configuration option is set to Off, the badge and a chart appear. The type of chart depends on the object type that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning risk alerts generated over time, if the monitored object is a group.
- A trend line displays the risk status of the monitored object for all other object types.

If the Badge Mode is set to On, only the badge appears.

You edit a risk widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

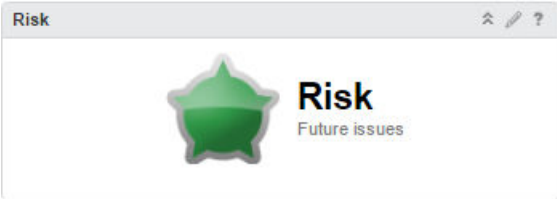


## Where You Find the Risk Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Risk Widget Display Options

The Risk Widget displays a risk badge. The widget also displays a risk trend chart when not in badge mode.

| Option | Description |
| --- | --- |
| Risk Badge | Status of the objects configured for this instance of the widget.<br>Click the badge to open the **Alerts** tab for the object that provides data to the widget. |
| Risk Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears. |

## Risk Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

| Option | Description |
| --- | --- |
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br>Select one of the following options:<br>■ On. Only the badge appears in the widget.<br>■ Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Rolling View Chart Widget

The Rolling View Chart widget cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget.

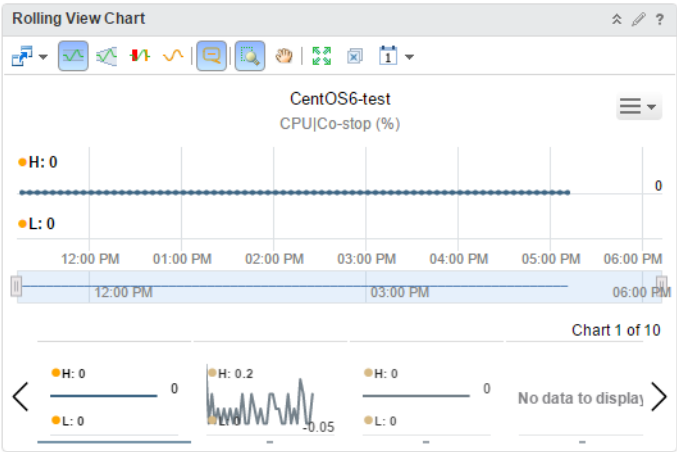


### How the Rolling View Chart Widget and Configuration Options Work

The Rolling View Chart widget shows a full chart for one selected metric at a time. Miniature graphs for the other selected metrics appear at the bottom of the widget. You can click a miniature graph to see the full graph for that metric, or set the widget to rotate through all selected metrics at an interval that you define. The key in the graph indicates the maximum and minimum points on the line chart.

You edit a Rolling View Chart widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

## Where You Find the Rolling View Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Rolling View Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
| --- | --- |
| **Trend Line** | Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points. |
| **Dynamic Thresholds** | Shows or hides the calculated dynamic threshold values for a 24-hour period. |
| **Show Entire Period Dynamic Thresholds** | Shows or hides dynamic thresholds for the entire time period of the graph. |
| **Anomalies** | Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below. |
| **Zoom to Fit** | Changes all graphs to show the entire time period and value range. |
| **Zoom the view** | Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section. |
| **Pan** | Click this icon and click and drag the hierarchy to show different parts of the hierarchy. |
| **Show Data Values** | After you click the **Show data point tips** icon to retrieve the data, click this icon and point to a graphed data point to show its time and exact value. In non-split mode, you can hover over a metric in the legend to show the full metric name, the names of the adapter instances (if any) that provide data for the resource to which the metric belongs, the current value, and the normal range. If the metric is currently alarming, the text color in the legend changes to yellow or red, depending on your color scheme. Click a metric in the legend to highlight the metric in the display. Clicking the metric again toggles its highlighted state. |
| **Date Controls** | Use the date selector to limit the data that appears in each chart to the time period you are examining.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br>**Dashboard Time** is the default option. |

## Rolling View Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Transition Interval | Time interval for a switch between charts in the widget. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |

| Option | Description |
| --- | --- |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |

| Option | Description |
| --- | --- |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br>1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2 Optionally, you can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

## Scoreboard Widget

The Scoreboard widget shows the current value for each metric of objects that you select.

## How the Scoreboard Widget and Configuration Options Work

Each metric appears in a separate box. The value of the metric determines the color of the box. You define the ranges for each color when you edit the widget. You can customize the widget to use a sparkline chart to show the trend of changes of each metric. If you point to a box, the widget shows the source object and metric data. Icons in the box indicate the level of criticality.

You edit a Scoreboard widget after you add it to a dashboard. The widget can display metrics of the objects selected during editing of the widget or selected on another widget. When the Scoreboard widget is not in Self Provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. It shows 10 predefined metrics if you do not select an XML file or if the type of the selected object is not defined in the XML file.

For example, you can configure the Scoreboard widget to use the sample Scoreboard metric configuration and to receive objects from the Topology Graph widget. When you select a host on a Topology Graph widget, the Scoreboard widget shows the workload, memory, and CPU usage of the host.

To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

## Where You Find the Scoreboard Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Scoreboard Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.<br>When the Scoreboard widget is not in self-provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. |
| Round Decimals | Select the number of decimal places to round the scores that the widget displays. |
| Box Columns | Select the number of columns that appear in the widget. |
| Layout Mode | Select a Fixed Size or Fixed View layout. |
| Fixed Size<br>Fixed View | Use these options to customize the size of the box for each object. |
| Old metric values | Select **Show** if you want the widget to show the previous value of the metric, if the current value is not available. Select **Hide** to hide the previous value of the metric, if the current value is not available. |
| Visual Theme | Select a predefined visual style for each instance of the widget. |
| Max Scores Count | Use these menus to customize the format of the scores that the widget displays. |

| Option | Description |
| --- | --- |
| Show | Select one or more of the following items to display in the widget:<br>■ Select **Object Name** to display the name of the object in the widget.<br>■ Select **Metric Name** to display the name of the metric in the widget.<br>■ Select **Metric Unit** to display the metric unit in the widget.<br>■ Select **Sparkline** to display the Sparkline chart for each metric. |
| Period Length | Select a length of time for the statistic information that the sparkline chart displays. |
| Show DT | Select an option to show or hide the dynamic threshold for the sparkline chart. |
| | |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when |

| Option | Description |
| --- | --- |
| | you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br><br>You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links open in a new tab. Examples of external links are URLs whose hostname does not match with the current VMware Aria Operations instance hostname. Internal links are URLs whose hostname matches the current VMware Aria Operations instance hostname or starts with *index.action*.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |

| Option | Description |
| --- | --- |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br>1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**. |

| Option | Description |
| --- | --- |
| | You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links will open in a new tab. Examples of external links are URLs whose hostname does not match with the current VMware Aria Operations instance hostname. Internal links are URLs whose hostname matches the current VMware Aria Operations instance hostname or starts with *index.action*.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

Click the **Near Real-Time Monitoring** icon to display near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days.

## Scoreboard Health Widget

The Scoreboard Health widget displays color-coded health, risk, efficiency, and custom metrics scores for objects that you select.

## How the Scoreboard Health Widget and Configuration Options Work

The icons for each object are color coded to give a quick indication of the state of the object. You can configure the widget to display the scores of common or specific metrics of the object. You can use the symptom state color code or you can define your criteria to color the images. If you configure the widget to show the metric for objects that do not have this metric, those objects have blue icons.

You can double-click an object icon to show the Object Detail page for the object. When you point to the icon, a tool tip shows the name of the object and the name of the metric.

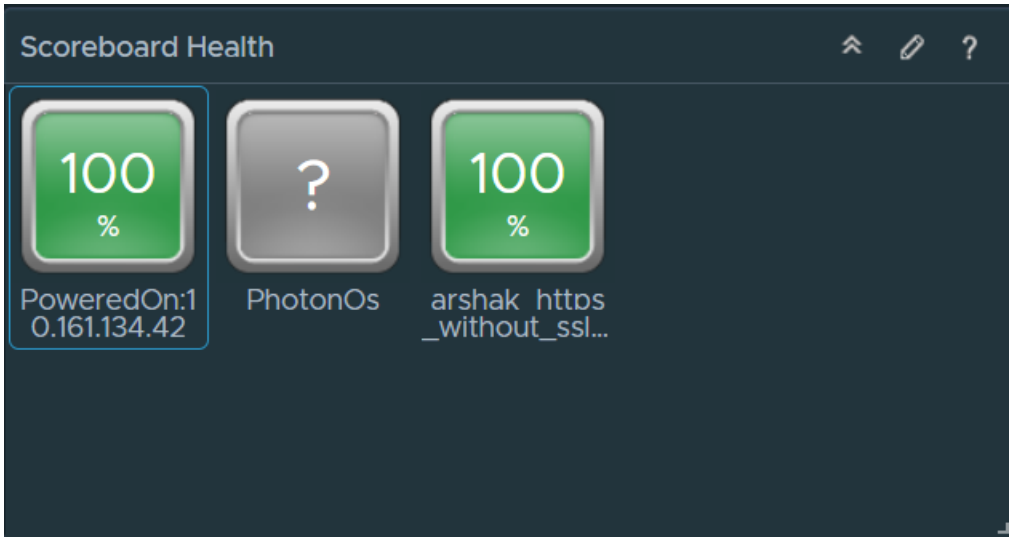


You edit a Scoreboard Health widget after you add it to a dashboard. To configure the widget, click the pencil at the upper-right corner of the widget window. The widget can display metrics of the objects that you select when you edit the widget, or that you select on another widget. For example, you can configure the widget to show the CPU workload of an object that you select on the Topology Graph widget. To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

## Where You Find the Scoreboard Health Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Scoreboard Health Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Image Type | Select an image type for the metrics. |
| Metric | Select the default or custom metric. |
| Pick Metric | Active only when you select **Custom** from the **Metric** menu.<br>Use to select a custom metric for the objects that the widget displays. Click **Pick Metric** and select an object type from the Object Type pane.<br>Use the Metric Picker pane to select a metric from the metric tree and click **Select Object** to check the objects from the type that you select on the Object Types pane. |
| Use Symptom state to color chart | Select to use the default criteria to color the image. |
| Custom ranges | Use to define custom criteria to color the image. You can define a range for each color. |

| Option | Description |
| --- | --- |
| **Input Data** | |
| | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |

## Sparkline Chart Widget

The Sparkline Chart widget displays graphs that contain metrics for an object in VMware Aria Operations. You can use VMware Aria Operations to create one or more graphs that contain metrics for objects that you add to your custom dashboards.

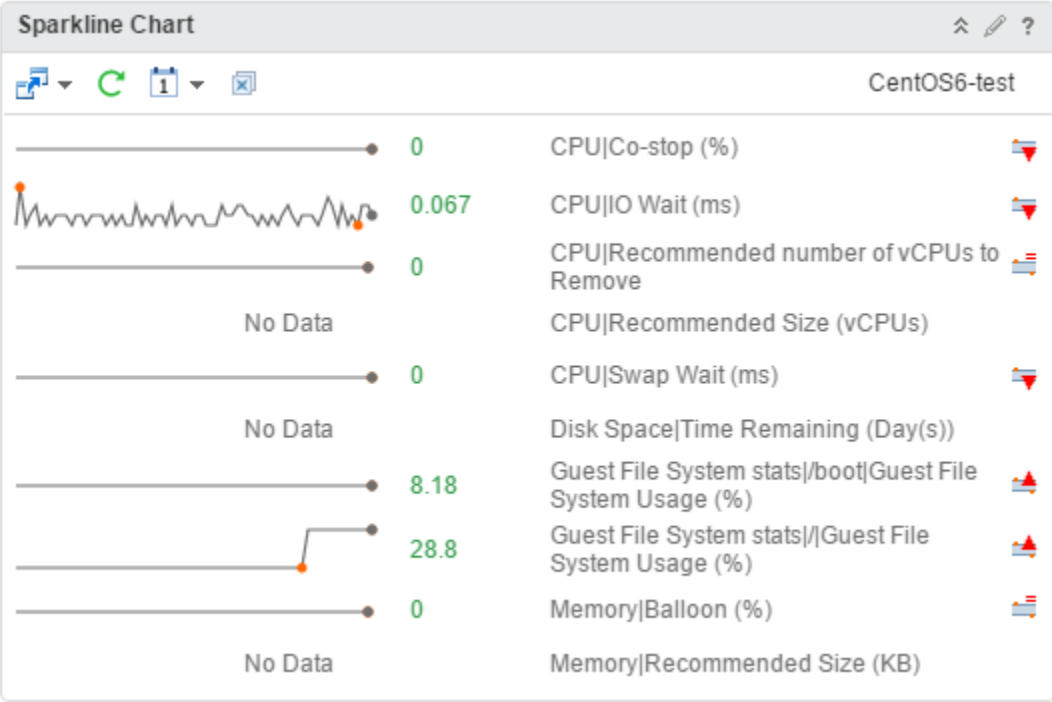


## How the Sparkline Chart Widget and Configurations Options Work

If the metrics in the Sparkline Chart are for an object that another widget provides, the object name appears at the top right of the widget. If you select a metric when you edit the widget configuration, the widget uses the metric and its corresponding object as the source for dashboard interactions. The line in the graphs represents the average value of the selected metric for the specified time period. The boxed area in the graph represents the dynamic threshold of the metric.

Point to a graph in the Sparkline Chart widget to view the value of a metric in the form of a tool tip. You can also view the maximum and minimum values on a graph. The values are displayed as orange dots.

You can add the Sparkline Chart widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The metrics shown in sparkline widget is the current value, to view the average value you can use transformations in list views or distribution charts to calculate an average. Another way to get to an average value is to double click on the sparkline to open the metric chart, click and drag to select a range, keep the mouse button depressed and hover for a few seconds, you should see a popup that has average value.

## Where You Find the Sparkline Chart Widget

The widget might be included on any of your custom dashboards. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

## Sparkline Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
| --- | --- |
| Dashboard Navigation | You can navigate to another dashboard when the object you select is also available in the dashboard to which you want to navigate. |
| Refresh | Refreshes the widget data. |
| Time Range | Select the range for the time period to show on the graphs. You can select a period from the default time range list or select start and end dates and times.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br>**Dashboard Time** is the default option. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Remove All | Removes all graphs. |

## Sparkline Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

| Option | Description |
| --- | --- |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Show Object Name | You can view the name of the object before the metric name in the Sparkline Chart widget.<br>■ **On**. Displays the name of the object before the metric name in the widget.<br>■ **Off**. Does not display the name of the object in the widget. |
| Column Sequence | Select the order in which to display the information.<br>■ **Graph First**. The metric graph appears in the first column in the widget display.<br>■ **Label First**. The metric label appears in the first column in the widget display. |
| Show DT | Select an option to show or hide the dynamic threshold for the sparkline chart. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1 Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br>Click the **Select All** icon to select all the metrics in the list.<br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you |

| Option | Description |
| --- | --- |
| | observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |

| Option | Description |
| --- | --- |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br>Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1 Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2 Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |

| Option | Description |
| --- | --- |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

# Tag Picker Widget

The Tag Picker widget lists all available object tags.

## How the Tag Picker Widget and Configuration Options Work

With the Tag Picker widget, you can check the list of the object tags. You can use the widget to filter the information that another widget shows. You can select one or more tags from the object tree or search for tags, and the destination widget displays information about the objects with this tag. For example, you can select **Object Types > Virtual Machine** on the Tag Picker widget to observe statistic information about the VMs on the Environment Status widget.

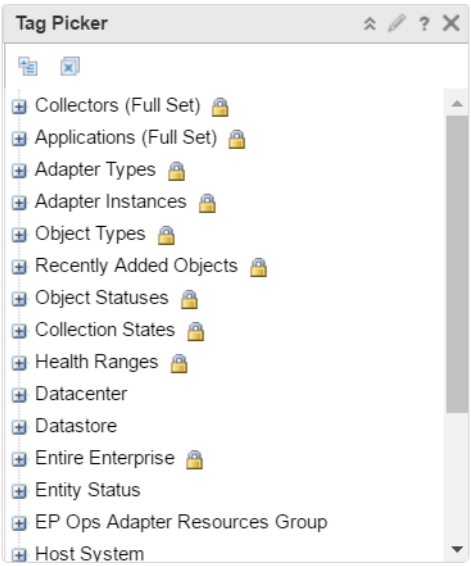

You edit a Tag Picker widget after you add it to a dashboard. To configure the widget, click the pencil in the upper right of the widget window. You can configure the Tag Picker widget to send information to another widget on the same dashboard or on another dashboard. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard. You can configure two Tag Picker widgets to interact when they are on different dashboards.

## Where You Find the Tag Picker Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Tag Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Collapse All | Close all expanded tags and tag values. |
| Deselect All | Remove all filtering and view all objects in the widget. |
| Tag Picker | Select an object from your environment. |
| Dashboard Navigation | **Note** Appears on the source widget and when the destination widget is on another dashboard.<br>Use to explore the information on another dashboard. |

## Tag Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

## Text Display Widget

You can use the Text Display widget to show text in the user interface. The text appears in the Text Display widget on the dashboard.

The Text Display widget can read text from a Web page or text file. You specify the URL of the Web page or the name of the text file when you configure the Text widget. To use the Text Display widget to read text files you must set a property in the *web.properties* file to specify the root folder that contains the file.

You can enter content in the Text Display widget in plain text or rich text format based on the view mode that you configure. Configure the Text Display widget in HTML view mode to display content in rich text format. Configure the Text Display widget in Text mode to display content in plain text format.

The Text Display widget can display websites that use the HTTPS protocol. The behavior of the Text Display widget with websites that use HTTP, depends on the individual settings of the websites.

**Note** If the webpage that you are linking to has **`X-Frame-Options`** set to **`sameorigin`**, which denies rendering a page in an iframe, the Text Display widget cannot display the contents of the webpage.

### How the Text Display Widget Configuration Options Work

You can configure the widget in the Text view mode or HTML view mode. In the HTML view mode, you can click **Edit** in the widget and use the rich text editor to add content.

If you configure the widget to use Text view mode, you can specify the path to the directory that contains the files to read or you can provide a URL. The content in the URL will be shown as text. If you do not specify a URL or text file, you can add content in the widget. Double-click the widget and enter content in plain text.

You can also use command-line interface (CLI) commands to add file content to the Text Display widget.

- To view a list of parameters, run the `file -h|import|export|delete|list txtwidget` command.
- To import text or HTML content, run the `import txtwidget input-file [--title title] [--force]` command.
- To export the content to the file, run the `export txtwidget all|title[{,title}] [output-dir]` command.
- To delete imported content, run the `delete txtwidget all|title[{,title}]` command.
- To view the titles of the content, run the `list txtwidget` command.

## Where You Find the Text Display Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Text Display Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

| Option | Description |
| --- | --- |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| View mode | Display text in text or rich text format. You can configure the widget in HTML view mode only when the **URL** and **File** fields are blank. |
| URL | Enter the URL. |
| File | Navigate to the file that contains the source text file by clicking the **Select** button.<br>To add, edit, and remove source text files, go to the **TxtWidgetContent** node in the Configuration Files page. From the left menu, click **Configure > Configuration Files**, from the VMware Aria Operations user interface. |
| Test | Validates the correctness of the text file or URL that you enter. |

## Time Remaining Widget

The Time Remaining widget displays how much time remains before the resources of the object are exhausted.

VMware Aria Operations calculates the percentage by object type based on historical data for the pattern of use for the object type. You can use the time remaining percentage to plan provisioning of physical or virtual resources for the object or rebalance the workload in your virtual infrastructure.

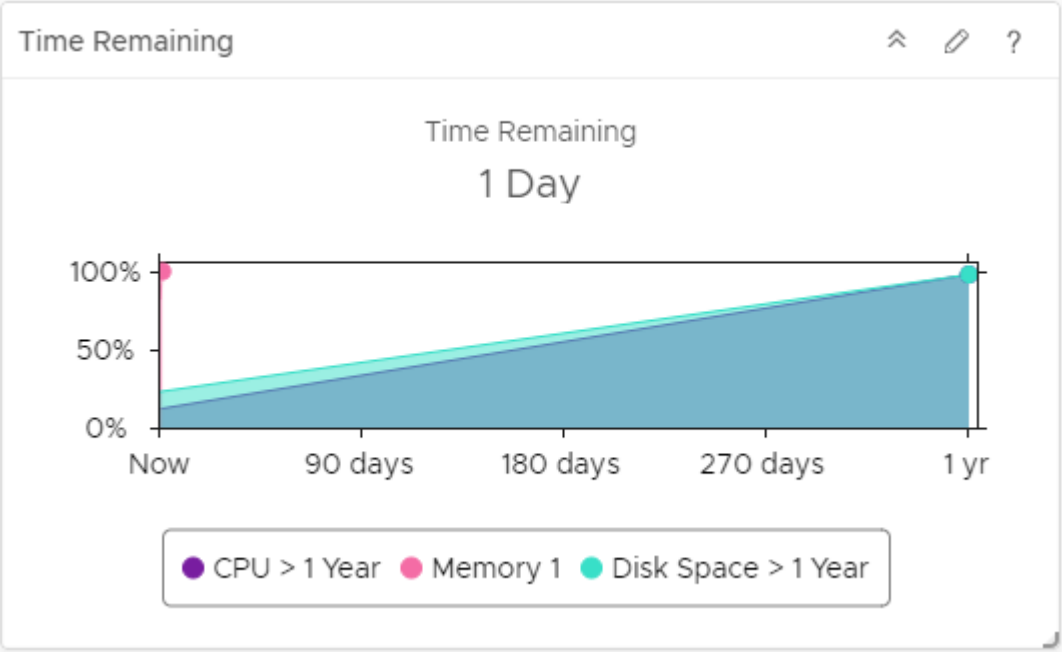


## Where You Find the Time Remaining Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Time Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Top Alerts Widget

Top alerts are the alerts with the greatest significance on the objects it is configured to monitor in VMware Aria Operations. These are the alerts most likely to negatively affect your environment and you should evaluate and address them.

## How the Top Alerts Widget and Configuration Options Work

You can add the top alerts widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a top alerts widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

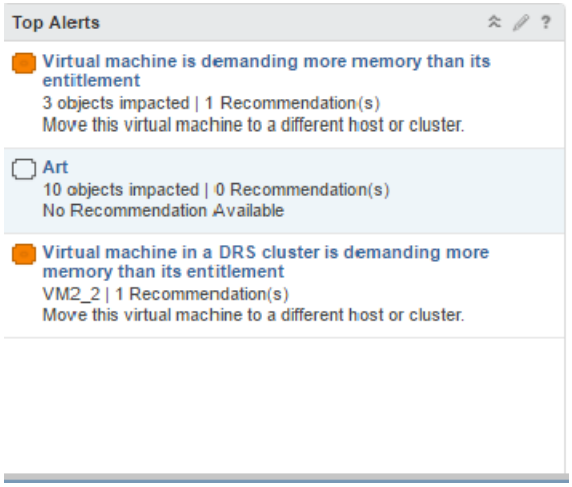


## Where You Find the Top Alerts Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize** > **Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize** > **Dashboards**. To create your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize** > **Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions** > **Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Top Alerts Widget Display Options

The Top Alerts widget includes the short description of alerts configured for the widget. The alert name opens a secondary window from which you can link to the alert details. In the alert details, you can begin resolving the alerts.

| Option | Description |
| --- | --- |
| Alert name | Name of the generated alert. Click the name to open the alert details. |
| Alert description | Number of affected objects, and the number of recommendations and the best recommendation to resolve the alert. |

## Top Alerts Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Impact Badge | Select the badge for which you want alerts to appear.<br>The affected badge is configured when you configure the alert definition. |
| Number of Alerts | Select the maximum number of alerts to display in the widget. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |

## Top-N Widget

The Top-N widget displays the top n results from analysis of an object or objects that you select.

## How the Top-N Widget and Configuration Options Work

You can select an object when you configure the Top-N widget or you can select an object on another widget. The widget shows an analysis of the applications, alerts, and metrics of an object and its child objects depending on how you configure the widget. The widget can show an analysis of the current values or values over a period of time. You can receive detailed information about each object on the widget. When you double-click an object, the Object Detail page appears.

You can configure a widget to receive data from another widget by selecting **Off** for Self Provider. You can configure a widget to display results from analysis of an object that you select on the source widget.

For example, you can select a host on a Topology widget and observe the metric analysis of the virtual machines on the host. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

## Where You Find the Top-N Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Top-N Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Icon | Description |
| --- | --- |
| Dashboard Navigation | Takes you to a predefined object. For example, when you select a datastore from the data grid and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client. |
| Select Date Range | Limits the alerts that appear in the list to the selected date range.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. |
| Object details | Select an object and click this icon to show the Object Detail page for the object. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |

## Top-N Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Redraw Rate | Set the redraw rate. |
| Bars Count | Select the number of top results. |
| Round Decimals | Select the number of decimals to round the scores displayed in the widget. |

| Option | Description |
| --- | --- |
| Filter old metrics | Select or deselect whether the analysis includes old metric values. |
| Application Health and Performance | ■ Top Least Healthy. The top n results from an analysis of the object or objects that are the least healthy.<br>■ Top Most Healthy. The top n results from an analysis of the object or objects that are the most healthy.<br>■ Top Most Volatile. The sorted list of values based on the standard deviation of values for several alerts over time.<br>Select the criteria for analysis of the objects. |
| Alert Analysis | Select the criteria for analysis of the alerts. |
| Metric Analysis | If you select this option, you must select a metric in the **Output Data** section.<br>■ Top Highest Utilization. A list of objects with similar object types that have the highest utilization on configuring usage metrics like CPU usage and memory usage.<br>■ Top Lowest Utilization. A list of objects with similar object types that have the lowest utilization on configuring usage metrics like CPU usage and memory usage.<br>■ Top Abnormal States. The objects are ordered by the duration of all alarms that are triggered on the selected metric for a selected interval.<br>■ Top Highest Volatility. The sorted list of values based on the standard deviation of values for several alerts over time.<br>Select the criteria for analysis of the metric that you select from the metric tree. |
| **Input Data** | |

| Option | Description |
| --- | --- |
| Objects | Select objects on which you want to base the widget data.<br>1 Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2 Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| | Select an object type in your environment on which you want to base the widget data.<br>1 Click the **Add Object Type** icon to search for and add an object type.<br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.<br>2 Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type.<br>If the objects have an input transformation applied, the transformed objects are the basis for the widget data. |
| Metric | Select a common metric or a metric for the selected object type in the list. The metric is the basis for the widget data. |

| Option | Description |
| --- | --- |
| Label | Type in a name that displays as a label for the metric.<br>You can add a label if you have selected **Metric Analysis > Top Highest Utilization** or **Metric Analysis > Top Lowest Utilization** as Top-N options in the **Configuration** section. |
| Unit | You can define measurement units for the metrics. Select a measurement unit in the **Unit** drop-down menu.<br>You can add a unit if you have selected **Metric Analysis > Top Highest Utilization** or **Metric Analysis > Top Lowest Utilization** as Top-N options in the **Configuration** section. |
| Maximum | Specify the maximum value based on which the bar size is calculated.<br>You can add a maximum value if you have selected any of the options under **Metric Analysis**. |
| Color Method | You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. If you do not want to use color, select None.<br>You can add color thresholds if you have selected **Metric Analysis > Top Highest Utilization**, **Metric Analysis > Top Lowest Utilization**, or **Metric Analysis > Percentile** as Top-N options in the **Configuration** section. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br>If the objects have an input transformation applied, you select tag values for the transformed objects. |

| Option | Description |
| --- | --- |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |
| **Additional Columns** | |
| | Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.<br>1 Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2 Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**. |

## Topology Graph Widget

The Topology Graph widget gives a graphical presentation of objects and their relationships in the inventory. You can customize each instance of the widget in your dashboard.

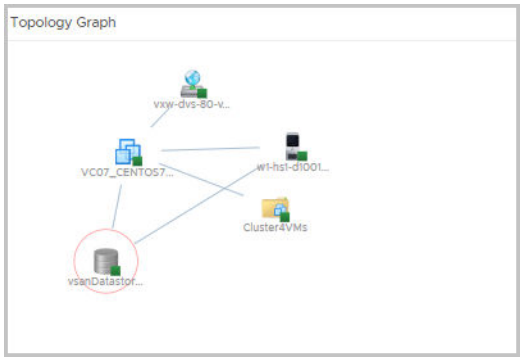


## How the Topology Graph Widget and Configuration Options Work

The Topology Graph widget helps you explore all nodes and paths connected to an object from your inventory. Connection between the objects might be a logical, physical, or network connection. The widget can display a graph that shows all of the nodes in the path between two objects, or that shows the objects related to a node in your inventory. You select the type of graph in the Exploration Mode when you configure the widget. You can select the levels of exploration between nodes in the displayed graph by using **Relationship** check boxes when you edit the widget. The widget displays all object types in the inventory by default, but you can select object types to view by using the Object View list during the configuration process. Double-clicking an object on the graph takes you to a detailed page about the object.

## Where You Find the Topology Graph Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Topology Graph Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
| --- | --- |
| Action | Use to select from predefined actions for each object type. To see available predefined actions, select an object in the graph and click the toolbar to select an action. For example, when you select a datastore object in the graph, you can click **Delete Unused Snapshots for Datastore** to apply this action to the object. |
| Dashboard Navigation | Takes you to a predefined object. For example, when you select a datastore from the graph and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client. |
| Pan | Use to move the entire graph. |
| Show values on point | Provides a tool tip with parameters when you point to an object in the graph. |
| Zoom in | Zooms in the graph. |
| Zoom out | Zooms out the graph. |
| Hierarchical View | Use to switch to hierarchical view. Hierarchical view is activated only for Node Exploration mode and with selected inventory tree. |
| Graph View | Use to switch to graph view. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Expand Node | Selects which object types related to your object to show on the graph. For example, if you select a virtual machine from the graph and click **Expand Node** toolbar icon and select **Host System**, the host on which the virtual machine is located is added to the graph. |
| Hide Node(s) | Use to remove a given object from the graph |
| Reset To Initial Object | Use to return to the initially displayed graph and configured object types. |
| Explore Node | Use to explore a node from a selected object in the graph. For example, if the graph displays a connection between a VM, a host, and a datastore, and you want to check the connection of the host with the other objects in the inventory, you can select the host and click **Explore Node**. |
| Status | Use to select objects based on their status or their state. |

## Topology Graph Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Exploration Mode | Use **Node Exploration mode** to observe a selected object from an object list and the objects related to it. For example, if you select a virtual machine and select node exploration mode, the widget shows the host where the VM is placed and the datastore storing the files of the VM.<br>Use **Path Exploration mode** to observe the relation between two objects. You must select them from the Select First Object list and the Select Second Object list. For example, if you select to explore the path between a VM and a vCenter Server, the graph shows you both objects and all nodes in the path between the VM and server as datastore, datastore cluster, and data center.<br>**Important** To select object view is mandatory for the widget to start working in path exploration mode. |

| Option | Description |
| --- | --- |
| Show Paths | Use **All** to observe connections between a node and nodes related to it as well as connections between the nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph shows a VM connected to its datastore and host and the connection between the host and datastore. |
| | Use **Discovered Only** to observe directly related nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph will shows the VM connected to its datastore and to its host, but without the connection between the host and datastore. |
| Configuration File | The default configuration includes parent and child relationship. Drop-down options depend on the installed Solutions. You can add a new type of relationship to the Relationship pane. |
| Metric Configuration | Specifies a list with attributes to display. |
| Layout | Select whether you want a graph view or hierarchical view for the topology graph. |
| Tree type | For a hierarchical layout, select whether you want a tree type view. |
| **Input Data** | |
| Selected object | From the object list, select an object on which you want to base the widget data. |
| Degree of separation | Available only when node exploration mode is selected. Use to define the levels of exploration in node exploration mode. The lowest degree configuration shows only directly related nodes rather than higher degrees that show the inventory in details. |
| Select First Object | Available only in path exploration mode. Select the first object from the object list. |
| Select Second Object | Available only in path exploration mode. Select the second object from the object list. |
| Object view | Use to select which types of objects to observe in the graph. |
| Relationship | Select the type of relationship between objects to observe in the graph, respectively the details about your inventory . The common relationships for all objects are parent and child, but the list of relationships can vary depending on added solutions to VMware Aria Operations. |

# View Widget

The View widget provides the VMware Aria Operations view functionality into your dashboard.

## How the View Widget and Configuration Options Work

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, supermetrics, properties, alerts, policies, and data from a different perspective.

You can add the View widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. List views can send interactions to other widgets.

## Where You Find the View Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

You can export the view as a CSV file for any view type.

## View Widget Toolbar Options

The View widget toolbar depends on the displayed view type.

| Option | Description |
| --- | --- |
| Export as CSV | You can export the view as a CSV file for any view type. |
| Open in External Application | Ability to link to another application for information about the object. For example, you have a List view with VMs. You can select any VM and select **Open in External Application** to open the VM in vSphere Web Client. |

| Option | Description |
| --- | --- |
| Time Settings | Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.<br>■ Relative Date Range. Select a relative date range of data transformation.<br>■ Specific Date Range. Select a specific date range of data transformation.<br>■ Absolute Date Range. Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: Hours, Days, Weeks, Months, and Years.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>■ Dashboard Time. Select this option to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Items per page | You can set the number of results that appear in the widget. Available for **List** view only. |
| Roll up interval | The time interval at which the data is rolled up. |
| Actions | An action on the selected object. Depends on the object type. |
| Filter | Limits the list to objects for a specific host, data center, and so on. You can drill-down in the hierarchical level. Available for **List**, **Trend**, and **Distribution** types of Views. |
| Filter by name | Limits the list to objects of a specific name. Available for **List** view only. |

## View Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Inventory trees | Select an existing predefined traversal spec to pick an object for the widget data. |
| Object | In self-provider mode, click the **Add Object** icon to select an object from the object list. The object list is displayed based on the inventory tree selection. You can also search for the object in this text box. |
| **Output Data** | |
| | A list of defined views available for the selected object is displayed.<br>You can create, edit, delete, clone, export, and import views directly from the View widget configuration options.<br>For more information, see Chapter 18 Views in VMware Aria Operations. |

| Option | Description |
| --- | --- |
| Auto Select First Row | Determines whether to start with the first row of data for list type views. |
| Show | Select one or more of the following items to display in the widget:<br>■ To display the list of legends in the widget, select **Legend**.<br>■ To display the name of the labels in the widget, select **Labels**. |

# Weather Map Widget

The Weather Map widget provides a graphical display of the changing values of a single metric for multiple resources over time. The widget uses colored icons to represent each value of the metric. Each icon location represents the metric value for particular resources. The color of an icon changes to show changes in the value of the metric.

## How the Weather Map Widget and Configuration Options Work

You can add the Weather Map widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

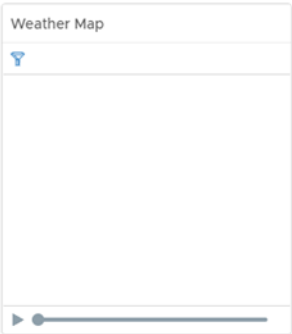


Watching how the map changes can help you understand how the performance of the metric varies over time for different resources. You can start or stop the display using the **Pause** and **Play** options at the bottom of the map. You can move the slider forwards or backwards to a specific frame in the map. If you leave the widget display and return, the slider remains in the same state.

The map does not show the real-time performance of the metrics. You select the time period, how fast the map refreshes, and the interval between readings. For example, you might have the widget play the metric values for the previous day, refreshing every half second, and have each change represent five minute's worth of metric values.

To view the object that an icon represents, click the object.

## Where You Find the Weather Map Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Weather Map Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains the icons that you can use to view the graph.

| Icon | Description |
| --- | --- |
| **Pause** and **Play** | Start or stop the display. The icon remains in the same state if you leave the widget display and return. |
| **Display Filtering Criteria** | View the current settings for the widget, including the current metric. |

## Weather Map Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

| Option | Description |
| --- | --- |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Redraw Rate | An interval at which cashed data is refreshed based on newly collected data.<br>For example, if you set metric history to **`Last 6 hours`** and image redraw rate to **`15 minutes`**, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes.<br>For example, if you set metric history to **`Last 6 hours`** and image redraw rate to **`15 minutes`**, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes. |
| Metric History | Select the time period for the weather map, from the previous hour to the last 30 days. |
| Metric Sample Increment | Select the interval between metric readings. For example, if you set this option to one minute and set the Metric History to one hour, the widget has a total of 60 readings for each metric. |
| Group by | Select a tag value by which to group the objects. |
| Sort by | Select **Object name** or **Metric value** to set the way to sort the objects. |
| Frame Transition Interval | Select how fast the icons change to show each new value. You can select the interval between frames and the number of frames per second (fps). |
| Start Over Delay | The number of seconds for the display to remain static when it reaches the end of the Metric History period, the most current readings, before it starts over again from the beginning. |
| Color | Shows the color range for high, intermediate, and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes.<br>If you leave the text boxes blank, VMware Aria Operations maps the highest and lowest values for the **Color By** metric to the end colors.<br>If you set a minimum or maximum value, any metric at or beyond that value appears in the end color. |
| **Output Data** | |

| Option | Description |
| --- | --- |
| | Select an object type in your environment on which you want to base the widget data.<br>1 Click the **Add Object Type** icon to search for and add an object type.<br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.<br>2 Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type. |
| Metric | Select a common metric or a metric for the selected object type in the list. The metric will be the basis for the widget data. The object corresponding to the metric is the selected object for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>1 In the first drop-down menu, select an object type.<br>2 In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3 In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4 To add more filter criteria, click **Add**.<br>5 To add another filter criteria set, click **Add another criteria set**. |

## Workload Widget

The Workload widget displays data indicating how hard a selected resource is working.

The Workload widget displays a graph depicting how hard the object that you selected is working. The Workload widget reports data on CPU usage, Memory usage, Disk I/O, and Network I/O.

## Where You Find the Workload Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

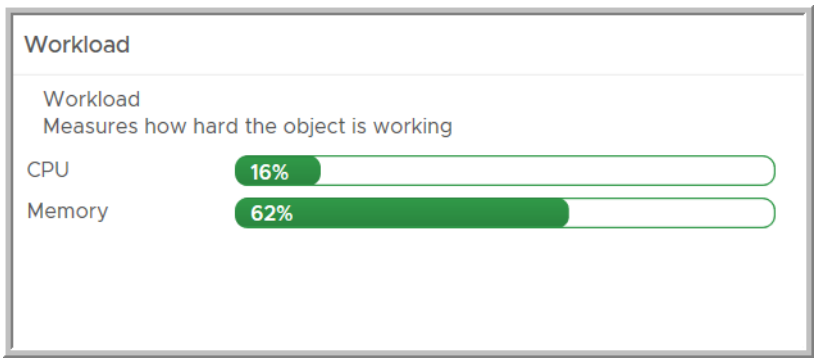


## About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.
- The metric named `devices|totalLatency_average` is not supported.

As a result, VMware Aria Operations does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

## Workload Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Workload Pattern Widget

The Workload Pattern widget displays a historical view of the hourly workload of an object.

## Where You Find the Workload Pattern Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Visualize > Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Visualize > Dashboards**. To create your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Visualize > Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions > Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

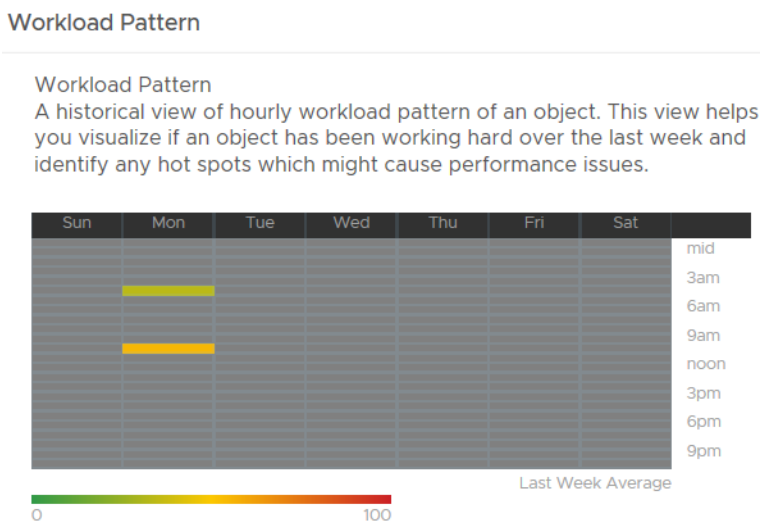


## Workload Pattern Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
| --- | --- |
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

| Option | Description |
| --- | --- |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>■ On. You define the objects for which data appears in the widget.<br>■ Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

# Reports in VMware Aria Operations

17

A report is a scheduled snapshot of views and dashboards. You can create reports in VMware Aria Operations to represent objects and metrics. The report can contain a table of contents, cover page, and footer.

With the VMware Aria Operations reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.

## Types of Report Templates

There are two types of report templates in VMware Aria Operations.

- Predefined report templates that are out of the box with VMware Aria Operations. For more information, see Accessing Report Templates .
- Custom report templates that you create based on your requirement. For more information, see Create a Report Template.

Read the following topics next:

- Create a Report Template
- Manage Report Templates

## Create a Report Template

You create a report to generate a scheduled snapshot of views and dashboards. You can track current resources and predict potential risks to the environment. You can schedule automated reports at regular intervals.

### Procedure

1 From the left menu, click **Visualize > Reports**. The **Report Templates** tab is in the right panel.

2 From the **Report Templates** tab, click **Add** to create a template.

3 Complete the steps in the left pane to:

   a Enter a name and description for the report template.

     Name and Description Details

   b Add a view or a dashboard.

     Views and Dashboards Details

   c Select an output for the report.

     Formats Details

   d Select the layout options.

     Layout Options Details

4 Click **Save**.

## Name and Description Details

The name and description of the report template as they appear in the list of templates on the **Report Templates** tab.

### Where You Add Name and Description

To create report templates, from the left menu, click **Visualize > Reports**. The **Report Templates** tab is in the right panel. From the **Report Templates** tab, click **Add**. From the **New Template** dialog box, in the workspace, on the left, click **Name and Description**.

Table 17-1. Name and Description Options in the Report Template Workspace

| Option | Description |
| --- | --- |
| Name | Name of the template as it appears on the **Report Templates** tab. |
| Description | Description of the template. |

## Views and Dashboards Details

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs.

### Where You Add Views and Dashboards

To create report templates, from the left menu, click **Visualize > Reports**. The **Report Templates** tab is in the right panel. From the **Report Templates** tab, click **Add**. From the **New Template** dialog box, in the workspace, on the left, click **Views and Dashboards**. If you create a template, complete the required previous steps of the workspace.

## How You Add Views and Dashboards

To add a view or a dashboard to your report template, select it from the list on the left pane and drag it to the main panel. You can drag the views and dashboards in the main panel to reorder them. You can select a portrait or landscape orientation for each view or dashboard from the drop-down menu next to its title.

### Table 17-2. Views and Dashboards Options in the Report Template Workspace

| Option | Description |
| --- | --- |
| Data type | Select **Views** or **Dashboards** to display a list of available views or dashboards that you can add to the template. |
| Filter | Search for views or dashboards by name. To see the complete list of views or dashboards, delete the search box contents and press Enter. |
| List of views | List of the views that you can add to the template. This list is available when you select **Views** from the **Data type** drop-down menu. |
| List of dashboards | List of the dashboards that you can add to the template. This list is available when you select **Dashboards** from the **Data type** drop-down menu. |
| Preview of views and dashboards | In the main panel, you see a preview of the views and dashboards that you add.<br>When you create a template in the context of an object from the environment, you see a live preview of the views and dashboards. |
| Colorization | You can activate or deactivate a colorized PDF output for each list view. This option is available from the right panel when you select **Views** from the **Data type** drop-down menu. |

# Formats Details

The formats are the outputs in which you can generate the report.

## Where You Add Formats

To create report templates, from the left menu, click **Visualize > Reports**. The **Report Templates** tab is in the right panel. From the **Report Templates** tab, click **Add**. From the **New Template** dialog box, in the workspace, on the left, click **Formats** to select a format for the report template. If you create a template, complete the required previous steps of the workspace.

Table 17-3. Formats Options in the Report Template Workspace

| Option | Description |
| --- | --- |
| PDF | With the PDF format, you can read the reports, either on or off line. This format provides a page-by-page view of the reports, as they appear in printed form. |
| CSV | In the CSV format, the data is in a structured table of lists. |

## Layout Options Details

The report template can contain layout options such as a cover page, table of contents, and footer.

### Where You Add Layout Options

To create report templates, from the left menu, click **Visualize > Reports**. The **Report Templates** tab is in the right panel. From the **Report Templates** tab, click **Add**. From the **New Template** dialog box, in the workspace, on the left, click **Layout Options**. If you create a template, complete the required previous steps of the template.

Table 17-4. Layout Options in the Report Template Workspace

| Option | Description |
| --- | --- |
| Cover Page | Can contain an image up to 5 MB.<br>The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page. |
| Table of contents | Provides a list of the template parts, organized in the order of their appearance in the report. |
| Footer | Includes the date when the report is created, a note that the report is created by VMware Aria Operations, and page number. |

# Manage Report Templates

Using VMware Aria Operations you can manage report templates by running, scheduling, and generating report templates. Additionally, you can create custom report templates, edit existing report templates, download report templates, and so on.

## Reports Workflow

The following flowchart describes a simple workflow for reports in VMware Aria Operations.

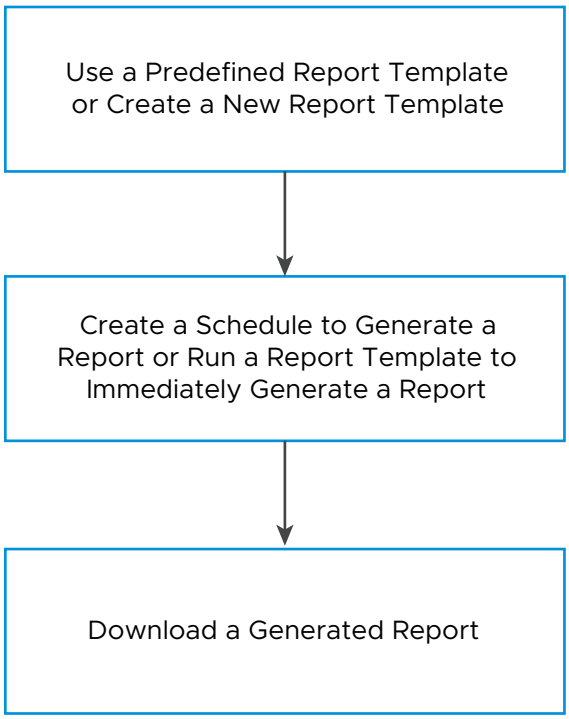


## Accessing Report Templates

A report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. VMware Aria Operations offers several predefined report templates that you can use based on your requirement.

**Where You Can Access Report Templates From**

From the left menu, click **Visualize > Reports**. The **Report Templates** page is in the right panel, or

From the left menu, click **Environment > Object Browser**. From the **Object Browser** panel, select an object and click **Reports > Report Templates** to access the **Reports Templates** tab.

The listed report templates are user-defined and predefined by VMware Aria Operations. You can order them by template name, description, subject, date they were modified, last run report, or the user who modified them. For each template, you can see the number of generated reports and schedules.

You can filter the reports based on the name of the report template, the subject, and the owner. You can click **Add** to create a report template. For information about creating a report template, see Create a Report Template.

You can select a report template from the list, click the vertical ellipsis against each report template, and select options such as run, edit, schedule, delete, clone, and export a report.

Table 17-5. Predefined Filter Groups

| Filter Group | Description |
| --- | --- |
| Name | Filter by the template name. For example, type **`my template`** to list all reports that contain the **`my template`** phrase in their name. |
| Subject | Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by the other objects. |
| Owner | Filter by the owner of the report template. |

The maximum number of reports per template is 10. After the tenth report is generated, VMware Aria Operations deletes the oldest report. vSphere users must be logged in until the report generation is complete. If you log out or your session expires, the report generation fails.

### Report Template Actions

You can select more than one report template and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

| Option | Description |
| --- | --- |
| Delete | Deletes the report template. |
| Export | Downloads the report template. |
| Import | Allows you to import a report template by selecting a report template in XML or zip file format.<br>To import a report template:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select a report template ZIP or XML file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the report template, and click **Done**. |
| Change default cover image | Allows you to change the default cover image of the report template. For more information, see Upload a Default Cover Page Image for Reports. |

## Schedule a Report

In VMware Aria Operations, the schedule of a report is the time and recurrence of report generation. To generate a report on a selected date, time, and recurrence, you create a schedule for the report template. You set the email options to send the generated report to your team.

The date range for the generated report is based on the time when the report is generated and not on the time when you schedule the report or when VMware Aria Operations places the report in the queue.

**Note** Only users created in VMware Aria Operations can add and edit report schedules.

**Note** You can alternatively schedule a report from **Visualize > Reports > Manage > Report Templates**. Continue from Step 4 below.

Prerequisites

- Download the generated report to verify the output.
- To activate sending email reports, you must have configured Outbound Alert Settings. See Notifications in VMware Aria Operations.

Procedure

1 From the left menu, click **Environment > Object Browser**.

2 From the **Object Browser** panel, navigate to the object.

3 Click the **Reports** tab and then click **Report Templates**.

4 Select the relevant report template from the list.

5 Click the vertical ellipsis and select **Schedule**.

6 Select an object and click **Next**.

This step is only applicable if you Schedule a report from **Visualize > Reports > Manage > Report Templates tab > Vertical Ellipsis > Schedule**.

7 Select the time zone, date, hour, and minutes (in the range of 0, 15, 30, and 45 minutes) to start the report generation.

VMware Aria Operations generates the scheduled reports in sequential order. Generating a report can take several hours. This process might delay the start time of a report when the previous report takes an extended period of time.

8 From the **Recurrence** drop-down menu, select one of the following options for report generation:

| Option | Description |
| --- | --- |
| **Daily** | You can set the periodicity in days. For example, you can set report generation to every two days. |
| **Weekly** | You can set the periodicity in weeks. For example, you can set report generation to every two weeks on Monday. |
| **Monthly** | You can set the periodicity in months. |

9 Select the **Email report** check box to send an email with the generated report.

Email a generated report to a predefined email group or to a network shared location. For more information about how to set up and configure the email options, see Add a Standard Email Plug-In for VMware Aria Operations Outbound Alerts.

a In the **Email addresses** text box, enter the email addresses that must receive the report. You can also add email addresses in the CC list and BCC list.

b Select an outbound rule.

An email is sent according to this schedule every time a report is generated.

#### What to do next

### Editing a Report Schedule

To edit the schedule of a report, click the link in the **Schedules** column against the report template from the **Report Templates** tab, and then from the **Scheduled Reports** dialog box, click **Edit Schedule**. You see the **Scheduled Reports** page.

Table 17-6. Scheduled Reports Toolbar Options

| Options | Description |
| --- | --- |
| New Schedule | You can create a schedule for the report. |
| Edit Schedule | You can edit an existing report schedule. |
| Delete Schedule | You can delete an existing report schedule. |
| Transport Report Schedule | You can assign a new owner for the selected report schedule. You can select a target user from the **Transfer Report Schedules** dialog box. |

**Note** You can edit, clone, and delete report templates. Before you do, familiarize yourself with the consequences of these actions. When you edit a report template and delete it, all reports generated from the original and the edited templates are deleted. When you clone a report template, the changes that you make to the clone do not affect the source template. When you delete a report template, all generated reports are also deleted.

## Generate and Regenerate a Report

To generate a report, use a predefined or custom report template.

#### Prerequisites

Create a report template.

**Note** You can alternatively generate a report from **Visualize** > **Reports** > **Manage** > **Report Templates**. Continue from Step 4 below.

### Procedure

1 From the left menu, click **Environment > Object Browser**

2 Navigate to the relevant object.

3 Click the **Reports** tab and then click **Report Templates**.

The listed report templates are associated with the current object.

4 Navigate to the relevant report template, click the vertical ellipsis, and select **Run**.

### Results

The report is generated and listed on the **Generated Reports** tab.

**Note** To regenerate the selected report, from the **Generated Reports** tab, click the vertical ellipsis against the generated report and select **Run**.

### What to do next

Download the generated report and verify the output.

## Accessing Generated Reports

You can view a list of report templates that have been generated in VMware Aria Operations.

### Where You Can Access Generated Reports From

From the left menu, click **Visualize > Reports**. From the **Reports** panel, click **Generated Reports**. The right pane contains all the generated reports, or

From the left menu, click **Environment > Object Browser**, and then from the **Object Browser** page, select an object and click **Reports > Generated Reports** to access the Generated Reports tab. All reports that are generated for a selected object are listed in the **Generated Reports** tab.

If the report is generated through a schedule, the owner is the user who created the schedule.

**Note** The maximum number of reports per template is 10. After the tenth report is generated, VMware Aria Operations deletes the oldest report.

To select a generated report from the list, click the vertical ellipsis against each generated report and select options such as run and delete. You can also select more than one generated report and select **Delete** from the **Actions** drop-down menu to delete a generated report.

You can filter the reports list by adding a filter from the upper-right corner of the panel.

Table 17-7. Predefined Filter Groups

| Filter Group | Description |
| --- | --- |
| Report Name | Filter by the report template name. For example, type **`my template`** to list all reports that contain the **`my template`** phrase in their name. |
| Template | Filter by the report template. You can select a template from a list of templates applicable for this object. |
| Completion Date/Time | Filter by the date, time, or time range. |
| Subject | Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by that second object. |
| Status | Filter by the status of the report. |

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

If you log in to VMware Aria Operations with vCenter Server credentials and generate a report, the generated report is always blank.

## Download a Report

To verify that the information appears as expected, you download the generated report.

### Prerequisites

Generate a report.

### Procedure

1 From the left menu, click **Environment > Object Browser**.

2 Navigate to the object for which you want to download a report.

3 Click the **Reports** tab and then click **Generated Reports**.

The listed reports are generated for the current object.

4 Click the PDF or the CSV icon in the Download column to download the report.

### Results

VMware Aria Operations saves the report file.

### What to do next

Schedule a report generation and set the email options, so your team receives the report.

# Upload a Default Cover Page Image for Reports

You can upload a common default image for the cover page of reports. You do not have to upload a cover page for each report. The cover pages of predefined reports are modified when you use this option. The cover pages of user-defined reports do not change.

## Where Do You Upload a Default Cover Page Image for Reports

To upload a default cover page for reports, from the left menu, click **Environment > Object Browser**, and then in the left pane navigate to an object and click the **Reports** tab. From the **Report Templates** tab, click the horizontal ellipsis next to the **Add** option and click the **Change default cover image** option.

## How Do You Upload a Default Cover Page Image for Reports

Browse for the image that you want to add to the cover page and click **Save**. You can also use the default product image that is available.

# Views in VMware Aria Operations 18

VMware Aria Operations provides several types of views. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective. Views also show information that the adapters in your environment provide.

You can configure VMware Aria Operations views to show transformation, forecast, and trend calculations.

- The transformation type determines how the values are aggregated.
- The trend option shows how the values tend to change, based on the historical, raw data. The trend calculations depend on the transformation type and roll up interval.
- The forecast option shows what the future values can be, based on the trend calculations of the historical data.

You can use VMware Aria Operations views in different areas of VMware Aria Operations.

- To manage all views, from the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage**.
- To see the data that a view provides for a specific object, navigate to that object, click the **Details** tab, and click **Views**.
- To see the data that a view provides in your dashboard, add the View widget to the dashboard. For more information, see View Widget.

Table 18-1. Options from the Views panel

| Options | Description |
| --- | --- |
| Manage | You can manage views by clicking **Visualize > Views**. From the Views panel, click **Manage**. |
| Create | Use this option to create a view. See Create and Configure a View. |
| Search | You can search for a view across the **Recents** and **All** folders in the **Views** panel. |

Table 18-1. Options from the Views panel (continued)

| Options | Description |
| --- | --- |
| Recent | The views are listed in the order in which you select them, with the most recent view that you selected, appearing at the top. Up to ten views can be displayed as Recent views.<br>If you do not pin the view and log out of the user interface, on logging back in, the view is removed from the Recents folder |
| All | Lists the views based on their type. You can use this menu for quick navigation through your views. When you navigate to a view using the **Visualize > Views** option, the views are listed in the left pane of the Views panel under **All**. You can also search for views using keywords and letters. |

# Views and Reports Ownership

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

Read the following topics next:

- Accessing Predefined Views
- Views Overview
- Views and Reports Ownership
- Create and Configure a View
- Editing, Cloning, and Deleting a View
- Share a View
- Including Deleted VMs in List View

- User Scenario: Create, Run, Export, and Import a VMware Aria Operations View for Tracking Virtual Machines

# Accessing Predefined Views

You can access some of the useful, predefined views from the **Views** home page.

To access these views, from the left menu, click **Visualize > Views**. From the **Views** panel, click **Home**.

The views are categorized as follows: Availability, Capacity, Configuration, Inventory, Performance, and Compliance. To easily access some of the useful, predefined dashboards under these categories, click on the drop-down button against the selected category and click on the specific dashboard.

# Views Overview

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective.

**How You Access the Views Page**

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage** to access the **Views** page.

## Manage and Preview Views

You can preview a view by clicking a view from the **Views** page. Add an object if necessary, by clicking **Select preview source** from the upper-right corner of the **Views** page. The preview of the view appears just below the **Views** option in the right pane.

You can select a view from the list, click the vertical ellipsis against each view, and select the various options such as edit, delete, clone, and export a view.

You can filter the views based on the name, type, description, subject, and owner. You can click the **Add** option to create a view. For information about creating a view, see Create and Configure a View.

Views are also categorized and listed in the **Views** panel based on the type of view and subject.

Table 18-2. Filter Groups

| Filter Group | Description |
| --- | --- |
| Name | Filter by the view name. For example, type `my view` to list all views that contain the `my view` phrase in their name. |
| Type | Filter by the view type. |

Table 18-2. Filter Groups (continued)

| Filter Group | Description |
| --- | --- |
| Description | Filter by the view description. For example, type **`my view`** to list all views that contain the **`my view`** phrase in their description. |
| Subject | Filter by the subject. |
| Owner | Filter by owner. |

# Datagrid Options

| Column Names | Description |
| --- | --- |
| Name | Displays the name of the view. |
| Type | Displays the type of view: list, summary, trend, distribution, text, or image. |
| Description | Displays the description of the dashboard. |
| Subject | Displays the base object type for which the view shows information. |
| Dashboard Usage | Displays the number of dashboards where the view is used. Click the number in the column to view the name of the dashboard/s. |
| Report Usage | Displays the number of reports where the view is used. Click the number in the column to view the name of the report/s. |
| Last Modified | Displays the date the view was last modified. |
| Modified By | Displays the user who last modified the view. |

## Views Actions

You can select more than one view and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

| Option | Description |
| --- | --- |
| Delete | Deletes the view. |
| Export | Downloads the view. |
| Import | Allows you to import a view by selecting a view in XML or zip file format.<br>To import a view:<br>■ Click the **Import** option from the horizontal ellipsis.<br>■ Click **Browse** and select a view XML or ZIP file to import.<br>■ Select if you want to Overwrite or Skip the file in case of a conflict.<br>■ Click **Import** to import the view, and click **Done**. |

# Views and Reports Ownership

The owner of views, reports, or templates might change over time.

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

# Create and Configure a View

To collect and display information for a specific object, you can create a custom view.

### Procedure

1. From the left menu, click **Visualize > Views**.
2. From the **Views** panel, click **Create**.
3. Select one of the following views from the right panel.
   - List View
   - Summary View

- Trend View
- Distribution View
- Text View
- Image View

4 At the end of each tab in the selected view, you can go to the previous or next tab. You can also cancel the creation of the view.

5 After you have added all the details, click **Create** to create the view.

## List View

List views provide tabular data about specific objects in the monitored environment that correspond to the selected view.

### Where You Find the List View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **List** from the right panel.

### Name and Configuration Tab

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Items per page | Select the number of items per page. Each item is one row and its metrics and properties are the columns. |
| Top result count | Select the top results. Restricts the number of results. For example, if you list all the clusters in a View, selecting 10 in this option displays the top 10 clusters with the relevant information. You can reduce the number of rows for the purposes of reporting. |
| Include Deleted Objects | Select to add deleted objects. |
| Show Objects | Select the type of object you want displayed in the view. You can select **Existing**, **Deleted**, or **All** objects. |
| Show Object Creation Date | Select to display the date the object was created. |
| Make the view available at > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available at > Report template creation and modification | Select if you want to make the view available in a report template. |

| Option | Description |
| --- | --- |
| Make the view available at > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click the subject for which you add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
| --- | --- |
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| Group By | You can group the results based on a parent object, by making a selection in the Group By drop-down option. If you generate a report based on the list view for which a group has been specified, the report displays group-based information for the selected object. You can also view summary calculations for the group of objects in the report, along with the total summary results for all the objects. |
| **Breakdown By** | |
| Add interval breakdown | Select this check box to see the data for the selected resources broken down in time intervals.<br>After you select this check box, you can enter a label, specify whether the values have to be sorted in ascending or descending order, and select a breakdown interval for the time range. |

| Option | Description |
| --- | --- |
| Add Instance breakdown | Select this check box to see the data for all instances of the selected resources.<br>After you select this check box, you can enter a label and select a metric group to break down all the instances in that group. Deselect **Show non-instance aggregate metric** to display only the separate instances. Deselect **Show only instance name** to display the metric group name and instance name in the instance breakdown column.<br>For example, you can create a view to display CPU usage by selecting the metric **CPU:0\|Usage**. If you add an instance breakdown column, the column **CPU:0\|Usage** displays the usage of all CPU instances on separate rows (0, 1, and so on). To avoid ambiguity, you can change the metric label of **CPU:0\|Usage** to **Usage**. |
| **Data Grid options** | |
| Data selection tree | Select a metric or a property. |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the **Units** drop-down menu, you can change the value to Hz, kHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order. |

| Option | Description |
| --- | --- |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>■ Minimum. The minimum value of the metric over the selected time range.<br>■ Maximum. The maximum value of the metric over the selected time range.<br>■ Average. The mean of all the metric values over the selected time range.<br>■ Sum. The sum of the metric values over the selected time range.<br>■ First. The first metric value for the selected time range.<br>■ Last. The last value of a metric within the selected time range.<br>■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>■ Standard Deviation. The standard deviation of the metric values.<br>■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>■ Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, **`sum/((max + min)/2)`**. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps.<br>■ Timestamp: You can choose between Absolute Timestamp or Relative Timestamp.<br>■ If applied to a numeric metric/property defined with a time-unit definition, the actual value is converted to a human readable timestamp. The metric value is rounded-off to an hour. |

| Option | Description |
| --- | --- |
| | ■ In the remaining cases, a timestamp is displayed when metrics and properties are added or modified. In this case, the behavior is the same as the Timestamp option selected for a non-Timestamp transformation.<br>Applicable for Absolute Timestamp and Relative Timestamp.<br>Available for List view and Minimum, Maximum, Current, First, and Last transformation |
| Ranges for metric coloring | You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the **Red Bound** field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br>This option is applicable to the Transformation configuration option. |

## Time Settings Tab

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data.

### Table 18-3. Time Settings Options

| Configuration Option | Description |
| --- | --- |
| Time Range Mode | In Basic mode, you can select date ranges.<br>In Advanced mode, you can select any combination of relative or specific start and end dates.<br>You can also activate the **Business Hour** option and select business hours/days for weekdays. |
| Relative Date Range | Select a relative date range of data transformation.<br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br>Available in Basic mode. |

Table 18-3. Time Settings Options (continued)

| Configuration Option | Description |
| --- | --- |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |
| Business Hours | Select business hours from Monday to Sunday by moving the sliders on the left and right sides to set the start and end time for each day of the week.<br>For example, as a VM owner, you can track the average utilization of VMs over a week (business days), during specified hours of the day (business hours).<br>This option is available for Minimum, Maximum, Average, Sum, and Percentile transformations<br>Available in Advanced mode for List Views. |

## Filter Tab

The filter option allows you to add additional criteria when the view displays too much information. For example, a List view shows information about the health of virtual machines. From the **Filter** tab, you add a risk metric less than 50%. The view displays the health of all virtual machines with risk less than 50%. For selected criteria you can also apply Business Hours, if the selected transformation type you add as a filter is supported by the business hours functionality.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

**Note** For Symptom based views, in the filter, if you select either **Symptom Definition Name**, **Alert Definition Name**, or **Alert Type**, it is recommended that you select a preview source that has a smaller number of objects.

Table 18-4. Filter Add Options

| Option | Description |
| --- | --- |
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

### Summary Tab

You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, from an existing or new view dialog box, click the **Summary** tab in the right pane. Click the plus sign to add a summary row.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

### Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

## Summary View

Summary views provide tabular data about the use of resources in the monitored environment.

### Where You Find the Summary View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **Summary** from the right panel.

## Name and Configuration Tab

| Option | Description |
| --- | --- |
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Items per page | Select the number of items per page. Each item is one row and its metrics and properties are the columns. |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click on the subject for which you want to add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
| --- | --- |
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Select a metric or property |

| Option | Description |
| --- | --- |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand(MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order. |

| Option | Description |
| --- | --- |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>■ Minimum. The minimum value of the metric over the selected time range.<br>■ Maximum. The maximum value of the metric over the selected time range.<br>■ Average. The mean of all the metric values over the selected time range.<br>■ Sum. The sum of the metric values over the selected time range.<br>■ First. The first metric value for the selected time range.<br>■ Last. The last value of a metric within the selected time range.<br>■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>■ Standard Deviation. The standard deviation of the metric values.<br>■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>■ Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, **`sum/((max + min)/2)`**. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them. |

| Option | Description |
| --- | --- |
| | This option is applicable to the Transformation configuration option. |

## Time Settings Tab

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. .

### Table 18-5. Time Settings Options

| Configuration Option | Description |
| --- | --- |
| Time Range Mode | In Basic mode, you can select date ranges.<br>In Advanced mode, you can select any combination of relative or specific start and end dates.<br>You can also activate the **Business Hour** option and select business hours/days for weekdays. |
| Relative Date Range | Select a relative date range of data transformation.<br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br>Available in Advanced mode. |

Table 18-5. Time Settings Options (continued)

| Configuration Option | Description |
| --- | --- |
| Specific End Date | Select a specific end date of data transformation.<br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

## Filter Tab

The filter option allows you to add additional criteria when the view displays too much information. For example, a view shows information about the health of virtual machines. From the **Filter** tab, you add a risk metric less than 50%. The view displays the health of all virtual machines with risk less than 50%.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 18-6. Filter Add Options

| Option | Description |
| --- | --- |
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

## Summary Tab

You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, from an existing or new view dialog box, click the **Summary** tab in the right pane. Click the plus sign to add a summary row.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

### Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

## Trend View

Trend views use historic data to generate trends and forecasts for resource use and availability in the monitored environment.

### Where You Find the Trend View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **Trend** from the right panel.

### Name and Configuration Tab

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| The maximum plot lines | Enter the maximum number of plot lines. Limits the output in terms of the objects displayed in the live preview of the view type on the left upper pane. The number you set as the maximum number of plot lines determines the plot lines.<br>For example, if you plot historical data and set the maximum at 30 plot lines, then 30 objects are displayed. If you plot historical, trend, and forecast lines, and set the maximum to 30 plot lines, then only 10 objects are displayed as each object has three plot lines. |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click on the subject for which you want to add data. Double-click the data from the tree in the left panel to add it to the view. For each subject the data available to add, might be different.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
| --- | --- |
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Select a metric or a property. |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the Units drop-down menu, you can change the value to Hz, KHz, or GHz. If you select Auto, the scaling is set to a meaningful unit. |
| Data Series | You can select whether to include historical data, trend of historical data, and forecast for future time in the trend view calculations. |
| Ranges for metric coloring | You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the **Red Bound** field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats. |

| Option | Description |
| --- | --- |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select `Sum` as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br>This option is applicable to the Transformation configuration option.<br>Available for all views. |
| Threshold Lines | You can set a threshold for a single metric:<br>■ None. You have not set a threshold.<br>■ By Symptom Definition. You can set a threshold value based on a symptom definition.<br>■ Custom. You can set the threshold value as **Warning**, **Critical**, or **Immediate**. These options are available only for the **Custom** option. |

## Time Settings Tab

Use the time settings to select the time interval of data transformation.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. .

### Table 18-7. Time Settings Options

| Configuration Option | Description |
| --- | --- |
| Time Range Mode | In Basic mode, you can select date ranges.<br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>Available in Basic mode. |

Table 18-7. Time Settings Options (continued)

| Configuration Option | Description |
| --- | --- |
| Relative Start Date | Select a relative start date of data transformation.<br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

## Filter

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 18-8. Filter Add Options

| Option | Description |
| --- | --- |
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

## Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

# Distribution View

Distribution views provide aggregated data about resource distribution in the monitored environment. When you add a distribution type of View to a dashboard, you can click a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment.

## Where You Find the Distribution View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **Distribution** from the right panel.

## Name and Configuration Tab

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Configuration** | |
| Visualization | You can view the data as a pie chart, a bar chart, or a donut chart. When you add a distribution type of View to a dashboard, you can click a section of the pie chart, or on one of the bars in the bar chart, or a section of the donut chart to view the list of objects filtered by the selected segment. You can select the display colors for single or multi-colored charts. |
| Coloring | |
| Colorize | The colors of the slices in the pie chart are displayed in the order of the colors in the color palette. |
| Select Color | Select the color that you want the chart to appear in. If there is more than one slice in a pie chart, the colors are chosen sequentially from the color palette. In a bar chart, the bars are all the same color. |
| Distribution Type > Dynamic Distribution | |
| Buckets Count | The number of buckets to use in the data distribution. |
| Buckets Size Interval | The bucket size is determined by the defined interval divided by the specified number of buckets. |
| Buckets > Size > Logarithmic bucketing | The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data. |
| Buckets > Size > Simple Max/Min bucketing | The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets. |

| Option | Description |
| --- | --- |
| Distribution Type > Manual Distribution | Specify the number of buckets and the minimum and maximum values of each bucket.<br>You can also select a color for each defined bucket that you specify. |
| Distribution Type > Discrete Distribution | Specify the number of buckets in which distributes the data.<br>If you increase the number of buckets, you can see more detailed data. |
| Distribution Type > Summary | |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a trend view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click the subject for which you add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
| --- | --- |
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |

| Option | Description |
| --- | --- |
| Data selection tree | Double-click to select a metric or a property. |
| Data column | Click the metric or property to enter configuration details in the Configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order. |

| Option | Description |
| --- | --- |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>■ Minimum. The minimum value of the metric over the selected time range.<br>■ Maximum. The maximum value of the metric over the selected time range.<br>■ Average. The mean of all the metric values over the selected time range.<br>■ Sum. The sum of the metric values over the selected time range.<br>■ First. The first metric value for the selected time range.<br>■ Last. The last value of a metric within the selected time range.<br>■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>■ Standard Deviation. The standard deviation of the metric values.<br>■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>■ Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, **`sum/((max + min)/2)`**. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select `Sum` as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them. |

| Option | Description |
| --- | --- |
| | This option is applicable to the Transformation configuration option. |

## Time Settings Tab

Use the time settings to select the time interval of data transformation.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data.

### Table 18-9. Time Settings Options

| Configuration Option | Description |
| --- | --- |
| Time Range Mode | In Basic mode, you can select date ranges.<br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br>Available in Advanced mode. |

Table 18-9. Time Settings Options (continued)

| Configuration Option | Description |
| --- | --- |
| Specific End Date | Select a specific end date of data transformation.<br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

## Filter Tab

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 18-10. Filter Add Options

| Option | Description |
| --- | --- |
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

## Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

# Text View

Text views allows you to insert provided text. The text can be dynamic and contain metrics and properties. You can format text to increase or decrease the font size, change the font color, highlight text, and align text to the left, right, or center. You can also make the selected text appear bold, in italics, or underlined. By default the text view is available only for report template creation and modification. You can change this in the **Visibility** option in the **Name and Configuration** tab.

## Where You Find the Text View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **Text** from the right panel.

## Name and Configuration Tab

| Option | Description |
| --- | --- |
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. This is the default option. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click the subject for which you add data. Double-click the data from the tree in the left panel to add it to the view. For each subject, the data available to add might be different. The data and configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
| --- | --- |
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Select a metric or property. |

| Option | Description |
| --- | --- |
| Data column | Click the metric or property to enter configuration details in the Configuration column.<br>You can also enter text to display in the view. |
| **Configuration column** | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand(MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |

| Option | Description |
| --- | --- |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>■ Minimum. The minimum value of the metric over the selected time range.<br>■ Maximum. The maximum value of the metric over the selected time range.<br>■ Average. The mean of all the metric values over the selected time range.<br>■ Sum. The sum of the metric values over the selected time range.<br>■ First. The first metric value for the selected time range.<br>■ Last. The last value of a metric within the selected time range.<br>■ Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>■ Standard Deviation. The standard deviation of the metric values.<br>■ Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>■ Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>■ Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>■ Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, **`sum/((max + min)/2)`**. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them. |

| Option | Description |
| --- | --- |
| | This option is applicable to the Transformation configuration option. |

## Time Settings

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. .

### Table 18-11. Time Settings Options

| Configuration Option | Description |
| --- | --- |
| Time Range Mode | In Basic mode, you can select date ranges.<br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br>Available in Advanced mode. |

Table 18-11. Time Settings Options (continued)

| Configuration Option | Description |
| --- | --- |
| Specific End Date | Select a specific end date of data transformation.<br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

## Filter

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

Table 18-12. Filter Add Options

| Option | Description |
| --- | --- |
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br>For instance metrics, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, and Sum. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

## Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

# Image View

Image views allow you to insert a static image. By default the image view is available only for report template creation and modification. You can change this in the **Visibility** option of the **Name and Configuration** tab.

### Where You Find the Image View

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Create**. Click **Image** from the right panel.

### Name and Configuration Tab

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. This is the default option. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

### Data Tab

### How to Add Data to a View

From the **Data** tab, browse and select an image from the left panel to add it to the view.

You can see a live preview of the view type when you select an image and then click **Select preview source**.

### Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

## Editing, Cloning, and Deleting a View

You can edit, clone, and delete a view. Before you do, familiarize yourself with the consequences of these actions.

## Edit a View

When you edit a view, all changes are applied to the report templates that contain it. To edit a view, from the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Edit**.

## Clone a View

When you clone a view, the changes that you make to the clone do not affect the source view. To clone a view, from the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Clone**.

## Delete a View

When you delete a view, it is removed from all the report templates that contain it. To delete a view, from the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Delete**.

# Share a View

You might want to present insights into your application or infrastructure performance to a wider audience, using tools such as Microsoft Power BI. By using the view share functionality, you can export a view to an external tool to ensure relevant stakeholders have the latest insights. You can use the provided API end point URL to import view data in a reporting tool. You can also export the configuration details of the view.

## Where You Can Access the Option to Share Views

From the left menu, click **Visualize > Views**. From the **Views** panel, click **Manage**. From the **Views** page on the right, click an existing view, select a preview source and then click the **Share View** icon in the top-right corner.

Table 18-13. Options in the Share View Dialog Box

| Option | Description |
| --- | --- |
| URL Tab | You can export view data using a URL. The resulting view data includes view configuration data, objects, and metrics metadata, per object metric data, summary and calculated data, and other view specific data. View data usually corresponds to the data shown in the View UI for the specified view configuration. You can copy and paste the URL in other reporting tools for integration with VMware Aria Operations.<br>■ Preview Source: You can select the last previewed source.<br>■ URL: You can copy the URL for the selected configured view. To use the URL as an authorized user, a CSP authentication token is required, you must generate a token separately on the CSP page. To generate a CSP authentication token, follow these steps:<br>a Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.<br>b In the **My Account** page, click **API Tokens**, and then click **Generate a New API Token**.<br>c Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.<br>d Click **Generate**.<br>e Copy or save the generated token.<br>**Note** The URL provided is the required REST call for fetching data. See the API Programming guide for details. |
| Export Tab | You can download and export the view configuration details. |

# Including Deleted VMs in List View

In VMware Aria Operations, you can view the deleted objects and the relationship of the objects in the list view. The objects can be VMs, deployments, projects, vApps, and edge gateways. You can also retain the relationship of the objects even after the objects are deleted from the system. The cost of the deleted virtual machines (VMs) is available until the retention period for that VM is over.

## Where You Find Global Settings for Deleted VMs

To specify for how long you want to retain the deleted virtual machines in VMware Aria Operations, from the left menu, click **Administration**, and then click the **Global Settings** tile. Navigate to **Data Retention** > **Deleted Objects**.

You can also specify the **Deletion Scheduling Interval** which specifies the number of hours between resource deletion scheduling.

Select **Object Deletion Schedule > Add Object Deletion Schedule** and select the virtual machine object from the **Object Kind** drop-down menu, specify the value, and click **Save**. The global setting value for the deleted virtual machine is updated in VMware Aria Operations.

For VMware Aria Automation, the price of the deleted VMs or deployments is added to the corresponding project object as a separate metric. If the deleted VM from VMware Aria Automation is associated with a cost-based pricing policy, then the price for that VM is not added to the corresponding project.

For vCloud Director, the price of deleted VMs, vApps, and Edge Gateways is added to the corresponding organization VDC object again as a separate metric. For vCenter Server, if VM is on unclustered Host, then deleted VM price is assigned to the Host, otherwise to the Cluster.

## How to Include Deleted VMs in List View

The deleted VMs are visible from **Environment** > **Inventory** > **Collection States** > **Not Existing**.

# User Scenario: Create, Run, Export, and Import a VMware Aria Operations View for Tracking Virtual Machines

As a virtual infrastructure administrator, you use VMware Aria Operations to monitor several environments. You must know the number of virtual machines on each vCenter Server instance. You define a view to gather the information in a specific order and use it on all VMware Aria Operations environments.

### Prerequisites

Verify that you have the necessary access rights to perform this task. Your VMware Aria Operations administrator can tell you which actions you can perform.

You will create a distribution view and run it on the main VMware Aria Operations environment. You will export the view and import it in another VMware Aria Operations instance.

### Procedure

1. Create a VMware Aria Operations View to view VM Memory Overhead

   To collect and display VM memory overhead vCenter Server, you create a custom view.

2. Run a View

   To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

3. Export a View

   To use a view in another VMware Aria Operations instance, you export a content definition XML file.

4 Import a View

To use views from other VMware Aria Operations environments, you import a content definition XML file.

## Create a VMware Aria Operations View to view VM Memory Overhead

To collect and display VM memory overhead vCenter Server, you create a custom view.

### Procedure

1 From the left menu, click **Visualize > Views**.

2 From the **Views** panel, click **Create** and then click **Distribution** from the right panel.

3 From the **New View** dialog box, in the **Name & Configuration** tab, enter **`Virtual Machines Memory Overhead`**, as the name for the view.

4 Enter a meaningful description for the view.

For example, **`A view showing the VM memory overhead`**.

5 From the **Configuration > Visualization** drop-down menu, select **Pie Chart**.

6 From the **Configuration > Distribution Type** option, select **Discrete distribution**.

Leave **Max number of buckets** deselected because you do not know the number of hosts on each vCenter Server instance. If you specify a number of buckets and the hosts are more than that number, one of the slices shows unspecified information labeled Others.

7 Click **Next**.

8 From the **Data** tab, click **Subject** to select the object type that applies to the view.

a From the drop-down menu, select **Host System**.

9 From the tree panel in the left, in the filter text box enter **`VM Overhead`**.

10 Double-click **Memory** > **`VM Overhead`** to add the metric.

11 Retain the default metric configurations and click **Create**.

## Run a View

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

### Prerequisites

Verify that you have the necessary access rights to perform this task. Your VMware Aria Operations administrator can tell you which actions you can perform.

### Procedure

1 From the left menu, click **Environment > Object Browser**.

2 From the **Object Browser** panel, navigate to a vCenter Server instance and click the **Details** tab.

   All listed views are applicable for the vCenter Server instance.

3 From the **All Filters** drop-down menu on the left, select **Type > Distribution**.

   You filter the views list to show only distribution type views.

4 Navigate to and click the **Virtual Machines Distribution** view.

   The bottom pane shows the distribution view with information about this vCenter Server. Each slice represents a host and the numbers on the far left show the number of virtual machines.

## Export a View

To use a view in another VMware Aria Operations instance, you export a content definition XML file.

If the exported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

### Prerequisites

Verify that you have the necessary access rights to perform this task. Your VMware Aria Operations administrator can tell you which actions you can perform.

### Procedure

1 From the left menu, click **Visualize > Views**.

2 From the **Views** panel, click **Manage**.

3 Select a view and click **Export** from the vertical ellipsis next to the selected view.

## Import a View

To use views from other VMware Aria Operations environments, you import a content definition XML file.

### Prerequisites

Verify that you have the necessary access rights to perform this task. Your VMware Aria Operations administrator can tell you which actions you can perform.

### Procedure

1 From the left menu, click **Visualize > Views**.

2 From the **Views** panel, click **Manage**.

3 Select a view and click the **Import** option from the horizontal ellipsis next to the **Add** option.

4 Browse to select the Virtual Machines Distribution content definition XML file and click **Import**.

If the imported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

**Note** The imported view overwrites if a view with the same name exists. All report templates that use the existing view are updated with the imported view.

# Configuring Super Metrics

19

The super metric is a mathematical formula that contains one or more metrics or properties. It is a custom metric that you design to help track combinations of metrics or properties, either from a single object or from multiple objects. If a single metric does not inform you about the behavior of your environment, you can define a super metric.

After you define it, you assign the super metric to one or more object types. This action calculates the super metric for the objects in that object type and simplifies the metrics display. For example, you define a super metric that calculates the average CPU usage on all virtual machines, and you assign it to a cluster. The average CPU usage on all virtual machines in that cluster is reported as a super metric for the cluster.

When the super metric attribute is activated in a policy, you can also collect super metrics from a group of objects associated with a policy.

Because super metric formulas can be complex, plan your super metric before you build it. The key to creating a super metric that alerts you to the expected behavior of your objects is knowing your own enterprise and data. Use this checklist to help identify the most important aspects of your environment before you begin to configure a super metric.

**Table 19-1. Designing a Super Metric Checklist**

| | |
|---|---|
| Determine the objects that are involved in the behavior to track. | When you define the metrics to use, you can select either specific objects or object types. For example, you can select the specific objects VM001 and VM002, or you can select the object type virtual machine. |
| Determine the metrics to include in the super metric. | If you are tracking the transfer of packets along a network, use metrics that refer to packets in and packets out. In another common use of super metrics, the metrics might be the average CPU usage or average memory usage of the object type you select. |
| Decide how to combine or compare the metrics. | For example, to find the ratio of packets in to packets out, you must divide the two metrics. If you are tracking CPU usage for an object type, you might want to determine the average use. You might also want to determine what the highest or lowest use is for any object of that type. In more complex scenarios, you might need a formula that uses constants or trigonometric functions. |

Table 19-1. Designing a Super Metric Checklist (continued)

| | |
|---|---|
| Decide where to assign the super metric. | You define the objects to track in the super metric, then assign the super metric to the object type that contains the objects being tracked. To monitor all the objects in a group, activate the super metric in the policy, and apply the policy to the object group. |
| Determine the policy to which you add the super metric. | After you create the super metric, you add it to a policy. For more information, refer to Policy Workspace in VMware Aria Operations. |

# What Else Can You Do with Super Metrics

- To create alert definitions to notify you of the performance of objects in your environment, define symptoms based on super metrics. For more information, refer to Symptom Definitions in VMware Aria Operations.
- Learn about the use of super metrics in policies. For more information, refer to Policy Workspace in VMware Aria Operations.
- Use OPS CLI commands to import, export, configure, and delete super metrics. For more information, refer to the OPS CLI documentation.
- To display metric-related widgets, create a custom set of metrics. You can configure one or more files that define different sets of metrics for a particular adapter and object types. This ensures that the supported widgets are populated based on the configured metrics and selected object type.

Read the following topics next:

- Create a Super Metric
- Enhancing Your Super Metrics
- Exporting and Importing a Super Metric
- Super Metrics Tab

# Create a Super Metric

Create a super metric when you want to check the health of your environment, but you cannot find a suitable metric to perform the analysis.

### Procedure

1. From the left menu, click **Configure** and then click **Super Metrics**.
2. Click **Add** .

   The **Create Super Metric** wizard opens.

3 (Optional) To edit a super metric, click the vertical ellipsis next to the super metric and select **Edit**. You can also edit the super metric using the **EDIT** option in super metrics page.

4 Enter a meaningful name for the super metric such as **`Worst VM CPU Usage (%)`** in the **Name** text box.

**Note** It is important that you have an intuitive name as it appears in dashboards, alerts, and reports. For meaningful names, always use space between words so that it is easier to read. Use title case for consistency with the out of the box metrics and add the unit at the end.

5 Provide a brief summary of the super metric in the **Description** text box.

**Note** Information regarding the super metric, like why it was created and by whom can provide clarity and help you track your super metrics with ease.

6 From the **Object Types** drop-down list, select the object to associate with the super metric and click **Next**.

7 Create the formula for the super metric.

For example, to add a super metric that captures the average CPU usage across all virtual machines in a cluster, perform the following steps.

a Select the function or operator. This selection helps combine the metric expression with operators and/or functions. In the super metric editor, enter **`avg`** and select the **avg** function.

You can manually enter functions, operators, objects, object types, metrics, metrics types, property, and properties types in the text box and use the suggestive text to complete your super metric formula.

Alternatively, select the function or operator from the **Functions** drop-down menu.

b To create a metric expression, enter **`Virtual`** and select **Virtual Machine** from the object type list.

c Add the metric type, enter **`usage`**, and select the **CPU|Usage (%)** metric from the metric type list.

**Note** The expression ends with depth=1 by default. If the expression ends with depth=1, that means that the metric is assigned to an object that is one level above virtual machines in the relationship chain. However, since this super metric is for a cluster which is two levels above virtual machine in the relationship chain, change the depth to 2.

The depth can also be negative, this happens when you need to aggregate the parents of a child object. For example, when aggregating all the VMs in a datastore, the metric expression ends with depth=-1, because VM is a parent object of datastore. But, if you want to aggregate all the VMs at a Datastore Cluster level, you need to implement 2 super metrics. You cannot directly aggregate from VM to Datastore Cluster, because both are parents of a datastore. For a super metric to be valid, depth cannot be 0 (-1+1=0). Hence, you need to create the first super metric (with depth=-1) for the aggregate at the datastore level, and then build the second super metric based on the first (with depth = 1).

The metric expression is created.

d To calculate the average CPU usage of powered on virtual machines in a cluster, you can add the `where` clause. Enter **`where=""`**.

**Note** The **`where`** clause cannot point to another object, but can point to a different metric in the same object. For example, you cannot count the number of VMs in a cluster with the CPU contention metric > SLA of that cluster. The phrase "SLA of that cluster " belongs to the cluster object, and not to the VM object. The right operand must also be a number and cannot be another super metric or variable. The where clause cannot be combined using AND, OR, NOT, which means you cannot have `where="VM CPU>4 and VM RAM>16"` in your super metric formula.

e Position the pointer between the quotation marks, enter **Virtual**, and select the **Virtual Machine** object type and the **System|Powered ON** metric type.

f To add the numeric value for the metric, enter **`==1`**.

g To view hints and suggestions, click **ctrl+space** and select the objects, object types, metrics, metrics types, property, and properties types to build your super metric formula.

h Select **This** option from the drop-down menu.

If **This** option is selected during the creation of a metric expression, it means that the metric expression is associated to the object for which the super metric is created.

8 Select the unit of the super metrics from the **Unit** drop-down.

**Note** The super metrics unit configured here can be changed in the metrics charts, widgets, and views.

9 Click **Validate** to verify that the super metric formula has been created correctly.

After you click the **Preview** button the system selects a random object and displays a metric graph showing values for the current super metric. For example, if you have selected `Host System` in the **Object Types** tab, after you click the **Preview** button it will randomly select a host system object from the list of the available objects and displays the graph for the selected host. Alternatively, you can also type in the object name in the **Object** text box, and the result will also depend on the pre-selected object type.

a Expand the **Preview** section.

A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

b Click **Next**.

The Policies page is displayed.

10 Select the policy which you want to associate with super metric and click **Update**.

The selected policy is applied to the super metric. You can view the super metric you created, the associated object type, and policy on the **Super Metrics** page.

# Enhancing Your Super Metrics

You can enhance your super metrics by using clauses and resource entry aliasing.

## Where Clause

The **`where`** clause verifies whether a particular metric value can be used in the super metric. Use this clause to point to a different metric of the same object, such as **`where=(${metric=metric_group|my_metric} > 0)`**.

For example: **`count(${objecttype = ExampleAdapter, adaptertype = ExampleObject, metric = ExampleGroup|Rating, depth=2, where =($value==1})`**

## IsFresh Function

Use the **`isFresh`** function in the **`where`** clause to check if the last value of the metrics is fresh or not.

For every metric published in VMware Aria Operations, the point with the latest publishing time is called as the last point of that metric. The value of that metric's last point is called the last value of that metric. A metric's last point is considered fresh when the time elapsed after the metric's last point is lesser than the estimated publishing interval of that metric.

The **`isFresh`** function returns true if the last value of the metrics is fresh. For example, in the following scenarios, the function:

- **`${this, metric=a|b, where=($value.isFresh())}`**, returns the last value of the metric a|b if the last value is fresh.

- **`${this, metric=a|b, where=($value == 7 && $value.isFresh())}`**, returns the last value of the metric a|b if it is equal to seven and is fresh.
- **`${this, metric=a|b, where=(${metric=c|d} == 7 && ${metric=c|d}.isFresh())}`**, returns the last value of the metric a|b only if the last value of the metric c|d is equal to seven and is fresh.

## Resource Entry Aliasing

Resource entries are used to retrieve metric data from VMware Aria Operations for computing super metrics. A resource entry is the part of an expression which begins with **`$`** followed by a **`{..} block`**. When computing a super metric, you might have to use the same resource entry multiple times. If you have to change your computation, you must change every resource entry, which might lead to errors. You can use resource entry aliasing to rewrite the expression.

The following example, shows a resource entry that has been used twice.

```
(min(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)/
(max(${adaptertype=VMWARE, objecttype=HostSystem, attribute=cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)"
```

The following example shows how to write the expressing using resource entry aliasing. The output of both expressions is the same.

```
(min(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)} as cpuload) + 0.0001)/
(max(cpuload) + 0.0001)"
```

Follow these guidelines when you use resource entry aliasing:

- When you create an alias, make sure that after the resource entry you write **`as`** and then **`alias:name`**. For example: **${...} as alias_name**.
- The alias cannot contain the ()[]+-*/%|&!=<>,.?:$ special characters, and cannot begin with a digit.
- An alias name, like all names in super metric expressions, is case-insensitive.
- Use of an alias name is optional. You can define the alias, and not use it in an expression.
- Each alias name can be used only once. For example: **`${resource1,…} as r1 + ${resource2,…} as R1`**.
- You can specify multiple aliases for the same resource entry. For example: **${...} as a1 as a2**.

## Conditional Expression `?:` Ternary Operators

You can use a ternary operator in an expression to run conditional expressions.

For example: **`expression_condition ? expression_if_true : expression_if_false`**.

The result of the conditional expression is converted to a number. If the value is not 0, then the condition is assumed as true.

For example: **`-0.7 ? 10 : 20`** equals 10. **`2 + 2 / 2 - 3 ? 4 + 5 / 6 : 7 + 8`** equals 15 (7 + 8).

Depending on the condition, either **`expression_if_true`** or **`expression_if_false`** is run, but not both of them. In this way, you can write expressions such as, **`${this, metric=cpu|demandmhz} as a != 0 ? 1/a : -1`**. A ternary operator can contain other operators in all its expressions, including other ternary operators.

For example: **`!1 ? 2 ? 3 : 4 : 5`** equals 5.

# Exporting and Importing a Super Metric

You can export a super metric from one VMware Aria Operations instance and import it to another VMware Aria Operations instance. For example, after developing a super metric in a test environment, you can export it from the test environment and import it use in a production environment.

If the super metric to import contains a reference to an object that does not exist in the target instance, the import fails. VMware Aria Operations returns a brief error message and writes detailed information to the log file.

Procedure

1 Export a super metric.
   a From the left menu, select **Configure** and then click **Super Metrics**.
   b Select the super metric to export, click horizontal ellipsis and then click **Export**.

      VMware Aria Operations creates a super metric file, for example, `SuperMetric.json`.
   c Download the super metric file to your computer.
2 Import a super metric.
   a From the left menu, select **Configure** and then click **Super Metrics**.
   b Click the horizontal ellipsis and then click **Import**.
   c (Optional). If the target instance has a super metric with the same name as the super metric you are importing, you can either overwrite the existing super metric or skip the import, which is the default.

# Super Metrics Tab

A super metric is a mathematical formula that contains a combination of one or more metrics for one or more objects. With super metrics you can assess information more quickly when you are observing fewer metrics.

## Where You Configure Super Metrics

From the left menu, click **Configure** and then click **Super Metrics**.

To view the details for a specific super metric, click the super metric from the list. The super metric details are displayed in the right-side panel. The super metric details include the assigned object types, formula and policies activated for the selected super metric.

### Table 19-2. Configuration Options for Super Metrics

| Option | Description |
| --- | --- |
| Toolbar | Use the toolbar selections to manage super metric options.<br>■ Add New Super Metric. Starts the Create Super Metric workspace.<br>■ Edit Selected Super Metric. Starts the Create Super Metric workspace.<br>■ Clone Selected Super Metric. Duplicates the super metric. Edit the clone or associate it with a different object type.<br>■ Delete Selected Super Metric.<br>■ Export Selected Super Metric. Exports a super metric to use in another VMware Aria Operations instance. See Exporting and Importing a Super Metric.<br>■ Import Super Metric. Imports a super metric to this VMware Aria Operations instance. See Exporting and Importing a Super Metric. |
| Super Metrics list | Configured super metrics listed by name and formula description. |

## Enhancements to the Super Metric Functions

In the earlier implementation of aggregate functions in super metrics, you had to explicitly specify the Adapter Kind and Resource Kind in the formula.

Old Formula

```
count(${adaptertype=VMWARE, objecttype=HostSystem,attribute=badge|health, depth=1})
```

The new implementation of aggregate function provides a way to define a super-metric without explicitly specifying the Resource Kind. You can use "objecttype=*" in the super-metric formula which indicates you to consider all Resource Kinds having the specified attribute.

New Formula

```
count(${adaptertype=VMWARE, objecttype=*,attribute=badge|health, depth=1})
```

**Note** The explicit specification of "adaptertype" remains mandatory. However, "*" can be used only to select all Resource Kinds for the given Adapter Kind.

# Manage Super Metric Workspace

You use the Manage Super Metric workspace to create or edit a super metric. The toolbar helps you to build the mathematical formula with the objects and metrics you select.

## Where You Configure Super Metrics

From the left menu, click **Configure** and then click **Super Metrics**.

### Table 19-3. Super Metrics Workspace Options

| Option | Description |
|---|---|
| Super Metric | ■ Name. The name you give to the super metric.<br>■ Description. The textual description you give about the super metric. |
| Object Types Pane | Use this page to associate the super metric with an object type. You can use this list to select the object type with the metrics to measure. The object type selection affects the list of objects, metrics, and attribute types displayed. |
| Formula | Define the formula you want to associate with the super metric. You can preview and validate the formula before you create it. Use drop-down menu to select the metrics to add to the formula.<br>■ Functions. Mathematical functions that operate on a single object or group of objects. See Super Metric Functions and Operators.<br>■ Operators. Mathematical symbols to enclose or insert between functions. See Enhancing Your Super Metrics.<br>■ This Object. Assigns the super metric to the object selected in the Object pane and displays `this` in the formula instead of a long description for the object.<br>■ Unformatted. Slide the toggle button to the right to view the unformatted version of the formula, slide the button to the left to view the formula in textual format.<br>■ Preview. Shows the super metric in a graph. Look at the graph so that you can verify that VMware Aria Operations is calculating the super metric for the target objects that you selected. |
| Policies | Displays the policies related to the object types you assigned your super metric to. |

# Super Metric Functions and Operators

VMware Aria Operations includes functions and operators that you can use in super metric formulas. The functions are either looping functions or single functions.

## Looping Functions

Looping functions work on more than one value.

### Table 19-4. Looping Functions

| Function | Description |
| --- | --- |
| avg | Average of the collected values. |
| combine | Combines all the values of the metrics of the included objects in a single metric timeline. |
| count | Number of values collected. |
| max | Maximum value of the collected values. |
| min | Minimum value of the collected values. |
| sum | Total of the collected values. |

## Looping Function Arguments

The looping function returns an attribute or metric value for an object or object type. An attribute is metadata that describes the metric for the adapter to collect from the object. A metric is an instance of an attribute. The argument syntax defines the desired result.

For example, CPU usage is an attribute of a virtual machine object. If a virtual machine has multiple CPUs, the CPU usage for each CPU is a metric instance. If a virtual machine has one CPU, then the function for the attribute or the metric return the same result.

### Table 19-5. Looping Function Formats

| Argument syntax example | Description |
| --- | --- |
| *funct*(${this, metric =*a*\|*b:optional_instance*\|*c*}) | Returns a single data point of a particular metric for the object to which the super metric is assigned. This super metric does not take values from the children or parents of the object. |
| *funct*(${this, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points for attributes of the object to which the super metric is assigned. This super metric does not take values from the child or parent of the object. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={id1=*val1*,id2=*val2*,...}, metric=*a*\|*b:instance*\|*c*}) | Returns a single data point of a particular metric for the *resname* specified in the argument. This super metric does not take values from the children or parents of the object. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={*id1=val1, id2=val2,...*}, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates attributes of the *resname* specified in the argument. This super metric does not take values from the child or parent of the object. |

### Table 19-5. Looping Function Formats (continued)

| Argument syntax example | Description |
|---|---|
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, metric=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates metrics of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects, where *depth* describes the object location in the relationship chain.<br>For example, a typical relationship chain includes a data center, cluster, host, and virtual machines. The data center is at the top and the virtual machines at the bottom. If the super metric is assigned to the cluster and the function definition includes depth = 2, the super metric takes values from the virtual machines. If the function definition includes depth = -1, the super metric takes values from the data center. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates attributes of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects. |

For example, `avg(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|usage_average, depth=1})` averages the value of all metric instances with the cpu|usage_average attribute for all objects of type `VirtualMachine` that the vCenter adapter finds. VMware Aria Operations searches for objects one level below the object type where you assign the super metric.

## Single Functions

Single functions work on only a single value or a single pair of values.

### Table 19-6. Single Functions

| Function | Format | Description |
|---|---|---|
| *abs* | abs(x) | Absolute value of x. x can be any floating point number. |
| *acos* | acos(x) | Arccosine of x. |
| *asin* | asin(x) | Arcsine of x. |
| *atan* | atan(x) | Arctangent of x. |
| *ceil* | ceil(x) | The smallest integer that is greater than or equal to x. |
| *cos* | cos(x) | Cosine of x. |
| *cosh* | cosh(x) | Hyperbolic cosine of x. |
| *exp* | exp(x) | e raised to the power of x. |
| *floor* | floor(x) | The largest integer that is less than or equal to x. |
| *log* | log(x) | Natural logarithm (base *x*) of x. |
| *log10* | log10(x) | Common logarithm (base 10) of x. |
| *pow* | pow(x,y) | Raises x to the y power. |

### Table 19-6. Single Functions (continued)

| Function | Format | Description |
|---|---|---|
| *rand* | rand() | Generates a pseudo random floating number greater than or equal to 0.0 and less than 1.0. |
| *sin* | sin(x) | Sine of x. |
| *sinh* | sinh(x) | Hyperbolic sine of x. |
| *sqrt* | sqrt(x) | Square root of x. |
| *tan* | tan(x) | Tangent of x. |
| *tanh* | tanh(x) | Hyperbolic tangent of x. |

## Operators

Operators are mathematical symbols and text to enclose or insert between functions.

### Table 19-7. Numeric Operators

| Operators | Description |
|---|---|
| + | Plus |
| - | Subtract |
| * | Multiply |
| / | Divide |
| % | Modulo |
| == | Equal |
| != | Not equal |
| < | Less than |
| <= | Less than, or equal |
| > | Greater than |
| >= | Greater than, or equal |
| \|\| | Or |
| && | And |
| ! | Not |
| ? : | Ternary operator. If/then/else<br>For example: **`conditional_expression ? expression_if_condition_is_true : expression_if_condition_is_false`**<br>For more information about ternary operators, see Enhancing Your Super Metrics. |

### Table 19-7. Numeric Operators (continued)

| Operators | Description |
| --- | --- |
| ( ) | Parentheses |
| [ ] | Use in an array of expressions |
| [x, y, z] | An array containing x, y, z. For example, min([x, y, z]) |

### Table 19-8. String Operators

| String Operators | Description |
| --- | --- |
| equals | Returns true if metric/property string value is equal to specified string. |
| contains | Returns true if metric/property string value contains specified string. |
| startsWith | Returns true if metric/property string value starts with the specified prefix. |
| endsWith | Returns true if metric/property string value ends with the specified suffix. |
| !equals | Returns true if metric/property string value is not equal to specified string. |
| !contains | Returns true if metric/property string value does not contain specified string. |
| !startsWith | Returns true if metric/property string value does not start with the specified prefix. |
| !endsWith | Returns true if metric/property string value does not end with the specified suffix. |

**Note** String operators are valid in 'where' condition only. For example: **`${this, metric=summary|runtime|isIdle, where = "System Properties|resource_kind_type !contains GENERAL"}`**

# Configuring Administration Settings

20

After VMware Aria Operations is installed and configured, you can use administration settings to manage your environment. You find most administration settings under the Administration selection of the VMware Aria Operations interface.

Read the following topics next:


- VMware Aria Operations Maintenance Schedules
- Manage Maintenance Schedules
- Managing Users and Access Control in VMware Aria Operations
- VMware Aria Operations Certificates
- Adding Physical Data Centers in VMware Aria Operations
- Modifying Global Settings
- Managing Content
- Managing Orphaned and Unassigned Content
- VMware Aria Operations Dynamic Thresholds
- Customizing Icons


## VMware Aria Operations Maintenance Schedules

Maintenance schedules identify objects that are in maintenance mode at specific times, which prevents VMware Aria Operations from showing misleading data based on those objects being offline or in other unusual states because of maintenance.

Many objects in the enterprise might be intentionally taken offline. For example, a server might be deactivated to update software. If VMware Aria Operations collects metrics when an object is offline, it might generate incorrect anomalies and alerts that affect the data for setting dynamic thresholds for the object attributes. When an object is identified as being in maintenance mode, VMware Aria Operations does not collect metrics from the object or generate anomalies or alerts for it. In addition, VMware Aria Operations cancels any active symptoms and alerts for the object.

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object in maintenance mode from midnight until 3 a.m. each Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can manually put an object in maintenance mode, or take it out of maintenance mode, even it if has an assigned maintenance schedule. For more information, see Manage Maintenance Schedules for Your Object Workspace.

## How Maintenance Schedules Work

Maintenance schedules require that you select the days and time-of-day when updates or other object maintenance occurs. Note that creating a maintenance schedule does not activate the schedule. A maintenance schedule must be part of a policy before the schedule can take effect. For more information, see Maintenance Schedule Details.

## Where You Find the Maintenance Schedules

From the left menu, click **Configure** > **Maintenance Schedules**.

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 20-1. Maintenance Schedule Toolbar Options

| Option | Description |
| --- | --- |
| Add | Open a window in which you can select the maintenance schedule settings for a new schedule. |
| Edit | Change the maintenance schedule settings for an existing schedule. |
| Delete | Remove the selected maintenance schedule. |

# Manage Maintenance Schedules

Add or edit a maintenance schedule to take an object offline. VMware Aria Operations does not collect data from an object that is offline.

## Where You Find Manage Maintenance Schedules

1 From the left menu, click **Configure** > **Maintainance Schedules**.

2 Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

Table 20-2. Manage Maintenance Schedule Add or Edit Options

| Option | Description |
| --- | --- |
| Schedule Name | Name that describes the maintenance schedule |
| Time Zone | Time zone in which you are currently located |
| Days | Number of days the maintenance period covers |

Table 20-2. Manage Maintenance Schedule Add or Edit Options (continued)

| Option | Description |
| --- | --- |
| Recurrence | Specify a maintenance schedule to run over a selected period<br>■ Once<br>■ Daily<br>■ Weekly<br>■ Monthly |
| Expire after | The number of times the schedule is run |
| Expire on | The date upon which the schedule stops running |

# Managing Users and Access Control in VMware Aria Operations

To ensure security of the objects in your VMware Aria Operations instance, as a system administrator you can manage some aspects of user access control. You can assign each user to be a member of one or more user groups, and assign roles to each user or user group to set their privileges.

Users must have privileges to access specific features in the VMware Aria Operations user interface. Access control is defined by assigning privileges to both users and objects. You can assign one or more roles to users, and activate them to perform a range of different actions on the same types of objects. For example, you can assign a user with the privileges to delete a virtual machine, and assign the same user with read-only privileges for another virtual machine.

## User Access Control

In order to use your corporate credentials to log on to VMware Cloud Service including VMware Aria Operations, you can federate Identity and Access Management to your corporate domain(s). For more information, see Setting Up Enterprise Federation with VMware Cloud Services in the VMware Cloud services Product Documentation.

## Roles and Privileges in VMware Aria Operations

VMware Aria Operations provides two predefined roles- the GeneralUser role and the Administrator role. These two roles are assigned to the user by the organization owner from VMware Cloud Services portal.

Each predefined role includes a set of privileges for users to perform

### Administrator

Includes privileges to manage the VMware Aria Operations instance and its objects. The Administrator can also customize the privileges associated with the GeneralUser role.

### GeneralUser

This role is defined role out of the box. This is the only default editable role, which has more permissions than ReadOnly user and less than Administrator.

**ReadOnly**

Users have read-only access and can perform read operations, but cannot perform write actions such as create, update, or delete.

## User Scenario: Manage User Access Control

As a system administrator or virtual infrastructure administrator, you manage user access control in VMware Aria Operations so that you can ensure the security of your objects. Your company just hired a new person, and you must assign a role to the account so that the new user has permission to access specific content and objects in VMware Aria Operations.

In this scenario you will learn how to assign roles to the user accounts to specify access privileges to views and objects. You will then demonstrate the intended behavior of the permissions on these accounts.

#### Prerequisites

Verify that the following conditions are met:

- VMware Aria Operations is installed and operating properly, and contains objects such as clusters, hosts, and virtual machines.
- One or more user groups are defined.

#### What to do next

Assign a new role.

### Create a New Role

You use roles to manage access control for user accounts in VMware Aria Operations.

In this procedure, you will add a new role and assign administrative permissions to the role.

#### Prerequisites

Verify that you understand the context of this scenario. See User Scenario: Manage User Access Control. For information about roles and associated permissions, see KB 59484.

#### Procedure

1. From the left menu, click **Administration**, and then click the **Access Control** tile.
2. Click the **Roles** tab.
3. Click the **Add** icon on the toolbar to create a role.

   The **Create Role** dialog box appears.

4 For the role name, type **`admin_cluster`**, then type a description and click **OK**.

   The admin_cluster role appears in the list of roles.

5 Click the **admin_cluster** role.

6 In the Details grid below, on the Permissions pane, click the **Edit** icon.

   The **Assign Permissions to Role** dialog box appears.

7 Select the **Administrative Access - all permissions** check box.

8 Click **Update**.

   This action gives this role administrative access to all the features in the environment.

#### What to do next

Create a user account, and assign this role to the account.

## Access Control in VMware Aria Operations

Each user must have a unique account with one or more roles assigned to enforce a role-based security when they use VMware Aria Operations. Contact your organization owner to create user accounts for VMware Aria Operations. Once created, you can assign the account to be a member of one or more user groups to allow the user to inherit the roles and scopes associated with the user group.

### Where You Find the Access Control Options

You can manage user accounts and their associated user groups, roles, scopes and passwords.

From the left menu, click **Administration**, and then click the **Access Control** tile.

Table 20-3. Access Control Tabs

| Option | Description |
| --- | --- |
| User Accounts | Edit VMware Aria Operations user accounts and manage user roles, their membership in groups, and the scopes assigned for association with the user. |
| User Groups | Add, edit, delete, or clone VMware Aria Operations user groups, update the members in a group and the associated roles and scopes that they can access. |
| Roles | For users to perform actions in VMware Aria Operations, they must be assigned specific roles. With role-based access, when you assign a role to a user, you are determining not only what actions the user can perform in the system, but also the objects upon which those actions can be performed while holding the role. For example, to import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management. |
| Scopes | Add, edit, clone, or remove scope associated with users or groups in VMware Aria Operations. Scope lets you limit the access of a user or a set of users to VMware Aria Operations. You can also define the scope for all the objects managed by VMware Aria Operations. |

## Access Control: User Accounts Tab

With access control, you manage roles, the objects a user can access while assigned a specific role, and the membership in user groups.

### Where You Manage User Accounts

From the left menu, click **Administration**, and then click the **Access Control** tile.

To view the details of the user account configured in VMware Aria Operations, click the user account, the user account details is displayed in the right-side panel. The user account details include the user groups associated with the user account along with the roles and scope of the selected user account.

**Note** For user accounts only the exclusive roles that is associated with the user account is displayed. The roles inherited from the user groups are not displayed.

Table 20-4. Access Control User Accounts Summary Grid

| Summary Grid Options | Description |
|---|---|
| User Accounts toolbar | To manage user accounts, use the toolbar options.<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>  ■ **Edit**. Edit the selected user account, and modify the details for the user account in the **Edit User Account** dialog box. |
| User Name | User name, without spaces, that you use to log in to VMware Aria Operations |
| First Name | User's first name, created when you create the user account. |
| Last Name | User's last name, created when you create the user account. |
| Email | User's email address, created when you create the user account. |
| Description | Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access privileges. |
| Imported | Indicates whether the user account is imported or not. |
| Source Type | Indicates whether the user account platform. |
| Activated | Indicates whether the user account is activated to use VMware Aria Operations features. An administrator can edit a user account to manually activate it, or deactivate it to prevent user access to VMware Aria Operations. |
| Locked | Indicates whether VMware Aria Operations has locked the user account. For example, a user account can get locked based on the password lockout policy, or if the user enters an incorrect password three times in the span of five minutes. |
| Access All Objects | Indicates whether the user account is allowed to access all the objects that are imported into the VMware Aria Operations instance. |
| Modified By | Indicates the last person to update the user account. |
| Last Modified | Indicates the last time the user account was updated. |

Table 20-4. Access Control User Accounts Summary Grid (continued)

| Summary Grid Options | Description |
| --- | --- |
| Last Login Time | Indicates the last time the user logged in. |
| Filters | Activates you to search the list of user accounts according to the following criteria:<br>■ User Name<br>■ First Name<br>■ Last Name<br>■ Description<br>■ Activated<br>■ Locked<br>■ Access All Objects<br>■ Modified By<br>■ Last Login Time |

Use the Details grid to view and edit which user accounts are assigned to user groups, and view the permissions assigned to the user account.

Table 20-5. Access Control User Accounts Details Grid

| Details Grid Options | Description |
| --- | --- |
| First Name | User's first name, created when you create the user account. |
| Last Name | User's last name, created when you create the user account. |
| Modified By | Displays the name of the user who last updated the user account. |
| Last Modified | Displays the date and time when the user account was last updated. |
| Roles and Scope | Displays the role and scope associated with the selected user.<br>■ Roles: Indicates the name of the role or roles assigned to the user.<br>■ Scope : Indicates the scope associated with the user account. |
| User Groups | Assigned user groups appear when you click a user in the summary grid. You can then view and modify which user groups the user is associated with. |

## Modify User Accounts and Assign Groups and Permissions

Contact your organization owners to add user accounts so that users can access the features of VMware Aria Operations and certain objects in the environment. After the organization owners add user accounts, you can assign them to one or more user groups, and assign roles and scopes. Assign the administrators role only to specific users who must access objects and perform actions in the entire environment.

### Where You Add or Edit User Accounts

1 To modify a user account, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 To edit a user account, click the vertical ellipsis and select **Edit**. You can also click the **EDIT ACCOUNT** button in the **User Details** page and edit the user account.

The **Edit User Information** page is displayed.

- By default, the user is added to the "Everyone" group. You can add or remove users to or from the user group when you create or edit the user group, . You can assign the permissions to the users by selecting the required role-scope pairs.
- When you add users to a user group, the user inherits the permissions (role-scope pairs) from that group. You also have the option of not assigning permissions to the user, in that case the user actions are restricted, and the user cannot login to VMware Aria Operations.
- You cannot change the user name once you create it. So, select your username appropriately.
- You have read only access for the Super Administrator account. You can change the password for the Super Administrator account from the user interface only.

**Assign Roles and Scope**

The Assign Roles and Scope option lets you select a role for each user and then assign a scope for that role.

3 From the **Select Role** drop-down menu, select a role for the user.

4 From the **Select Scope** drop-down menu, assign a scope for the selected user.

**Note** You can click the + sign to add multiple roles and then assign the required scope for each role.

**Assign User Groups**

The Assign User Group option lets you select user groups that the user will be a member of.

5 Select the user groups for which you want the user to be member of.

6 Click **Save**.

**Note** You cannot manually add users to groups imported from LDAP and SSO.

## Access Control: User Groups Tab

You can manage the user groups associated with the users and objects in your environment.

### Where You Manage User Groups

1 To manage user groups, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Click the **User Groups** tab.

To view the details of the user groups configured in VMware Aria Operations, click the user group, the user group details are displayed in the right-side panel. The user group details include the user accounts associated with the selected user group along with the roles and scope associated with the user group.

### Table 20-6. Access Control User Groups Summary Grid

| Option | Description |
| --- | --- |
| User Groups toolbar | To manage user groups, use the toolbar options.<br>■ Click the **Add** button to add a user group, and provide the details for the user group in the **Group Information** page.<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>  ■ **Edit**. Edit the selected user group, and modify the details for the user group in the **Edit User Group** dialog box.<br>  ■ **Delete**. Delete a user group.<br>  ■ **Clone**. Clone a user group, and enter a name and description for the cloned user group.<br>  ■ **Export**. Click export to export a user group, and provide the details to export the user group in the **Export Authentication Source** dialog box. |
| Group Name | Name of the user group. |
| Description | Description of the group, indicating its purpose. |
| Members | Number of members in the group. |
| Group Type | Type of group like a local user group |
| Distinguished Name | Names of users. |
| Access All Objects | Indicates if the user group account is allowed to access all the objects that are imported into the VMware Aria Operations instance. |
| Modified By | Indicates the last person to update the user group. |
| Last Modified | Indicates the last time the user group was updated. |
| Filters | Enables you to search the list of user groups according to the following criteria:<br>■ Group Name<br>■ Description<br>■ Access All Objects<br>■ Modified By |

After you select a user group in the summary grid, view details about associated users in the Details pane.

### Table 20-7. Access Control User Groups Details Grid

| Option | Description |
| --- | --- |
| User Group Name | User Group name, created when you create the user group. |
| Modified By | Displays the name of the user who last updated the user group. |
| Last Modified | Displays the date and time when the user group was last updated. |
| Description | Description of the group, indicating its purpose. |

Table 20-7. Access Control User Groups Details Grid (continued)

| Option | Description |
| --- | --- |
| Roles and Scope | Displays the role and scope associated with the selected user.<br>■ Role: Indicates the name of the role or roles assigned to the user.<br>■ Scope: Indicates the scope associated with the user group. |
| User Accounts | Displays the user accounts associated with the selected user group. |

## Add User Groups and Assign Members and Roles

You can view and modify the details for user groups, including users, roles, and scope.

### Where You Add User Groups

1 To add a user group, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Select the **User Groups** tab and then click **Add**

3 To edit a user group, click the vertical ellipsis and select **Edit**. You can also click the **EDIT GROUP** button in the **Group Details** page and edit the user group.

Table 20-8. Add or Edit Group Information - Name and Description Page

| Option | Description |
| --- | --- |
| Group Name | Name of the user group that is created manually. |
| Description | Description of the user group, indicating its purpose. |

4 **Assign Roles and Scope**

The Assign Roles and Scope option lets you select a role for the user group and then assign a scope for each role.

5 From the **Select Role** drop-down menu, select a role for the user group.

6 From the **Select Scope** drop-down menu, assign scope for the selected user group.

**Note** You can click the + sign to add multiple roles and then assign the required scope for each role.

**Assign Users**

The Assign Users option lets you add members to the group based on the membership criteria that you have defined.

7 Select the users who would be the members of this group.

8 Click **Save**.

## Export and Import of User Groups

You can export User Group configurations from one VMware Aria Operations and import into any VMware Aria Operations.

### Export User Groups

1 From the left menu click **Administration**, and then click **Access Control** tile.

2 In the **User Groups** tab, select the user groups to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.

The user groups `.json` file is downloaded to the default download location.

### Import User Groups

1 From the left menu click **Administration**, and then click **Access Control** tile.

2 In the **User Groups** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.

3 Click **Browse** and select the user groups `.json` file.

4 In case of a conflict, select either **Overwrite existing User Groups** or **Skip User Groups**.

5 Click **Import**.

### Important Points

- User group's scopes are matched in the target VMware Aria Operations, which means user group's scopes references are exported, then assigned to the user group while importing only if those scopes exist on the target VMware Aria Operations.
- User group's roles are matched in the target VMware Aria Operations, which means user group's role references are exported, then assigned to the user group while importing only if those roles exist on the target VMware Aria Operations.
- User group's members are matched in the target VMware Aria Operations, which means user group's member references are exported, then assigned to the user group while importing only if those users exist on the target VMware Aria Operations.
- Import of the external user groups fails if their authentication source is not found in the target VMware Aria Operations.
- The traversal specifications are not imported or exported in VMware Aria Operations.
- An error message is displayed with details for failed imports.
- User groups of type "Cloud Services Platform" can not be exported in VMware Aria Operations (SaaS).

## Access Control: Roles Tab

You can assign users-specific roles to perform actions and view features and objects in VMware Aria Operations. With role-based access, users can only perform the actions that their permissions allow.

### Where You Manage User Roles

1 To manage user roles, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Click the **Roles** tab.

To view the details of roles configured in VMware Aria Operations, click the role, the role details are displayed in the right-side panel. The role details include the permissions, user accounts, and user groups associated with the selected role.

You can view and edit details about a role, by selecting a role in the summary grid, and clicking the **Edit** icon in the Roles toolbar.

### Table 20-9. Access Control Roles Summary Grid

| Option | Description |
| --- | --- |
| Roles toolbar | To manage roles, use the toolbar icons.<br>■ Click the **Add** icon. to add a user role, provide the name and description for the role, and assign permissions in the **Role Information** page.<br>**Note** To assign all the available permissions to the selected role, select the **Select all Permissions** check box.<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>  ■ **Edit**. Edit the selected user role, and modify the details for the role in the **Edit Role** dialog box.<br>  ■ **Delete**. Delete a user role.<br>  ■ **Clone**. Clone the selected user role<br>  ■ **Export**. Export a user role.<br>■ Click the **Horizontal Ellipses** to perform any one of the following actions:<br>  ■ **Delete**. Click delete to Delete a role.<br>  ■ **Export**. Click export to export a role.<br>  ■ Click **Import** to import a role from the **Import Roles** page. You can import existing roles, overwrite existing roles, and skip roles. |
| Filters | Activates you to search the list of roles according to the following criteria:<br>■ Role Name<br>■ Role Description<br>■ Modified By |

You can view details for the user accounts and user groups associated with a selected role in the Details panes.

### Table 20-10. Access Control Roles Details Panes

| Option | Description |
| --- | --- |
| Role Name | Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the role, indicating its purpose. |
| Modified By | Displays the name of the user who last updated the role. |
| Last Modified | Displays the date and time when the role was last updated. |

Table 20-10. Access Control Roles Details Panes (continued)

| Option | Description |
| --- | --- |
| Permissions | Displays the permissions assigned to the role according to the categories: Administration, Automation Central, Configure, Data Sources, Environment, Home, Optimize, Plan, Troubleshoot, and Visualize. Expand the tree of each category to view all the assigned permissions.<br>You can edit the permissions assigned to the role by clicking the **Edit Role** icon.<br>■ Click the **Expand All** button to expand the trees of all the categories, and select the check boxes to apply permissions for the selected role. |
| User Accounts | The roles are assigned to the users. You can only view users who have been assigned roles directly, the users belonging to a user group are not displayed here. The information in this pane is based on the data entered when you created the user, or imported with the user.<br>■ First Name. Indicates the first name of each user who is assigned this role.<br>■ Last Name. Indicates the last name of each user who is assigned this role.<br>■ User name , without spaces, that will log in to VMware Aria Operations<br>■ Email. Indicates the email address for each user who is assigned this role. |
| User Groups | The roles are assigned to the selected user groups.<br>■ Group Name: Name of each group that is associated with the selected role.<br>■ Members: Number of members in each group. |

## Add or Edit Roles and Assign Permissions

You can add or edit user-specific roles to perform actions and view features and objects in VMware Aria Operations. With role-based access, you can assign roles to users and let them only perform the actions that their permissions allow.

### Where You Add or Edit User Roles

1 To manage user roles, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Select the **Roles** tab and click **Add**.

3 To edit a role, select the role and then click **Edit Role**. You can also click the **EDIT ROLE** button in the **Role Details** page and edit the role.

Table 20-11. Add or Edit Role Information Page

| Option | Description |
| --- | --- |
| Name | Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the role, indicating its purpose. |
| **Assign Permissions** | |
| **Note** You can set default permissions for the selected role using the **Apply Default Permissions** option. The Apply Default Option is available only when you edit an Out-of-the-box role. | |

Table 20-11. Add or Edit Role Information Page (continued)

| Option | Description |
| --- | --- |
| Select all Permissions | Click **Select all Permissions** to apply the permissions for the role you create or select. You can click the arrow next to the categories to set specific permissions to the roles you define. |
| Expand All | Click the **Expand All** button to expand the trees of all the categories and select the check boxes to apply or modify the permissions for the selected role. |
| Collapse All | Click the **Collapse All** button to Close the tree view and display only the high-level categories. |

4 Click **Save**.

### Export and Import of Roles

You can export User Roles configurations from one VMware Aria Operations and import into any VMware Aria Operations.

#### Export User Roles

1 From the left menu click **Administration**, and then click **Access Control** tile.

2 In the **Roles** tab, select the user roles to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.

The roles `.json` file is downloaded to the default download location.

#### Import Roles

Invalid privilege keys of a role are ignored during import.

1 From the left menu click **Administration**, and then click **Access Control** tile.

2 In the **Roles** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.

3 Click **Browse** and select the roles `.json` file.

4 In case of a conflict, select either **Overwrite existing Roles** or **Skip Roles**.

5 Click **Import**.

#### Important Points

- Invalid permissions are ignored during import based on target VMware Aria Operations.
- An error message is displayed with details for failed imports.

## Access Control: Scopes Tab

You can add, edit, or remove scope associated with users and groups in VMware Aria Operations. Scope provides you the ability to limit access of a user or a set of users to VMware Aria Operations. The All Objects scope is created by default. Assign the All Objects scope to the user to access all objects discovered by VMware Aria Operations.

### Where to Manage Scope

1 To manage scope, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Click the **Scopes** tab.

To view the details for a specific scope defined in VMware Aria Operations, click the scope. The scope details are displayed in the right-side panel. The scope details include the user accounts and user groups associated with the scope.

#### Table 20-12. Access Control Scope Summary Grid

| Option | Description |
|---|---|
| Scope Toolbar | To manage scopes, use the toolbar options.<br>■ Click **Add** to add the scope and provide the name and description for the scope, and assign members to objects in the **Scope Information** page.<br>Select **Object Hierarchies** and then select the object to assign members to. You can select one or more objects.<br>**Note** For vSphere Storage object hierarchy, select the **Propagate to children** checkbox to include all future discovered objects in the scope.<br>■ Click the **Vertical Ellipses** to perform any one of the following actions:<br>■ **Edit**. Edit a specific scope and modify the details for the scope in the **Edit Scope Information** page.<br>■ **Clone**. Clone a scope and enter a name and description for the cloned scope.<br>■ **Delete**. Delete a scope.<br>■ Click the **Horizontal Ellipses** and select **Delete**. Delete a scope. |
| Filters | Activates you to search the list of scopes according to the following criteria:<br>■ Scope Name<br>■ Description<br>■ Modified By |

#### Table 20-13. Access Control Scope Details Grid

| Option | Description |
|---|---|
| Scope Name | Displays the name of the selected scope. |
| Description | Displays the description of the scope, indicating its purpose. |
| Modified By | Indicates the last person to update the scope. |
| Last Modified | Indicates the time when the scope was last updated. |
| User Accounts | Displays the user accounts associated with the selected scope. |
| User Groups | Displays the user groups associated with the selected scope. |

### Modify Scopes and Assign Member Objects

You can add and modify the scope defined in VMware Aria Operations. Scope lets you create a package and associate the objects and object hierarchies to the package. You can select the

object hierarchy and assign the required objects to the selected object hierarchy. After you define a scope, you can associate the scope to user accounts and user groups.

### Where You Add or Edit Scope

1 To manage scope, from the left menu, click **Administration**, and then click the **Access Control** tile.

2 Click the **Scopes** tab.

3 In the **Scopes** tab, click **Add**.

4 To edit a scope, select the scope and then click **Edit Scope**. You can also click the **EDIT SCOPE** button in the **Scope Details** grid and edit the scope.

Table 20-14. Access Control Scopes Details Grid

| Option | Description |
| --- | --- |
| Scope Name | Name of the scope to apply to a specific user or a set of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the scope, indicating its purpose. |

5 To add or edit the scope, select the required **Object Hierarchy**, and then select the **Object** to be associated.

6 Click **Save**.

After you save the scope, you can assign the scope to user accounts and user groups.

## Audit Users and the Environment in VMware Aria Operations

At times, you might need to provide documentation as an evidence of the sequence of activities that took place in your VMware Aria Operations environment. Auditing allows you to view the users, objects, and information that is collected. To meet audit requirements, such as for business critical applications that contain sensitive data that must be protected, you can generate reports on the activities of your users, the privileges assigned to users to access objects, and the counts of objects and applications in your environment.

Auditing reports provide traceability of the objects and users in your environment.

### User Activity Audit

Run this report to understand the scope of user activities, such as logging in, actions on clusters and nodes, changes to system passwords, activating certificates, and logging out.

### User Permissions Audit

Generate this report to understand the scope of user accounts and their roles, access groups, and access privileges.

### System Audit

Run this report to understand the scale of your environment. This report displays the counts of configured and collecting objects, the types and counts of adapters, configured and collecting metrics, super metrics, applications, and existing virtual environment objects. This report can help you determine whether the number of objects in your environment exceeds a supported limit.

**VC Pricing Migration Audit**

Run this report to understand the details of VC pricing migration.

**Chargeback Migration Audit**

Run this report to understand the details of Chargeback migration.

## Reasons for Auditing Your Environment

Auditing in VMware Aria Operations helps data center administrators in the following types of situations.

- You must track each configuration change to an authenticated user who initiated the change or scheduled the job that performed the change. For example, after an adapter changes an object, which is associated with a specific object identifier at a specific time, the data center administrator can determine the principal identifier of the authenticated user who initiated the change.
- You must track who made changes to your data center during a specific range of time, to determine who changed what on a particular day. You can identify the principal identifiers of authenticated users who were logged in to VMware Aria Operations and running jobs, and determine who initiated the change.
- You must determine which objects were affected by a particular user during a time-specific range of time.
- You must correlate events that occurred in your data center, and view these events overlayed so that you can visualize relationships and the cause of the events. Events can include login attempts, system start up and shutdown, application failures, watchdog restarts, configuration changes of applications, changes to security policy, requests, responses, and status of success.
- You must validate that the components installed in your environment are running the latest version.

## User Activity Audit

The user activity report helps you understand the scope of user activities in your VMware Aria Operations instance, such as when users logged in, actions they took on clusters and nodes, changes they made to system passwords, when they activated certificates, and when they logged out.

### Where You Audit User Activity

To audit user activity, from the left menu, click **Administration**, and then click the **Audit** tile. The activities that users performed in the environment appear on the page.

Table 20-15. User Activity Audit Actions

| Option | Description |
| --- | --- |
| Download | Download the user activity audit information to a report in PDF or XLS format. |
| Date Range | Display the list of user activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

## User Permissions Audit

A user permissions audit report provides an overview of the local users in your VMware Aria Operations instance, and a list of groups to which each user belongs. This report helps you understand the scope of the user accounts and their roles, access groups, and access privileges in your environment.

The report displays the access group associated with each local user and the access privileges granted to the user in each access group. This report does not include vCenter Server users, roles, or privileges.

When a user is a member of a specific user group, the associated access group could provide the user with access to configuration, dashboards, and templates, or to specific navigation areas in the user interface such as Administration. The access rights associated with the access group include actions for each access group, such as the ability to add, edit, or delete dashboards, or to view, configure, or manage objects.

### Where You Audit User Permissions

1 To audit user permissions, from the left menu, click **Administration**, and then click the **Audit** tile.

2 Click the **User Permissions Audit** tab.

The permissions assigned to users, and their associated access groups and access privileges, appear on the page.

Table 20-16. User Permissions Audit Actions

| Option | Description |
| --- | --- |
| Download | Download the user permissions audit information to a report in PDF or XLS format. |

## System Audit for VMware Aria Operations

A system audit report provides an overview of the counts of objects, metrics, super metrics, applications, and custom groups in your VMware Aria Operations instance. This report can help you understand the scale of your environment.

The system audit report displays the types and number of objects that VMware Aria Operations manages. Reported objects include those that are configured and collecting data, the types of objects, object counts for adapters, the metrics that are configured and being collected, super metrics, VMware Aria Operations generated metrics, the number of applications used, and the number of custom groups.

You can use this report to help determine whether the number of objects in your environment exceeds a supported limit.

### Where You Audit the System

1 To audit the objects, metrics, applications, and custom groups in your environment, from the left menu, click **Administration**, and then click the **Audit** tile.

2 Click the **System Audit** tab.

The objects and their associated counts appear in the report.

Table 20-17. System Audit Actions

| Option | Description |
| --- | --- |
| Download | Download the system information to a report in PDF or XLS format. |

## VC Pricing Migration Audit

The vCenter Pricing Migration Audit provides details of the migration of pricing cards to VMware Aria Operations policies.

### Where You Audit vCenter Pricing Migration Components

1 To audit vCenter Pricing Migration components, from the left menu, click **Administration** and then click the **Audit** tile.

2 From the **Audit** page, click **Chargeback Audit tab > VC Pricing Migration Audit tab**.

Table 20-18. vCenter Pricing Migration Audit Actions

| Option | Description |
| --- | --- |
| Download | Download the vCenter Pricing migration information to a report in PDF or XLS format. |
| Date Range | Display the list of vCenter Pricing migration activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

**Note** For more information about vCenter pricing migration, see Migration of vCenter Pricing Cards to vCenter Pricing Policies

## Chargeback Migration Audit

The Chargeback Migration Audit provides details of the migration of VMware Chargeback configurations and content to VMware Aria Operations.

### Where You Audit Chargeback Migration Components

1 To audit Chargeback Migration components, from the left menu, click **Administration** and then click the **Audit** tile.

2 From the **Audit** page, click **Chargeback Audit tab** > **Chargeback Migration Audit tab**.

Table 20-19. Chargeback Migration Audit Actions

| Option | Description |
| --- | --- |
| Download | Download the Chargeback migration information to a report in PDF or XLS format. |
| Date Range | Display the list of Chargeback migration activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |

Table 20-19. Chargeback Migration Audit Actions (continued)

| Option | Description |
| --- | --- |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

**Note** For more information about Chargeback migration, see Migration of Chargeback Data to VMware Aria Operations in Case of an Upgrade .

# VMware Aria Operations Certificates

VMware Aria Operations includes a central page where you can review authentication certificate contents.

## How the Certificates Page Works

The Certificates page lets you examine certificate contents without the need to open the certificate outside of VMware Aria Operations.

## Where You Find Certificates

In the menu, click **Administration**, and then in the left pane, click **Management > Certificates**.

## Certificate Tabs

The certificate tab describes columns of exceptions tabs.

**Note** The CRL tab is activated only when you select the **Activate Standard Certificate Validation** under **Global Settings**.

Table 20-20. Certificate Tabs

| Tabs | Description |
| --- | --- |
| Exceptions | Lists the certificate that is accepted by the VMware Aria Operations administrator but is not certified by the Certificate Authority (CA). |
| CRL | A Certificate Revocation List (CRL) is a list of digital certificates that have been revoked by the issuing Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. Click the Add icon to upload the certificates. |

## Certificate Options

The options include a data grid for examining certificate contents.

### Table 20-21. Certificate Options

| Option | Description |
| --- | --- |
| Certificate Thumbprint | Unique alphanumeric string associated with the certificate |
| Issued By | Content associated with the issuer of the certificate, such as organization name and location |
| Issued To | Typically, content associated with the issuer, plus the certificate object Identifier (OID) |
| Expires | The date after which the certificate cannot be used for successful authentication |

## Importing CA Certificates

Certificate Authority (CA) or root certificates are used for establishing the outgoing connections from VMware Aria Operations. CA Certificates imported by the users will be used in the following VMware Aria Operations domains: Authentication Sources (Active Directory (AD), Open LDAP, VMware Identity Manager), Outbound Plugins, and Adapter Endpoint.

### Procedure

1. In the menu, click **Administration**, and then in the left pane click **Management > Certificates**.
2. Click **Import**.

   The Import CA Certificate(s) dialog box appears. You can only import certificates that are encoded in the PEM format.
3. Click **Browse**.
4. Locate the certificate `.pem` file and click **Open** to load the file in the Import CA Certificate(s) dialog box.

   The certificate information box appears with the certificate thumbprint, issued by, issued to and expiry date. For example, if you select a certificate that will expire in 10 days, you will receive a notification that the certificate is expiring soon.

   **Note** If a certificate is close to its expiry date, a corresponding notification is displayed on the **Home > Quick Start** page.
5. Click **Import**.
6. (Optional) Click the **Vertical Ellipsis** to delete a certificate.

## Removing an Adapter Certificate

If you want to delete an old or expired certificate associated with an adapter, perform the following steps:

### Procedure

1 Log in with the administrator user name and password.

2 In the menu, click **Administration**, and in the left pane click **Management > Certificates**.

3 In the certificate window, select the certificate that has to be removed.

4 Click **Delete** to remove the certificate.

5 If the certificate is being used by the adapter, then the following message comes up:

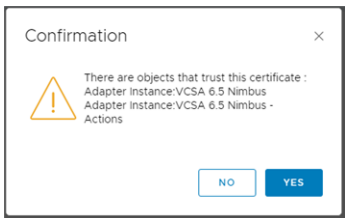


A certificate can be configured for one or more adapters if it is the same destination system.

6 If you delete a certificate which is already being used by another adapter, the adapter fails to connect or start. As a workaround, perform the following steps:

a On the left pane, click **Data Sources > Integrations**.

The Accounts tab appears.

b Select the particular integration and click the vertical ellipsis and then click **Edit**.

The Account Information page opens.

c Click **Validate Connection**.

d A prompt comes up asking the user to import the associated certificate. Click **OK**.

e Restart the adapter from the **Accounts** tab.

# Adding Physical Data Centers in VMware Aria Operations

In VMware Aria Operations, you can create and manage physical data centers and associate cloud accounts with them.

### Procedure

1 From the left menu, click **Administration** and then click the **Physical Data Centers** tile.

2 On the Physical Data Centers page, click **Add**.

Add Physical Data Centers page opens.

3 Enter a **Name** for the physical data center.

4 Select a location on the map.

   You can use the search bar to find a specific location. You can also zoom in or zoom out for a better view of the world map and find a suitable location.

5 Double-click on the selected location to create the data center.

   The location name and coordinates are reflected in the **Location** field.

6 From the **Associate Cloud Account(s) with this Physical Data Center** section, select the cloud account(s) you want to assign to this physical data center.

   **Note** The Associate Cloud Account(s) with this Physical Data Center section only lists the configured vCenter, VMware Cloud on Dell EMC, and VMware Cloud Foundation cloud accounts. You can assign multiple cloud accounts to a physical data center.

7 Click **Save**.

   The physical data center is created.

8 View the new physical data center on the Physical Data Center page.

| Option | Description |
| --- | --- |
| Name | Name of the physical data center you entered at the time of creation. |
| Number of associated cloud accounts | Displays the total number of cloud accounts associated with the physical data center. |
| Filter | Filters the list of physical data centers based on Name. |

9 (Optional) Click the vertical ellipsis and select **Edit** to edit the physical data center.

10 (Optional) Click the vertical ellipsis or the horizontal ellipsis and then select **Delete** to remove a physical data center.

# Modifying Global Settings

The global settings control the system settings for VMware Aria Operations, including data retention and system timeout settings. You can modify one or more of the settings to monitor your environment better. These settings affect all your users.

The global settings do not affect metric interactions, color indicators, or other object management behaviors. These behaviors are configured in your policies.

Settings related to managing objects with VMware Aria Operations are available on the **Inventory** page.

## Global Settings Best Practices

Most of the settings pertain to how long VMware Aria Operations retains collected and process data.

The default values are common retention periods. You might need to adjust the time periods based on your local policies or disk space.

## Access Global Settings

With global settings, you set times to delete objects, set timeouts, store historical data, use dynamic threshold and capacity calculations, and determine how vCenter Server users log in. For automated actions, you can select whether to allow actions to be triggered from alert recommendations automatically.

### Procedure

1. From the left menu, click **Administration**, and then click the **Global Settings** tile.
2. The global settings are categorised into Data Retention, Cost/Price, User Access, System Settings. Select one of the categories to edit the global settings.
3. If you cannot find a settings, type the name of the setting in the search bar.
4. When you edit a setting, **Save** and **Cancel** buttons appear. Make your selection to complete the editing process.

## List of Global Settings

The global settings determine how VMware Aria Operations retains data, keeps connection sessions open, and other settings. These are system settings that affect all users. Some of these settings are not editable. The global settings are grouped into four types. .

## Settings in the Data Retention Category

### Table 20-22. Global Setting Default Values and Descriptions

| Setting | Default Value | Description |
|---|---|---|
| Action History | 30 days | Number of days to retain the recent task data for actions.<br>The data is purged from the system after the specified number of days.<br>You can edit this setting from a minimum of one day to a maximum of 30 days in VMware Aria Operations. |
| Deleted Objects | 144 hours | This setting determines the number of hours for which objects in the inventory must be retained when:<br>■ the objects are deleted from an adapter data source or server before deleting them from VMware Aria Operations<br>■ A cloud account is deleted from VMware Aria Operations but the checkbox to delete related objects is not selected in the **Delete All Accounts** dialog box.<br>An object deleted from an adapter data source is identified by VMware Aria Operations as not existing and VMware Aria Operations can no longer collect data about the object. Whether VMware Aria Operations identifies deleted objects as not existing depends on the adapter. This feature is not implemented in some adapters.<br>This setting applies to objects deleted from the data source or server, not to any objects you delete from VMware Aria Operations on the Inventory page. |
| Deletion Scheduling Interval | 24 hours | Determines the frequency to schedule deletion of resources. This setting works with the Deleted Objects setting to remove objects that no longer exist in the environment. VMware Aria Operations transparently marks objects for removal that have not existed for the length of time specified under Deleted Objects. VMware Aria Operations then removes the marked objects at the frequency specified under Deletion Scheduling Interval. |

### Table 20-22. Global Setting Default Values and Descriptions (continued)

| Setting | Default Value | Description |
|---|---|---|
| Object History | 60 days | Number of days to retain the history of the object configuration, relationship, and property data.<br>**Note** The retention decision for each property value is based on the end date of that value instead of the start date. So any property value with an end date within the last 60 days interval (default setting) is retained. For example, if property value A was first published 1 year ago (which is long before the object history retention interval of 60 days) and property value B was published 3 days ago, and if no more points for that property exists, then the end date for the value A will be 3 days ago and for value B now. So both end dates will be within the recent 60 day limit and both values will be retained.<br>The configuration data is the collected data from the monitored objects on which the metrics are based. The collected data includes changes to the configuration of the object.<br>The data is purged from the system after the specified number of days.<br>You can edit this setting for a value from 10 days to 60 days in VMware Aria Operations. |
| Generated Reports Retention | 90 days | Number of days to retain generated reports. |
| Symptoms/Alerts | 30 days | Number of days to retain canceled alerts and symptoms.<br>The alerts and symptoms are either canceled by the system or by a user.<br>You cannot edit this setting in VMware Aria Operations. |
| Time Series Data Retention | 3 months | Number of months that you want to retain the collected and calculated metric data for the monitored objects. This setting is set to 3 months by default for 5 minutes interval data retention. |
| Additional Time Series Retention | 12 months | The number of months that the roll-up data extends beyond the regular period. The roll-up data is available starting from the end of the regular period and until the end of the roll-up data retention period.<br>You can edit this setting for a value of 0 months to 12 months in VMware Aria Operations. If you specify 0 as the value, then this will effectively disable the Additional Time Series Data Retention time and only data specified in Time Series Retention is stored. This setting ensures that after 3 months of normal retention for 5 minutes, the fourth month data is rolled up into a one hour roll up. |
| Near Real-Time Monitoring Data Retention | 3 days | The number of days to retain the near real-time data collected from the vCenter Server in VMware Aria Operations.<br>You cannot edit this setting in VMware Aria Operations. |

### Table 20-22. Global Setting Default Values and Descriptions (continued)

| Setting | Default Value | Description |
|---|---|---|
| Deleted Users | 100 days | You can specify the number of days to keep custom content created by a user who has been removed from VMware Aria Operations or by the automatic synchronization of LDAP. For example, the custom dashboards created by a user. |
| External Event Based Active Symptoms | disabled | The number of days to retain the external event-based active symptoms. |
| Maintain Relationship History | disabled | You can maintain a history of all the relationships of all the monitored objects in VMware Aria Operations<br>. |

## Settings in the Cost/Price Category

### Table 20-23. Global Setting Default Values and Descriptions

| Setting | Default Value | Description |
|---|---|---|
| Currency | | You can specify the currency unit that is used for all the cost calculations. You can select the type of currency from the list of currency types by clicking **Choose Currency**. From the **Set Currency**, select the required currency and confirm your action by clicking the check box, and set the currency. |
| Cost Calculation | activated | The host time at which cost calculations are run. |
| Cluster Utilization Ceiling Factor | 5 | Ceiling for Expected Utilization when running on Actual Utilization. |
| Tag Based Costing Metrics | Disabled | When activated, VMware Aria Operations additional cost metrics per tag. |
| Tag Based Pricing Metrics | Disabled | When activated, VMware Aria Operations additional pricing metrics per tag. |

## Settings in the User Access Category

### Table 20-24. Global Setting Default Values and Descriptions

| Setting | Entry | Description |
|---|---|---|
| Concurrent UI login sessions | activated | Allows concurrent UI login sessions per user. Once changed, this setting affects the subsequent login sessions. |

## Settings in the Reclamation Category

Settings in the reclamation category display information about powered -off VMs, idle VMs, snapshots, and orphaned disks.

### Table 20-25. Global Settings Default Values and Descriptions

| Setting | Entry | Description |
| --- | --- | --- |
| Include Powered off VMs for Reclamation | activated | VMs that have been continuously powered off during the defined time period.<br>The total storage capacity used is reclaimable. Total storage reclaimable cost is computed by multiplying storage rate with storage utilization. The direct cost of VM is also attributed. |
| Idle VM Management | | |
| Include Idle VMs for Reclamation | activated | You can configure the following parameters based on which VMware Aria Operations calculates idle VMs:<br>■ How many days the VM has been idle to be classified as a reclaimable idle VM.<br>■ The number of days before which the provisioned VMs are not considered in the classification of reclaimable idle. |
| VM Idleness Criteria | 100 MHz of CPU for 100% of Time | Defines the criteria to identify VMs that remain idle during the defined time period.<br>You can set the following values:<br>■ MHz of CPU consumption of VMs, and percentage of the time VMs have less than defined MHz of CPU within each day. |
| Include Snapshots for Reclamation | activated | VM snapshots that have existed for the entire defined time period.<br>Snapshots of a VM use storage space and such storage is reclaimable. The reclaimable cost is computed by multiplying the storage rate with the reclaimable storage value. |
| Orphaned Disk Management | | |

Table 20-25. Global Settings Default Values and Descriptions (continued)

| Setting | Entry | Description |
| --- | --- | --- |
| Include Orphaned Disks for Reclamation | disabled | Orphaned disks are VMDKs that are associated with a VM which are not in inventory, but still available in a data store. You can configure the minimum number of days for which VMDKs not related to any existing VM will be reported as orphaned and appear under the Orphaned Disks in the Reclaim page. |
| Orphaned Disks Collection Time | 8:00 pm | Host time to collect orphaned disks.<br>**Note** VMware Aria Operations checks for orphaned VMDKs in the vSphere Client instances based on the set time. The **Orphaned Disks Collection** and the **Cost Calculation** settings are interrelated. It is recommended to schedule orphaned disks collection before cost calculation so that the cost is calculated based on the recent list of reclaimable orphaned VMDKs. The default value for Orphaned Disks Collection is set to 8:00 PM, and the default value for Cost Calculation is set to 9:00 PM.<br>You can navigate to the **Cost/ Price** section under **Administration > Global Settings** and change the value of the **Cost Calculation** time. |

## Settings in the System Settings Category

### Table 20-26. Global Setting Default Values and Descriptions

| Setting | Entry | Description |
| --- | --- | --- |
| Session Timeout | 120 minutes | If your connection to VMware Aria Operations is idle for the specified amount of time, you are logged out of the application.<br>You must provide credentials to log back in. |
| Dynamic Threshold Calculation | activated | Determines whether to calculate normal levels of threshold violation for all objects.<br>If the setting is disabled, the following area of VMware Aria Operations does not work or are not displayed:<br>■ Alert symptom definitions based on dynamic thresholds will not work<br>■ Metric charts that display normal behavior are not present<br>Disable this setting only if you have no alternative options for managing resource constraints for your VMware Aria Operations system. |
| Customer Experience Improvement Program | activated | Determines whether to participate in the Customer Experience Improvement Program by having VMware Aria Operations send anonymous usage data to https://vmware.com. Deactivate this setting only if you want to opt out of CEIP. |
| Automated Actions | activated | Determines whether to allow VMware Aria Operations to automate actions. When an alert triggered, the alert provides recommendations for remediation. You can automate an action to remediate an alert when the recommendation is the first priority for that alert. You activate actionable alerts in your policies. |

Table 20-26. Global Setting Default Values and Descriptions (continued)

| Setting | Entry | Description |
| --- | --- | --- |
| Activate Standard Certificate Validation | disabled | This option activates certificate verification to Test Connection in the Create or Modify AI screen, using a standard verification flow.<br>The option checks CA authority.<br>■ Certificate Subject DN<br>■ Subject alternative name<br>■ Certificate validity period<br>■ Revocation list<br>This option also presents dialogs to user if one of those checks fail. It is up to the adapter implementation on how the adapter checks source certificate validity during a normal collection cycle. On a usual scenario, adapters just perform a thumb-print verification. However, in case this flag is activated, Test connection validates certificates in full scale and accepts certificates that are matching all criteria without any user dialogs. |
| Threshold For Adapters Certificate Expiration Alert | 5 for critical<br>14 for immediate<br>30 for warning | Set the number of days before which the system must raise certification expiration alerts.<br>To change the values of Critical, Immediate and Warning alerts, click the corresponding icon and move them along the slider. Alternatively, set the values manually.<br>Click **Save** after you make your changes.<br>**Note** For critical alerts, an alert banner is displayed under **Home > Quick Start** on the day of certification expiry. For example, if the critical alert is set to 10, then on the 10th day, an alert banner is displayed in the Home page. |

# Managing Content

As a VMware Aria Operations administrator, you can migrate all your content from VMware Aria Operations to VMware Aria Operations. You can take regular backups of your custom and the out-of-the-box content to manage your operational or regulatory needs. If there is a wrong edit or if the need to recover data arises, then you can use the recent backup to restore your content.

**Note** Any user with the "Content Management" permission can export the content. However, only a admin user or the user assigned to the Administrator role has the privilege to export all the content, including the content owned by other users, for example, custom dashboards.

## Creating a Backup

You can create regular backups of your custom and the out-of-the-box content in VMware Aria Operations. You can use this backup to restore your content or export the content while setting up another environment.

You can select the content or the configuration type that you want to export. Some content types have dependencies on other types. The dependencies, when not selected, will not be exported. Select all the required types to ensure that are no missing items when importing.

### Procedure

1. From the left menu, click **Administration**, and then in the right pane, click **Content Management**.
2. In the **Export** tab, select the content and configurations that you want to export. You can take a backup of the following content types and configurations available in VMware Aria Operations.

   **Note** Click **Include out-of-the-box content** under the Content section to include the out-of-the-box content while exporting. Click **Select All** under both Content and Configuration sections to select all the content and configurations respectively, while exporting.

   | Content | |
   |---|---|
   | ■ Dashboards | ■ Recommendations |
   | ■ Views | ■ Notifications |
   | ■ Reports | ■ Payload Templates |
   | ■ Report Schedules | ■ Policies |
   | ■ Configuration Files | ■ Custom Groups |
   | ■ Alert Definitions | ■ Custom Metric Groups |
   | ■ Symptom Definitions | ■ Super Metrics |
   | | ■ Compliance Scorecards |

| Configuration | |
|---|---|
| ■ User Groups<br>■ User Roles<br>■ Integration Accounts<br>■ Outbound Settings<br>■ Cost Drivers<br>■ Custom Services | ■ Service Based Applications<br>■ Rule Based Applications<br>■ Application Definition Assignments<br>■ Custom Profiles<br>■ Global Settings |

**Note** VMware Aria Operations SaaS does not support the export/import of Authentication Sources, Users, and Http Proxies.

3 Click **Export** to create a backup.

**Note** For configurations such as Integration Accounts, Http Proxies, Outbound Settings, Users, and Authentication Sources that have sensitive information, you must set up a new password to export data. The password should be at least 14 characters long.

The system compresses the content into one ZIP file. Once the export is complete, the Download ZIP file link is available in the Export tab.

4 Click the **Download ZIP file** link to download the backup content.

You can use the downloaded content to restore your content or export it to a different setup.

**Note** The ZIP file generated by a certain user will not overwrite that of another user, that is, if User A generates the ZIP file after User B, the latter's ZIP file will not be overwritten.

## Importing Content

You can take regular backups of your custom and the out-of-the-box content and import it to a different environment.

### Prerequisites

- Ensure that you have downloaded the backup ZIP file. For details, see Creating a Backup.

**Note** Items owned by a particular user (for example, dashboards and report schedules), except those owned by the admin user, will be skipped while importing the content.

**Note** Only one export or import operation can take place at a time.

### Procedure

1 From the left menu, click **Administration**, and then in the right pane, click **Content Management**.

2 Click the **Import** tab and then, click **Browse** to select the downloaded ZIP file with the exported content.

   The data included in the ZIP file is displayed in the 'Data Available' table.

3 If there is a conflict while importing the content, you can select to either **Overwrite existing content** or **Skip item(s)**.

   The import report with the timestamp is displayed after the import operation is complete. You can view this information under the **Results** section in the same page.

4 For content types with sensitive information, enter the password that you had set while exporting the content.

5 Click **Import**.

   After the import is completed, the content is available in the destination setup.

**Note** VMware Aria Operations SaaS does not support the export/import of Authentication Sources, Users, and Http Proxies.

## User Scenario: Migrating Content from VMware Aria Operations On-prem to VMware Aria Operations SaaS

As a VMware Aria Operations administrator, you can migrate your content from VMware Aria Operations to VMware Aria Operations SaaS. Migrating to VMware Aria Operations SaaS allows you to access content anywhere and anytime, and provides automatic upgrades to ensure you have the latest version and a secure one-way connection using Cloud Proxies.

While migrating from VMware Aria Operations on-prem to SaaS, you can migrate your custom as well as historic data.

### Migrating Custom Content

1 Log in to the VMware Aria Operations instance from which you want to export the content.

2 From the left menu, click **Administration**, and then in the right pane, click **Content Management**.

3 In the **Export** tab, select the content and configurations that you want to export, and click **Export**.

   For details, see Creating a Backup.

   The system compresses the content into one ZIP file.

**Note** For configurations such as Integration Accounts, Http Proxies, Outbound Settings, Users, and Authentication Sources that have sensitive information, you must set up a new password to export data. The password should be at least 14 characters long.

**Note** VMware Aria Operations SaaS does not support the export/import of Authentication Sources, Users, and Http Proxies.

4 Click the **Download ZIP file** link to download the content bundle.

   You can use the downloaded content to restore your content or export it to a different setup.

5 Log in to VMware Aria Operations SaaS.

6 From the left menu, click **Administration**, and then in the right pane, click **Content Management**.

7 Click the **Import** tab and then, click **Browse** to select the downloaded ZIP file with the exported content.

8 If there is a conflict while importing the content, you can select to either **Overwrite existing content** or **Skip item(s)**.

9 For content types with sensitive information, enter the password that you had set while exporting the content.

10 Click **Import**.

   After the import is completed, the content is available in VMware Aria Operations SaaS. The import report with the timestamp is displayed after the import operation is complete. You can view this information under the **Results** section in the **Import** page.

**Migrating Historic Data**

You have the following options if you want to migrate your historic data from VMware Aria Operations to VMware Aria Operations SaaS.

- If you want to collect the data for up to 30 days, then request to run SaaS along with On-prem free of charge for 30 days.
- If you want to retain historical data for a longer duration (more than 30 days), then contact your TAM or the Support team.

For details on Data Retention, refer to the 'List of Global Settings' topic in the *Configuring VMware Aria Operations Guide*.

## Best Practices for Migrating Content

Follow the below practices to ensure that your content is successfully migrated.

- Use the *admin* user account or a user assigned to the Administrator role to export all of the content, including other users' custom content, such as dashboards and report schedules.
- Before importing the content, ensure that the Management Packs to which the content is related is installed on the destination setup.
- Use a user from the CSP Admin group to import all content.

# Managing Orphaned and Unassigned Content

Dashboards, report schedules, and credentials created by deleted users or that are unassigned in VMware Aria Operations are stored in the Orphaned and Unassigned Content page. As an

administrator, you are the new owner of these dashboards, reports schedules, and credentials. You can transfer the ownership of dashboards and report schedules and assign the credentials to new users or to yourself.

## From Where You Can Transfer Ownership of Dashboards, Report Schedules, and Credentials

In the menu, click **Administration**. From the left pane, select **Management > Orphaned and Unassigned**.

## Orphaned and Unassigned Page

You can view a list of all deleted users or unassigned users from the **Deleted Users** panel in the left pane of the **Orphaned and Unassigned** page. Based on your selection in the **Deleted Users** panel, the dashboards, report schedules, and credentials for the deleted user are displayed under the **Dashboard**, **Report Schedules**, **Credentials** tabs in the **Orphaned and Unassigned** page.

As an admin user, you can take ownership, assign ownership, or discard orphaned dashboards, report schedules and credentials, from the **Actions** menu in the **Dashboards**, **Report Schedules**, and **Credentials** tabs. Enter the name or part of the name of a dashboard, report schedule, or credential in the **Filter** option and click **Enter**. The relevant dashboard, report schedule, or credential is displayed.

Table 20-27. Actions Menu Options

| Actions | Options |
| --- | --- |
| Take Ownership | You can take ownership of the selected dashboard, report schedule, or credential. |
| Assign Ownership | You can assign a new owner for the selected dashboards or report schedules. You can select a target user from the **Transfer Dashboards/Report Schedule/Credential** dialog box. |
| Discard | You can permanently delete the dashboard, report schedule, and unassigned credentials that are not in use in any adapter instance. You cannot delete an active credential that is being used by an adapter instance even if it is unassigned. For more information on credentials, see Configuring Credentials in Integrations. |

# VMware Aria Operations Dynamic Thresholds

A threshold marks the boundary between normal and abnormal behavior for a metric. In addition to fixed thresholds, VMware Aria Operations supports dynamic thresholds for a metric, calculated based on historical and incoming data.

## How Dynamic Thresholds Work

By default, dynamic thresholds are refreshed on a regular schedule, but you can recalculate dynamic thresholds outside of the schedule if you want to capture the most recent data.

## Where You Find Dynamic Thresholds

In the menu, click **Administration**, and then select **Global Settings > System Settings > Dynamic Thresholds**.

## Dynamic Threshold Options

The dynamic threshold feature includes options to activate or deactivate the calculation process and to review associated values.

Table 20-28. Dynamic Threshold Options

| Option | Description |
| --- | --- |
| Activate | The dynamic threshold calculation process ia activated by default. |
| Deactivate | Stop the dynamic threshold calculation currently in progress. |
| Calculation progress | Percentage completion of the current dynamic threshold calculation. |
| Host time to calculate normal behavior of all objects | Timestamps and metric counts associated with the last dynamic threshold calculation, and the time for the next scheduled calculation. |

# Customizing Icons

Every object or adapter in your environment has an icon representation. You can customize how the icon appears.

VMware Aria Operations assigns a default icon to each object type and adapter type . Taken collectively, object types and adapter types are known as objects in your environment. Icons represent objects in the UI and help you to identify the type of object. For example, in the Topology Graph widget on a dashboard, labeled icons show how objects are connected to one other. You can quickly identify the type of object from the icon.

If you want to differentiate objects, you can change the icon. For example, a virtual machine icon is generic. If you want to pictorially distinguish the data that a vSphere virtual machine provides from the data that a Hypervisor virtual machine provides, you can assign a different icon to each.

## Customize an Object Type Icon

You can use the default icons that VMware Aria Operations provides, or you can upload your own graphics file for an object type. When you change an icon, your changes take effect for all users.

### Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

### Procedure

1 From the left menu, click **Administration**, and then click the **Icons** tile.

2 Click the **Object Type Icons** tab.

3 Assign the Object Type icon.

   a Select the object type in the list with the icon to change.

   By default, object types for all adapter types are listed. To limit the selection to the object types that are valid for a single adapter type, select the adapter type from the drop-down menu.

   b Click the **Upload** icon.

   c Browse to and select the file to use and click **Done**.

4 (Optional) To return to the default icon, select the object type and click the **Assign Default Icons** icon.

   The original default icon appears.

## Object Type Icons Tab

VMware Aria Operations obtains data from different sources. Data sources are classified by the type of object or object type. In UI locations where metric data appears for objects, VMware Aria Operations includes an icon to show the object type. To graphically distinguish the different types of objects, you can customize the icon.

### Where You Customize Object Type Icons

From the left menu, click **Administration**, and then click the **Icons** tile. Click the **Object Type Icons** tab.

Table 20-29. Object Type Icons Options

| Option | Description |
| --- | --- |
| Adapter Type | Icons for all adapters are listed by default. To list a subset of the object types that are valid for one type of adapter, select the adapter type. |
| Toolbar options | Manages the selected icon.<br>■ **Upload** uploads a PNG file to uniquely identify the object type.<br>■ **Assign Default icons** returns the selection to the original icon. |
| Search | Search for objects with a particular name to narrow the selection of object types displayed. |

Table 20-29. Object Type Icons Options (continued)

| Option | Description |
| --- | --- |
| Object Type | Name of the type of object. |
| Icon | Pictorial representation of the type of object. |

## Customize an Adapter Type Icon

You can use the default icons that VMware Aria Operations provides, or you can upload your own graphics file for an adapter type. When you change an icon, your changes take effect for all users.

### Prerequisites

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

### Procedure

1 From the left menu, click **Administration**, and then click the **Icons** tile.

2 Click the **Adapter Type Icons** tab.

3 Assign the Adapter Type icon.

   a Select the adapter type in the list with the icon to change.

   b Click the **Upload** icon.

   c Browse to and select the file to use and click **Done**.

4 (Optional) To return to the default icon, select the adapter type and click the **Assign Default Icons** icon.

   The original default icon appears.

## Adapter Type Icons Tab

Adapters collect and provide data to VMware Aria Operations. Adapters are classified by the type of adapter or adapter kind. To graphically distinguish the different types of adapters, you can customize the icon.

### Where You Customize Adapter Type Icons

From the left menu, click **Administration**, and then click the **Icons** tile. Click the **Adapter Type Icons** tab.

### Table 20-30. Adapter Type Icons Options

| Option | Description |
| --- | --- |
| Toolbar options | Manages the selected icon.<br>■ **Upload** uploads a PNG file to uniquely identify the adapter type.<br>■ **Assign Default icons** returns the selection to the original icon. |
| Name | Name of the type of adapter. |
| Icon | Pictorial representation of the type of adapter. |