



vRealize Operations Best Practices

Supplemental Guide
Version 8.1

APRIL 2020
VERSION 1.6



Table of Contents

Introduction.....	5
Best Practices Concepts.....	5
Areas of Best Practices.....	5
Platform Best Practices.....	6
Sizing.....	6
Storage Approach.....	6
General Guidelines.....	6
Architecture.....	7
High Availability (HA).....	7
Continuous Availability (CA).....	9
Remote Collectors.....	10
Load Balancers.....	11
Deployment.....	11
Upgrade.....	11
Cluster.....	13
Backup & Restore.....	14
Backup.....	14
Restore.....	15
Disaster Recovery.....	16
Self-Monitoring.....	16
API and Integration.....	16
End Point Operations.....	17
Sizing.....	17
Deployment.....	17
Content Best Practices.....	17
Metrics.....	17
Alerts & Symptoms.....	18
Review Out-Of-The-Box (OOTB).....	18
Dashboards.....	19
Views and Reports.....	20
Views.....	20
Reports.....	20
Super Metrics.....	20
Policies.....	21
Account and Roles.....	21
Maintenance Schedule.....	21



Grouping.....	21
Work Load Placement (WLP).....	22
Predictive Distributed Resource Scheduler (pDRS).....	22
Operations Best Practices	23
SDDC Monitoring.....	23
Additional Best Practices.....	24
Documentation Links	24
Additional Links	24



Revision History

DATE	VERSION	DESCRIPTION
April	1.6	Updates with vRealize Operations 8.1
October 2019	1.5	Updates with vRealize Operations 8.0
September 2018	1.4	Updates with vRealize Operations Manager 7.0
April 2018	1.3	Updates
March 2018	1.2	Updates
March 2017	1.1	Updates
February 2017	1.0	Initial version



Introduction

This document describes the best practices and recommendations for VMware vRealize Operations. This document is not a deployment guide, but a guide that supplements the vRealize Operations installation and configuration documentation available in the [vRealize Operations Documentation Center](#).

There are additional best practices outlined in the product documentation; therefore, existing information may not be displayed in this document. Please refer to the product documentation for additional best practices.

This information is for the following products and versions.

PRODUCT	VERSION	DOCUMENTATION
vRealize Operations	8.0, 8.1	vRealize Operations Documentation Center
vRealize Operations Manager	6.6.1, 6.7, 7.0, 7.5	

Best Practices Concepts

This document provides information based on development, test, field, and customer interaction. Each environment is unique and the way vRealize Operations is used may vary; hence, this information provides general principles or techniques that, when applied, will produce results that are superior to those achieved by other means or by standard use.

In certain cases, it may not be practical to apply best practice methods nor is there a requirement to use all available best practices. The areas of best practice should be applied appropriately based on the environment, the user(s) and the way that vRealize Operations is being utilized.

Following are the advantages of applying best practices with vRealize Operations:

- Proven Results
- Enhanced Performance
- Consistency
- Improved Usability
- Greater Stability

Areas of Best Practices

Applying best practices for vRealize Operations focuses on three key areas:

- Platform (product)
The technical portion of the product, which includes architecture and sizing, deployment, cluster, high availability, continuous availability, remote collector, API, interoperability and integration, backup & restore, and disaster recovery.
- Content (product)
The functional part of the product, meaning the content that “sits on” the platform. Content includes policies, dashboards, alerts, reports, super metrics, groups, and actions.
- Operations
The how you use the product in your operations. This includes working with other roles in Operations (e.g. NOC, Storage, and Management). Examples of Operations are processes, roles, groups, tenants.



Platform Best Practices

The Platform is the technical portion of the product. The best practices applied here help to provide the most optimal options for the platform to provide the most stable running environment for daily operational use. Before deployment of vRealize Operations, the first requirement is to size the environment. This section will cover sizing and recommendations post deployment of the product. Additional best practices are included for administration tasks such as backup & restore or disaster recovery. These best practices will help to ensure that the platform, vRealize Operations, is properly sized to run and handle the monitoring load efficiently.

Sizing

Storage Approach

- Size the deployment with twelve to eighteen months of infrastructure growth

When an environment outgrows the original deployment size, performance degradation and usability problems may become present. Planning for infrastructure growth of twelve to eighteen months will allow the system to continue functioning without the need to immediately resize or scale out the deployment. For example, if you anticipate a 10% annual growth, increase the initial sizing by 15% to obtain an eighteen-month sizing recommendation.

- Review the sizing guidelines frequently and often during the growth of the environment (resizing)

To keep the environment running with optimal parameters, it is important to review the sizing guidelines and resize the deployment as necessary. Even with expected growth, frequently reviewing the sizing guidelines regularly will proactively prevent performance and usability problems typically associated with undersized environments.

- [vRealize Operations Sizing Guidelines](#)
- [vRealize Operations Sizing](#)

General Guidelines

- Validate the sizing guidelines with your actual environment

The sizing guidelines provide general estimates and requires confirmation with the actual environment. For example, the data entered in the sizing calculator may yield additional objects not captured in the actual environment or vice versa.

- Calculate only the components which will be monitored

It is possible that some components do not need to be monitored; therefore, exclude those components in the sizing calculations.

- Size the Cluster

There are multiple sizes for analytics nodes, extra small, small, medium, large and extra-large. It is best to use the least number of nodes when possible. For example, if the recommendation is to have 10 large nodes or 4 extra-large nodes, use the lesser number of nodes to minimize the amount of communication across more nodes.

- Size the Remote Collectors

There are two sizes for default remote collectors, standard and large. Use the appropriately sized remote collector based on how much data will be collected. If necessary, use multiple remote collectors to ensure the proper sizing of remote collectors for the environment.



- Adjust the time series data retention to keep data for a timeline which data is critically needed
The default setting for data retention is six months. If only three months of data is required, lower the default value. Understand what you gain when using longer data retention periods. It may not necessarily help having longer retention periods. Depending on your operational needs, configure the retention period to suit your requirements.
- Consider additional storage and IO requirements for longer data retention
For those times when longer data retention periods are required, consider additional storage and increased IO requirements. For example, retail businesses may need to keep more than one year of data to account for seasonal peaks.
- Leverage the additional time series retention to keep longer historical data while minimizing the time series data retention period.
The default setting for additional time series retention is thirty-six months. Adjust the default value to a necessary period and lower the time series data retention period to save on the amount of data being retained.
- Add VMDK instead of extending
Increase storage by adding a VMDK to minimize impact to existing
- Only install Management Packs that are available on the [VMware Solution Exchange](#)
There are several management packs available for vRealize Operations; however, only management packs certified and supported by VMware are available on the [VMware Solution Exchange](#).
- Confirm VMware product compatibility support before installing or upgrading components
Refer to the [VMware Product Interoperability Matrix](#) for all VMware product and management packs supported with vRealize Operations
- Validate management packs created by VMware partners are supported
The 3rd party authored Management Packs that are supported are listed in the [VMware Compatibility Guide](#).
- Before adding Management Packs, verify the additional metrics they will provide
The metric names may look correct but may not always mean what you really want. Be sure that the metrics from added management packs are what you really need and used properly; otherwise, disable the unnecessary metrics.

Architecture

High Availability (HA)

- Understand what High Availability (HA) provides (or does not provide) before enabling (or disabling)
Enabling HA may require double the resources, as data is stored redundantly in two nodes as opposed to only on one node when HA is disabled. Since the data is being stored in two nodes, this limits the total capacity by approximately 50%.
For example, a deployment of **extra-large nodes** will support the maximum number of objects:

vRealize Operations	Number of XL nodes	HA Disabled	HA Enabled
vRealize Operations 8.1	8	320,000	180,000

vRealize Operations 8.0	6	240,000	120,000
vRealize Operations Manager 7.5, 7.0	6	240,000	120,000
vRealize Operations Manager 6.7, 6.6.1	6	180,000	90,000

- HA will allow losing only one data node for the cluster to remain functional
It is important to understand and weigh the cost of the extra resources to the benefits that HA provides.
- Enable HA only after all nodes in the cluster have been added and are online
Add all data nodes to the cluster before enabling HA. On new deployments, add data nodes to build the cluster to fit the appropriate sizing and then enable HA. If adding new data nodes to an existing cluster, add as many data nodes as necessary, then enable HA. The goal is to minimize the number of times for enabling HA; the process to enable HA can be very disruptive so perform only when necessary.
- Deploy all analytics nodes for a single vRealize Operations cluster in the same data center
It is required to have all analytics nodes in the same data center to ensure latency requirements are consistently met for providing efficient cross node communication and optimal cluster performance.
- Deploy analytics cluster nodes on separate hosts for redundancy and isolation
If possible, establish a 1:1 mapping for nodes to hosts. This will protect the cluster if one host goes down, then only one node is lost, and the cluster remains functional. If it is not possible to establish a 1:1 mapping for nodes to host, make sure to separate the master node and master replica node on different hosts. This will safeguard the cluster if one of these hosts were to go down.
- Use anti-affinity rules that keep nodes on specific hosts in the vSphere cluster
To keep nodes separately on different hosts, use anti-affinity rules to prevent grouping of nodes on specific hosts. The idea is to prevent multiple nodes from going down if hosted on one node.
- Name nodes independent of role
Roles may change for nodes so statically naming a node a specific name may be confusing. For example, a node named 'Master' may no longer be the actual master node after promoting the replica node. This will avoid user confusion associated with poor naming convention.
- HA is not a substitute for a backup and recovery (B&R) plan
HA allows the cluster to remain functional only when one node is lost so a separate B&R solution should be used. See [vRealize Suite Documentation](#) for supported backup utilities and procedures.
- HA is not a Disaster Recovery (DR) strategy
HA for vRealize Operations is not a disaster recovery mechanism so a separate DR solution must be used. See [vRealize Suite Documentation](#). HA will allow the cluster to continue running if either the master node, the replica node or one data node fails. The entire cluster does not recover if multiple nodes fail at the same time.
- Hosts need to be on the same storage
For performance and consistency, use of the same storage is required.



Continuous Availability (CA)

- Understand what Continuous Availability (CA) provides (or does not provide) before enabling (or disabling)

Like HA, enabling CA may require double the resources, as data is stored redundantly in node pairs as opposed to only on one node when CA is disabled. Since the data is being stored in two nodes, this limits the total capacity by 50%.

For example, a deployment of **extra-large nodes** will support the maximum number of objects:

vRealize Operations	Number of XL nodes	CA Disabled	CA Enabled
vRealize Operations 8.1	8	320,000	200,000 (with 10 XL nodes)
vRealize Operations 8.0	6	240,000	160,000 (with 8 XL nodes)

- Deploy the witness node prior to enabling CA
The witness node must be deployed and added to the cluster in order to enable CA.
- Deploy the witness node in a separate datacenter
The witness node serves as a tiebreaker when a decision must be made regarding availability of vRealize Operations when the network connection between the two fault domains is lost. Keeping the witness node separate will ensure cluster availability if one of the datacenters are lost.
- Please make sure that the witness node has a reliable connection to both fault domains
The latency between witness node and fault domains should be as good as between the fault domains and it should be the same for both fault domains.
- CA must have an even number of analytics nodes before enabling CA
If the current cluster size consists of an odd number of analytics nodes, deploy one additional analytics node and add to the cluster. The added node must be the same version and size of the existing analytics nodes.
- Deploy fault domains into the highest object level as possible
Having fault domains separated into the highest object level in order of datacenters, then clusters, then hosts will ensure the highest level of availability during failures.
- CA will allow losing one fault domain for the cluster to remain functional
It is important to understand and weigh the cost of the extra resources, and placement of fault domains, to the benefits that CA provides.
- Enable CA only after all nodes in the cluster have been added and are online
Add all even number of data nodes and witness node to the cluster before enabling CA. On new deployments, add data nodes to build the cluster to fit the appropriate sizing and then enable CA. If adding new data nodes to an existing cluster, add as many even numbered of data nodes as necessary, then enable HA. The goal is to minimize the number of times for enabling CA; the process to enable CA can be very disruptive so perform only when necessary.
- Deploy all analytics nodes in the same data center for each fault domain
It is required to have all analytics nodes in the same data center for each fault domain, to ensure latency requirements are consistently met for providing efficient cross node communication and optimal cluster performance.

- Deploy analytics cluster nodes on separate hosts in each fault domain
If possible, establish a 1:1 mapping for nodes to hosts. This will minimize the impact to the fault domain if one host goes down.
- Use anti-affinity rules that keep nodes on specific hosts in the vSphere cluster
To keep nodes separately on different hosts, use anti-affinity rules to prevent grouping of nodes on specific hosts. The idea is to prevent multiple nodes from going down if hosted on one node.
- Name nodes independent of role
Roles may change for nodes so statically naming a node a specific name may be confusing. For example, a node named 'Master' may no longer be the actual master node after promoting the replica node. This will avoid user confusion associated with poor naming convention.
- CA is not a substitute for a backup and recovery (B&R) plan
CA allows the cluster to remain functional without data loss while at least one node from all node pairs is available so a separate B&R solution should be used. See [vRealize Suite Documentation](#) for supported backup utilities and procedures.
- CA is not a Disaster Recovery (DR) strategy
CA for vRealize Operations is not a disaster recovery mechanism so a separate DR solution must be used. See [vRealize Suite Documentation](#). CA allows the cluster to be stretched across two fault domains, with the ability to experience up to one fault domain failure and to recover without causing cluster downtime. The entire cluster does not recover if multiple node pairs, across fault domains, fail at the same time.
- Hosts need to be on the same storage in each fault domain
For performance and consistency, use of the same storage is required.

Remote Collectors

- Consider using Remote Collectors for local collections with larger vCenter servers (>7K objects)
Using remote collectors will help to reduce bandwidth across data centers and reduce the load on the vRealize Operations analytics cluster.
- Create collector groups when using multiple Remote Collectors
When utilizing multiple remote collectors for one vCenter server, create a collector group to provide collector high availability and redundancy. Collector groups can be configured to fault domains when CA is enabled.
- Deploy or update Remote Collectors to the same version of the Analytics nodes
Do not utilize mixed versions of Remote Collectors and Analytics nodes. Not only is a cluster running mixed versions unsupported, it may exhibit potential problems.
- Use Remote Collectors when using Management Packs or End Point Operations (EPOps) agents
Use remote collectors to isolate collection from Management Packs or End Point Operations agents to reduce the load on the vRealize Operations analytics cluster.
- Size Remote Collectors based on the number of collecting objects/metrics
Size remote collectors using the default sizing of standard and large nodes to accommodate the number of objects and metrics, which it will be collecting.
- Remote Collectors are recommended, but not required, to be included in the backup strategy
Include all remote collectors when taking a backup to restore the entire cluster health.



Load Balancers

- Review the latest API updates to use for node status
Starting with vRealize Operations 8.0, the node status API has been updated to use an optional set of services to get the aggregated statuses of the node. See [vRealize Operations Manager Load Balancing](#) for latest information.
- Use load balancers to provide a single UI entry for users
Use of a load balancer to provide multiple users a single URL for accessing the vRealize Operations cluster alleviates the need for users to remember logging into separate node names and accessing specific nodes.
- Use load balancers to provide high availability for remote collectors with End Point Operations agents
Use a load balance to group multiple remote collectors when using End Point Operations agents to provide high availability and redundancy.

Deployment

- Deploy vRealize Operations to a supported infrastructure
Ensure you are deploying vRealize Operations to a supported infrastructure as earlier versions may no longer be supported. Refer to the [VMware Product Interoperability Matrices](#) for platforms supported with vRealize Operations
- Do not modify or install third party applications on the appliance
When using the virtual appliance, installation or modifications of third-party applications is unsupported and may cause problems to vRealize Operations.
- Deploy the VA with FQDN
Register a fully qualified domain name for the vRealize Operations node. Simply using hostname may not properly resolve and may experience communication problems with the node.
- Use Thick Provisioning Eager Zeroed
When deploying nodes, set disk provisioning to “Thick Provision Eager Zeroed” for most optimum performance.
- When deploying Medium size nodes, increase the VM hardware level
The default hardware is set to “7” and limits the number of vCPUs per node. To increase the number of vCPUs when scaling a medium node to a large node, as example, the HW level must be set to a higher value.
- Leverage Remote Collectors
Use remote collectors where possible to navigate firewalls, reduce bandwidth across data centers, connect to remote data sources, or reduce the load on the vRealize Operations analytics cluster.

Upgrade

- Run the appropriate versioned pre-upgrade assessment tool on your current vRealize Operations before performing the upgrade to view the possible impact of your custom content to plan appropriate maintenance efforts for adjusting impacted custom content.
See [Using the Upgrade Assessment Tool for vRealize Operations Manager 8.1](#) and [vRealize Operations Upgrade Center](#) for latest information.



- Verify existing functionality before upgrading

Ensure the environment is fully functional before starting an upgrade. It is recommended to make a list of what works (or does not work) to confirm the same functionality post upgrade.
- Backup customized content before upgrade

Customized content should be backed up and saved for any potential overwrites or losses during upgrade.
- Snapshot VMs with the cluster offline before upgrading

After verifying functionality and backing up customized content, snapshot all the analytics VMs within the cluster for failsafe in event of an upgrade failure.
- Check interoperability of management packs before upgrade

It may be possible that some management packs will not be supported in the new product version and render the management pack inoperable. Before encountering this situation, confirm interoperability of management packs with the new product version.

See [VMware Product Interoperability Matrix](#) and e [VMware Compatibility Guide](#) for supported management pack versions.
- Perform the upgrade outside of DT / QIC / Costing or Backup processing

Perform the upgrade of the vRealize Operations cluster outside of dynamic threshold or capacity calculations or costing or during backups to avoid capturing high stress states.
- Setup blackout for maintenance to avoid false alerts

When performing maintenance, such as an upgrade or resizing the cluster, schedule a maintenance window to account for the performed activity to avoid receiving false alerts and notifications.
- Examine the recommendations from the validation checks before performing the upgrade

There is a pre-check upgrade validation script that runs before performing the actual upgrade. Address any failures and warnings before continuing to upgrade or the upgrade may fail.
- Enable the option to reset Default Content

Select the option to reset default content and bring in new content. This will overwrite existing content to a newer version provided by the update. User modifications to DEFAULT Alert Definitions, Symptoms, Recommendations, Policy Definitions, Views, Dashboards, Widgets and Reports will be overwritten; therefore, clone or backup the content before you proceed.
- Upgrade the OS PAK prior to upgrading the virtual appliance (VA) PAK for vRealize Operations 7.5 and lower.

To ensure a solid base OS before upgrading vRealize Operations, upgrade the OS of the virtual appliance first before upgrading the vRealize Operations.

Note: There is only one PAK when upgrading to vRealize Operations 8.0 and higher
- Use the appropriate vRealize Operations upgrade PAK file

Starting with vRealize Operations 8.1, there are two PAK files available for upgrade:

 1. Upgrade PAK file includes the OS upgrade files from SUSE to Photon and the vApp upgrade files for upgrading from vRealize Operations Manager 7.5 and lower.
 2. Upgrade PAK file includes the OS upgrade files from Photon to Photon and the vApp upgrade files for upgrading from vRealize Operations 8.0.



- Pre-distribute PAK files to minimize downtime during upgrade

One of the longest steps of the upgrade process is the distribution of the PAK files across all the nodes. To minimize this time, pre-distribute the PAK files to all nodes before starting the upgrade. See [How to reduce vRealize Operations Manager update time by pre-copying software update PAK files](#).
- Upgrade in order of vRealize Operations platform → EPOps agents → Management Packs

Upgrade the vRealize Operations platform first before upgrading the End Point Operations agents. Upgrade the End Point Operations agents from the admin UI using a PAK file. Lastly, upgrade any corresponding management packs.
- Verify functionality after the upgrade

Validate that the same functionality exists after the upgrade completed as compared to before the upgrade started.
- Remove VM snapshots when the upgrade is completed and successfully verified

Remove all VM snapshots post upgrade and verification of the environment as maintaining snapshots will cause performance problems.
- Be mindful when upgrading remote collectors

Remote collectors may be located in distant locations to the vRealize Operations cluster so consider potential latency and performance issues before performing an upgrade. Ensure that the remote collectors meet the latency requirements of less than 200ms. If they do not meet latency requirements, remove those remote collectors from the cluster one-by-one.

To remove high latency remote collectors, bring the cluster offline and take snapshots prior to removing the remote collectors. Then bring the cluster back online and remove each impacted remote collector one-by-one using the UI. After removal of all high latency remote collectors, follow the upgrade process. Once the upgrade is completed, install new remote collectors with the same product version to replace previously removed remote collectors and join the cluster.

Cluster

- Deploy all nodes on identical performance hardware

Deploy all vRealize Operations nodes on identical performance hardware to maintain consistency across nodes and for highest performance.
- Use ESXi with same specifications

Do not mix ESXi specifications as this can cause performance problems with specific nodes causing the vRealize Operations cluster to underperform.
- Use datastores backed by the same hardware resource

Mixing datastores backed by different hardware resources can affect the stability of the vRealize Operations cluster.
- All analytics nodes must be of the same size using out-of-the-box (OOTB) size

Deploy identical analytics nodes based on out-of-the-box sizes (small, medium, large, extra-large). Mixing sizes for different analytics nodes may cause instability and performance problems.
- If an analytics node requires additional compute or storage resources, apply equivalent updates to all other analytics nodes

All analytics nodes must have the same resources with each other; therefore, if upgrading (scaling up) one node, all other analytics nodes in the cluster must also be scaled up equally.

- Size Remote Collectors independently from Analytics nodes sizes
Size remote collectors independently from the analytics nodes within the vRealize Operations cluster using out-of-the-box sizes of standard or large. Mix remote collector sizes between standard and large but size them accordingly for the data they will collect.
- Distribute multiple cluster nodes across multiple hosts
A 1:1 mapping is ideal between hosts and nodes. For example, if a cluster will have eight nodes, use eight hosts. If a 1:1 mapping between hosts and nodes is not possible, use the highest number of available hosts for all nodes.
- Use Cluster DRS affinity rules to separate cluster nodes on hosts
Configure anti-affinity rules to keep as many nodes separated across available hosts.
- Storage DRS should be disabled
- Deploy cluster nodes in a single physical datacenter
It is an unsupported configuration to deploy nodes across multiple data centers even if they are collocated. Keep nodes on a single datacenter to maintain performance and easier maintenance.
- Add only one node at a time
Do not add multiple nodes at the same time as this will cause an unnecessary load on the vRealize Operations cluster.
- Let the node addition complete before adding another node
Allow vRealize Operations to process fully the addition of a single node before adding another node.
- Bring the cluster online only after adding all new nodes
Only bring the cluster online after adding all the planned nodes. Bringing the cluster online after adding each node will cause an unnecessary load on processing.

Backup & Restore

Backup

- It is highly recommended to take only backups during quiet periods
Since a snapshot-based backup happens at the block level, it is important that they are limited, or no changes being performed on the cluster configuration. This will help to ensure a healthy backup.
- It is best to take the cluster offline before backups
This will ensure the data consistency across the cluster and internally within the nodes. You can either shut down the VM before the backup or disable quiescing if the VM cannot be powered off.
- Do not quiesce the file system when the cluster remains online
If the cluster remains online, backup your vRealize Operations multi-node cluster by using vSphere Data Protection or other backup tools, disable quiescing of the file system. Snapshots with quiesce enabled is unsupported and may cause problems when restoring.
- Use resolvable host names and static IP addresses for all nodes
The hostname must be resolvable to ensure consistent communication between nodes. If the hostname fails to resolve or the IP has changed, problems may result.



- All nodes must be powered off and accessible during backups

All nodes in the cluster should be in the same powered state when taking backups to maintain a consistent state when restored. If nodes cannot be powered off, disable quiescing.
- Backup the entire cluster to include all VMs

Restoring only part of the cluster is unsupported and may cause synchronization problems preventing the cluster from going online.
- All VMDK files must be backed up that are part of the virtual appliance

Include all VMDK files in the backup; otherwise, the node may not properly connect to the cluster when restored.
- Backup of all nodes must be performed at the same time

Initiate backups of all nodes (master, replica, data, witness and remote collector) at the same time to maintain synchronization across nodes. Each node may complete their backup at a different time but starting the backup process at the same time minimizes the time differential between nodes when restored.
- Perform backups outside of vRealize Operations internal operations. By default, the following processes run:
 - Dynamic Threshold (DT) Calculation at 2:00 am
 - Capacity Calculation (CIQ) at 9:00 pm (vRealize Operations 6.6.1 and earlier)
 - predictive Distributed Resource Scheduler (pDRS) at 6:00 pm
 - Cost Calculation at 9:00 am (introduced in vRealize Operations 6.7)

Avoid processing overhead of the cluster by performing backup when DT, CIQ, pDRS or Costing are not running. These default times can be modified to avoid conflicts during backups
- Backup at different times from infrastructure backups

If there is a process which maintains a separate backup of the infrastructure, avoid taking backups of the vRealize Operations cluster at the same time.
- Do not backup remote collectors if they are already removed from the vRealize Operations cluster

Remove backing up remote collectors if they have been removed from the vRealize Operations cluster to prevent cluster confusion when restored.

Restore

- Power off and delete the existing cluster before restoring to the same infrastructure

If restoring a backup cluster to the same infrastructure, power off and delete the existing cluster to avoid potential MAC and IP address conflicts.
- Remove remote collectors and deploy new instances if unavailable

Remove remote collectors that report as down or no longer available to bring the cluster online, then add the replacement remote collectors as needed.
- Change the IP Address of Nodes After Restoring a Cluster on a Remote Host if IPs change before bringing cluster online

After you have restored a vRealize Operations cluster to a remote host, change the IP address of the master nodes and data nodes to point to the new host. See [Change the IP Address of Nodes After Restoring a Cluster on a Remote Host](#).



Disaster Recovery

- Use Site Recovery Manager (SRM) for disaster recovery.
VMware Site Recovery Manager is the only supported tool for disaster recovery.
See Disaster Recovery by Using Site Recovery Manager at [vRealize Suite Documentation](#).
- Migrate or recover vRealize Operations virtual machines to an identical network configuration
The recovery site should consist of an identical network configuration, if possible, to minimize transition changes when the recovery site becomes active.
- Change the IPs of the nodes when the recovery site does not have an identical network configuration
See [Change the IP Address of a vRealize Operations Manager Multi Node Deployment](#) or [Change the IP Address of a vRealize Operations Manager Single Node Deployment](#).
- Regularly test recovery plans and always clean up the executed recovery tests
To ensure a reliable recovery, test frequently to ensure latest updates are applied to the recovery site and clean up when tests have been completed.

Self-Monitoring

- Enable Self Service Monitoring Dashboards to help troubleshoot vRealize Operations
In vRealize Operations 6.6.1 and earlier, the self-service monitoring dashboards are not visible by default. Enabling these self-service monitoring dashboards will help provide a quick view of the health of the cluster.
In vRealize Operations 6.7 and later, the Self-Monitoring dashboards are enabled by default and can be found under the vRealize Operations group
- Examine syslog when something goes wrong with vRealize Operations
Viewing syslog will provide additional information to help diagnose potential issues with the cluster.
- Send syslog to vRealize Log Insight (vRLI), if integrated
Sending syslog messages to vRealize Log Insight will allow for faster message viewing and easier identification.
- Enable alerts for vRealize Operations
Enabling alerts will provide immediate notification of issues with the cluster.

API and Integration

- API
Use the API when there is a need to automate a well-defined workflow, such as repeating the same tasks to configure access control for new vRealize Operations users. The API is also useful when performing queries on the vRealize Operations data repository, such as retrieving data for particular assets in your virtual environment. In addition, use the API to extract all data from the vRealize Operations data repository and load it into a separate analytics system.

- 
- **SNMP**
Use manual discovery to perform a port scan through an IP range as an SNMP adapter does not know the location of the SNMP devices that you want to monitor.
 - **Email**
Use the [Realize Operations Email Template Manager](#) to customize the email template, as manual method is error prone.

End Point Operations

Sizing

- **Size End Point Operations agent collection with added Management Packs for Applications**
Understand what adding End Point Operations agents will bring into the environment and size (and resize) the cluster to accept the additional load.
- **Use Remote Collectors for End Point Operations agents**
Funnel data collection through remote collectors to reduce bandwidth from multiple End Point Operations agents and reduce the load on the vRealize Operations analytics cluster.
- **Use external Load Balancers for HA of Remote Collectors for End Point Operations agents**
Use load balancers to group multiple remote collectors when using End Point Operations agents to provide high availability and redundancy.

Deployment

- **Break deployments into groups of OS and platform**
- **Assemble multiple End Point Operations agent targets into groups of OS type and platform (i.e. Linux 64-bit) to allow multiple deployment installations using one installer.**
- **Deploy single End Point Operations agents using simultaneously approach**
To deploy a single End Point Operations agent, use the same approach for installing multiple End Point Operations Agents to build consistency and familiarity. See [Upgrade the End Point Operations Management Agent](#).

Content Best Practices

The Content is the functional part of the product; meaning “sits on” the platform vRealize Operations. This section will cover content such as policies, dashboards, alerts, reports, super metrics, groups, and actions. These best practices will help ensure effective use of the platform for displaying collected data, reporting, and notification.

Metrics

- **Use the metrics that are valuable**
There are many out-of-the-box metrics enabled by default so disable any metrics that are not valuable to reduce the amount of unsolicited noise.

- If you can't figure out what the metric numbers mean, don't use it
Only metrics which are understood provide the most value. If a metric does not make sense, the value is limited and only creates additional noise. Always verify that each metric makes sense.
- Super Metrics
Use consistent naming convention to easily identify the super metrics. Always preview or test the super metric before applying. Enable super metrics on specific objects. Disable super metrics from the policy and remove super metric from object type before deleting the super metric.

Alerts & Symptoms

Review Out-Of-The-Box (OOTB)

- Disable the alerts you do not need
There are many out-of-the-box alerts enabled by default so disable the alerts that are not valuable to minimize alert storm.
If alerts that are not required are not disabled, they may cause potential performance issues over time
- Create simple and straight forward alerts
Keep the combination of symptoms as simple and straightforward as possible to make them easily understood and more precise. Use a series of symptom definitions to describe the incremental levels of concern: warning, immediate, and critical. Create actionable alerts for better remediation.
- Use the Wait Cycle and Cancel Cycle to change sensitivity
Configure wait cycle and cancel cycle to avoid overlapping and gaps between alerts.
- Use actionable recommendations
Using actionable recommendations help resolve the issue quicker by providing the ability to have one-click actions to respond to infrastructure issues.
- Out-of-the-box (OOTB) alerts could be overwhelming
Select the alerts not needed and disable what is non-actionable.
- Minimize the number of alerts
Too many alerts become noise and the users will lose interest.
- Management Pack alerting
Disable any new alerts generated by management packs, which are non-actionable
- Non-actionable alerts
If Alerts are not actionable, they should be on dashboards or reports and not in a mailbox.
- Do not modify out-of-the-box alerts
Clone out-of-the-box content to create your own Symptoms, Recommendations & Alert Definitions before making any changes. An out-of-the-box alert may change after upgrading vRealize Operations or upgrading / installing management packs.
- Use multi-symptom alerts
Using multi-symptom alerts will help negate false positives.



Dashboards

- Dashboards should be quickly identifiable within 5 seconds
Create dashboards to keep the information precise and specific making the dashboards more valuable. Containing too much information in one view can lead to information overload. Do not mix scope.
- Use top-line header as summary
Allows to quickly identify what the content displays by having an informative summary in the header
- Divide dashboards into sections
Separate similar content into sections for quick viewing of data and related information
- Top-N data should not exceed 1 day
Top-N value is best looked at from one day for the most current information.
- Do not mix monitoring and troubleshooting
Keep monitoring separate from troubleshooting to maintain specific information.
- Use color
Color helps to emphasize content within the dashboard and points out more important items.
- Use View List Widgets
View list displays best aggregation of data.
- Naming convention
Use consistent naming conventions throughout dashboards and widgets to make items easily identifiable and understood.
- Tab groups
Groups similar like dashboards and unclutters the Dashboard List to provide quick navigation.
- Uncheck all dashboards not heavily used from dashboard list
Deselecting any dashboards not heavily used from the dashboard list will help avoid rendering performance.



Views and Reports

Views

- Utilize views that are available out-of-the-box (OOTB)
Leverage the many out-of-the-box views that provide much of the needed information.
- Clone views to make changes and rename with your company's naming convention
If minor tweaks are needed from an out-of-the-box view, clone the out-of-the-box view before making changes and save with a naming convention that identifies the company, so it can be easily identified and exported for later use.
- Create customized views
Customize views based on what dashboards and reports need to show with precise information. Use your customized views for your customized dashboards and customized reports

Reports

- Utilize reports that are available out-of-the-box (OOTB)
Leverage the many out-of-the-box reports that provide much of the needed information.
- Clone reports if needed and rename with your company's naming convention
If needing minor tweaks from an out-of-the-box report, clone the out-of-the-box report before making changes and save with a naming convention that identifies the company, so it can be easily identified and exported for later use.
- Create customized reports
Customize reports based on the report user's requirements to show specific and related information.

Super Metrics

- Design super metrics for performance
Avoid calculating large objects or using world level metrics that works on all VMs. Apply to only relevant objects and never apply to all objects.
- Make super metrics reusable
Use Depth greater than 1 to allow vRealize Operations to expand higher levels. Use clear naming convention without including specific object name but use Function Object Metric Units.
- Enabling super metrics only on relevant policy
Enable super metrics for the relevant policy, not the base policy.
- Use group instead of where clause.
Using group instead of where clause is easier to understand.



Policies

- Use policies sparingly
If you can, use groups.
- Clone policies to edit and make changes
If edits are required, clone the policy before making any changes.
- Do NOT change or edit on default policy
At any time, do not directly edit or change the default policy.

Account and Roles

- Avoid using the local 'admin' user
All out-of-the-box content is associated with the 'admin' account. If the 'admin' user is being used, there is no tracking of changes for audit purposes. For POC, create a local account with administrator privilege. For production, integrate with AD/LDAP.
- Utilize service accounts for connection credentials
Use service accounts with meaningful names, not coded convention that is easy to make mistakes. For example, SG-D-VM-MG-01 is not user friendly and easy to make human error.
- Create the roles and accounts to identify specific memberships
Creating specific roles helps identify personas such as storage team, network team, NOC, tenants, and IT Management.
- Grant specific roles
Do not always grant Administrator role to users; use specific roles to limit the permissions.
- Avoid enabling vCenter login when authenticating with AD/LDAP
Minimize authentication options to avoid confusion and translated permissions from vCenter.

Maintenance Schedule

- Specify your regular maintenance so they are excluded from calculation
Prevent skewing of results with reports, views, and dashboards by including regular maintenance schedules.

Grouping

- Group objects
There are four ways objects are grouped: vCenter tags, vCenter folders, vRealize Operations groups, and vRealize Operations tags.

- 
- vRealize Operations also provides Application
 - Application is a group, but with a specific purpose and limitation. The strength is to do multi-tier applications with just one group. The limitation is that there is no dynamic membership.
 - Use Groups for dynamic membership
 - For multi-tier apps, use Application
 - Naming Convention
 - Use consistent naming conventions throughout dashboards and widgets to make items easily identifiable and understood.
 - Do not create too many groups
 - Too many groups will cause added noise and make use more complicated; keep usage to a minimum.

Work Load Placement (WLP)

- Create shared datastores that comply with vMotion best practices
- Ensure hosts in WLP cluster are homogeneous

Predictive Distributed Resource Scheduler (pDRS)

- Enabling pDRS requires actions to be enabled.
- The Action credentials must have administrative permissions on the cluster which is enabling pDRS.
- pDRS may not be enabled for every cluster in vCenter
- vCenter can only receive pDRS data from one vRealize Operations instance.
 - There should only be one vRealize Operations to one vCenter relationship.
- Be careful, if you add another vRealize Operations to the same vCenter
 - Adding another vRealize Operations to an existing vCenter will overwrite the existing vRealize Operations.
- Always check pDRS scale numbers
 - For vRealize Operations, do not enable in clusters > 4k VMs
- You need vSphere 6.5 to enable pDRS functionality
 - It is required to use vSphere 6.5 to enable pDRS functionality when using vRealize Operations



Operations Best Practices

This is the how you use the product, vRealize Operations, in your operations. This includes working with other roles in operations (e.g. NOC, Storage, and Management). This section will detail operations such as processes, roles, groups, and tenants. These best practices will help give the user, the best experience when using the content and platform as part of vRealize Operations.

SDDC Monitoring

- Understand the level of monitoring and the metrics required

There are three levels of monitoring: business, application and infrastructure. It must be clear what tools will monitor what datatypes. For example, syslog needs a log analysis tool like vRealize Log Insight and network flow needs its own tool such as vRealize Network Insight.
- Plan for each role independently

Understand who needs to see what data and how they will see it to make vRealize Operations more effective.
- Plan separate dashboards for each role

Dashboards cannot be generic and consumed across roles as each role looks at data from a different viewpoint.
- Think big but start small

Begin with a small piece and grow from there. For example, start with vSphere and be on top of it since everything else sits on top of it. Then expand deeper into infrastructure and further into applications. Take small steps towards getting big.
- Define the needs

Be clear on what you are looking for and defining it. If you cannot define it, you cannot expect any tool to define it for you.



Additional Best Practices

Documentation Links

The product documentation has several places which also mentions best practices.

SECTION	CHAPTER
Reference Architecture	Best Practices for Deploying vRealize Operations
Cluster Requirements	vRealize Operations Cluster Node Best Practices
Configuring Alerts	Alert Definition Best Practices
Endpoint Operations Management Agent	Security Best Practices for Running End Point Operations Management Agents

Additional Links

There are additional best practice links which may be helpful.

- [vRealize Operations Best Practices](#)