Reference Architecture

NOVEMBER 2023
VMware Aria Operations 8.14

vmware®

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

https://docs.vmware.com/

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

Reference Architecture Overview

1

The *VMware Aria Operations Architecture Guide* provides recommendations for deployment topology, hardware requirements, interoperability, and scalability for VMware Aria Operations.

For information about software requirements, installation, and supported platforms see the VMware Aria Operations documentation.

Best Practices for Deploying VMware Aria Operations

2

Implement all the best practices when you deploy a production instance of VMware Aria Operations.

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

Note The master node is now referred to as the primary node. The master replica node is now referred to as the primary replica node.

- Deploy analytics nodes in the same vSphere Cluster except when activating Continuous Availability.
- Deploy analytics nodes with the same disk size on storage of the same type.
- When activating Continuous Availability, separate analytics nodes into fault domains based on their physical location.
- Depending on the size and performance requirements for analytics nodes, apply Storage DRS Anti-Affinity rules to ensure that nodes are on separate datastores.
- Set Storage DRS to manual for all VMware Aria Operations analytics nodes.
- If you deploy analytics nodes into a highly consolidated vSphere cluster, configure the resource reservation to ensure optimal performance. Ensure that the virtual CPU to physical CPU ratio is not negatively impacting the performance of analytics nodes by validating CPU ready time and CPU co-stop.
- Analytics nodes have a high number of vCPUs to ensure performance of the analytics computation that occurs on each node. Monitor CPU Ready time and CPU Co-Stop to ensure that analytics nodes are not competing for CPU capacity.
- If the sizing guideline provides several configurations for the same number of objects, use the configuration which has the least number of nodes. For example, if the number of collecting is 120,000, configure the cluster with four extra-large nodes instead of 12 large nodes.
- Deploy an extra even number of nodes to activate Continuous Availability. If the current configuration is an odd number of analytics nodes, deploy an extra analytics node to create an even pairing.

Witness Nodes

A witness node is required when continuous availability is activated to manage the analytics nodes in the fault domains. VMware Aria Operations can have only one witness node in its cluster.

- Deploy the witness node before activating continuous availability.
- Deploy the witness node using the witness configuration.
- Deploy the witness node in a different cluster separate from the analytics nodes.

Cloud Proxy

Using cloud proxies in VMware Aria Operations, you can collect and monitor data from your remote data centers. You can deploy one or more cloud proxies in VMware Aria Operations to create a one-way communication between your remote environment and VMware Aria Operations. The cloud proxies work as one-way remote collectors and upload data from the remote environment to VMware Aria Operations. Cloud proxies can support multiple vCenter Server accounts.

Cloud Proxy and Telegraf Agents

- Deploy Cloud Proxy in the same vCenter Server as the end point VMs on which you want to deploy the Telegraf agents. For Cloud Proxy sizing information see the Sizing Guidelines (KB 2093783).
- Ensure that your operating system platform is supported by Cloud Proxy, and the most recent versions of Windows and Linux OS are supported.
- System times must be synchronized between cloud proxy, end point VMs, the vCenter Server, ESX host, and VMware Aria Operations. To ensure synchronized time, use Network Time Protocol (NTP).
- Ensure that all the prerequisites are met. For more information, see Prerequisites.
- Disable UAC on Endpoint VMs before installing the Telegraf agent. If you cannot do this due to security restrictions, see KB 70780 for a work around script.
- Ensure that the version later than 12.3.5 of VMware Tools is installed on the end point VM on which you want to deploy the Telegraf agent.
- To deploy Telegraf agents onto end point VMs, ensure that the following prerequisites are met for the user account being used for deployment:

  Windows - The user account must be either:

  - An administrator account
  - A non-administrator account that is a member of the built-in administrator group

Linux - The user account must be either:

- A root user with all privileges
- A non-root user with all privileges
- A non-root user with specific privileges

For more information, see User Account Prerequisites in the *VMware Aria Operations Configuration Guide*.

Management Packs and Adapters

Various management packs and adapters have specific configuration requirements. Ensure that you are familiar with all prerequisites before you install a solution and configure the adapter instance.

- Utilize collector groups to separate data collection into fault domains when continuous availability is activated.

Deployment Formats

Deploy VMware Aria Operations with the same VMware Aria Operations vApp version for the following node types:

- Primary
- Primary Replica
- Data
- Witness

See the topic, Installing VMware Aria Operations in the *Getting Started with VMware Aria Operations* guide for more information.

Initial Considerations for Deploying VMware Aria Operations

3

For the production instance of VMware Aria Operations to function optimally, your environment must conform to certain configurations. Review and familiarize yourself with these configurations before you deploy a production instance of VMware Aria Operations.

Sizing

For ionformation about the number of monitored resources and how many analytics nodes are supported by VMware Aria Operations, refer to the Sizing Guidelines (KB 2093783). You must plan sizing of your VMware Aria Operations instance to ensure performance and support.

Environment

Deploy analytics nodes in the same vSphere cluster and use identical or similar hosts and storage. If you cannot deploy analytics nodes in the same vSphere cluster, you must deploy them in the same geographical location.

When continuous availability is activated, deploy analytics nodes in fault domains in the same vSphere cluster and use identical or similar hosts and storage. Fault domains are supported on vSphere stretched clusters.

Analytics nodes must be able to communicate with one another always. The following vSphere events might disrupt connectivity.

- vMotion
- Storage vMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

Due to a high level of traffic between analytics nodes, all analytics nodes must be on the same VLAN and IP subnet, and that VLAN is not stretched between data centers, when continuous availability is not activated.

When the continuous availability is activated analytics nodes within each fault domain should be on the same VLAN and IP subnet, and they should be able to communicate with each other.

Latency between analytics nodes cannot exceed 5 milliseconds, except when continuous availability is activated, where latency between fault domains cannot exceed 10 milliseconds but analytics nodes, within each fault domain, still cannot exceed 5 milliseconds. The bandwidth must be equal to or faster than 10 GB per second.

If you deploy analytics nodes into a highly consolidated vSphere cluster, configure resource reservations. See the Sizing Guidelines (KB 2093783) for more information. If you experience performance issues, review the CPU ready and co-stop to determine if the virtual to physical CPU ratio is the cause of the issues. For more information about how to troubleshoot VM performance and interpret CPU performance metrics, see KB 1017926.

You can deploy the witness node behind a firewall. You cannot use NAT between the witness node and analytics nodes.

Multiple Data Centers

VMware Aria Operations can be stretched across data centers only when continuous availability is activated. The fault domains may reside in separate vSphere clusters; however, all analytics nodes across both fault domains must reside in the same geographical location.

For example, the first data center is located in Palo Alto but is configured in two different buildings or in different locations of the city (downtown and mid-town) will have latency that is less than 5 milliseconds. The second data center is located in Santa Clara so the latency between the two data centers is greater than 5 milliseconds but less than 10 milliseconds. See the Sizing Guidelines (KB 2093783) for network requirements.

If VMware Aria Operations is monitoring resources in additional data centers, you must use cloud proxies and deploy the cloud proxies in the remote data centers. You might need to modify the intervals at which the configured adapters on the cloud proxies collect information depending on latency.

It is recommended that you monitor collections to validate that they are completing in less than five minutes. See the Sizing Guidelines (KB 2093783) for latency, bandwidth and sizing requirements. If all requirements are met and collections are still not completing within the default 5 minutes time limit, increase the interval to 10 minutes.

Certificates

A valid certificate signed by a trusted Certificate Authority, private, or public, is an important component when you configure a production instance of VMware Aria Operations. Configure a Certificate Authority signed certificate against the system before you configure agents.

You must include all analytics nodes, witness nodes, and load balancer DNS names in the Subject Alternative Names field of the certificate.

Adapters

It is recommended that you configure adapters to cloud proxies in the same data center as the analytics cluster for large and extra-large deployment profiles. Configuring adapters to cloud proxies improves performance by reducing load on the analytics node. As an example, you might decide to configure an adapter to cloud proxies if the total resources on a given

analytics node begin to degrade the node's performance. You might configure the adapter to a large cloud proxy with the appropriate capacity.

Configure adapters to cloud proxies when the number of resources the adapters are monitoring exceeds the capacity of the associated analytics node.

Authentication

You can use the Platform Services Controller for user authentication in VMware Aria Operations. For more information about deploying a highly available Platform Services Controller instance, see Deploying the vCenter Server Appliance in the *VMware vSphere Documentation*. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

Load Balancer

For more information about load balancer configuration, see the VMware Aria Operations Load Balancing guide.

Scalability Considerations

4

Configure your initial deployment of VMware Aria Operations based on the anticipated use.

For more information about sizing, see the Sizing Guidelines (KB:2093783).

Analytics Nodes

Analytics nodes consist of a primary node, a primary replica node, and data nodes.

For enterprise deployments of VMware Aria Operations, deploy all nodes as medium, large or extra-large deployments, depending on sizing requirements and your available resources.

Scaling Vertically by Adding Resources

If you deploy analytics nodes in a configuration other than large, you can reconfigure the vCPU and memory. It is recommended to scale up the analytics nodes in the cluster before scaling out the cluster with additional nodes. VMware Aria Operations supports various node sizes.

Scaling Vertically by Increasing Storage

You can increase storage independently of vCPU and Memory.

To maintain a supported configuration, data nodes deployed in the cluster must be the same node size.

For more information about increasing storage, see the topic, Add Data Disk Space to a Aria Operations vApp Node in the *Getting Started* guide. You cannot modify the disks of virtual machines that have a snapshot. You must remove all snapshots before you increase the disk size.

Scaling Horizontally (Adding nodes)

To see the number of t extra-large analytics nodes in a cluster, or the number of extra-large nodes in a cluster when continuous availability is activated, see the Sizing Guidelines (KB:2093783).

To maintain a supported configuration, analytics nodes deployed in the cluster must be the same node size.

Witness Node

VMware Aria Operations provides a single size regardless of the cluster size since the witness node does not collect nor process data.

Remote Collectors

Note Fresh deployment of remote collectors is not supported in VMware Aria Operations, starting from version 8.10. Remote collectors are available only if you had deployed them in a previous version of VMware Aria Operations. If you require a new agent to collect data, you must deploy a cloud proxy. For more information on how to deploy a cloud proxy, see the topic, Installing Cloud Proxy in the Getting Started guide.

VMware Aria Operations supports two sizes for remote collectors, standard and large. The maximum number of resources is based on the aggregate resources that are collected for all adapters on the remote collector. In large scale VMware Aria Operations monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed.

Cloud Proxy

VMware Aria Operations supports two sizes for Cloud Proxy, small and large. The maximum number of resources is based on the aggregate resources that are collected for all adapters on the Cloud Proxy. In large scale VMware Aria Operations monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed. Install a remote collector Cloud Proxy in areas when the latency is more than what is prescribed in the Sizing Guidelines. See (KB 2093783) for more information.

High Availability Considerations

5

High availability creates a replica for the VMware Aria Operations primary node and protects the analytics cluster against the loss of a node.

Cluster Management

Clusters consist of a primary node, a primary replica node, and data nodes.

Activating High Availability within VMware Aria Operations is not a disaster recovery solution. When you activate High Availability, information is stored (duplicated) in two different analytics nodes within the cluster. This doubles the system's compute and capacity requirements. If either the primary node or the primary replica node is permanently lost, then you must deactivate, and then reactivate High Availability to reassign the primary replica role to an existing node. This process, which includes a hidden cluster rebalance, can take a long time.

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you activate High Availability, you protect VMware Aria Operations from data loss when only a single node is lost. If two or more nodes are lost, there may be permanent data loss. Deploy each analytics node to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the VMware Aria Operations nodes remain on separate hosts.

Collector Group

In VMware Aria Operations, you can create a collector group. A collector group is a collection of nodes (cloud proxy, and analytics nodes). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

Note A collector group must contain the same type of nodes. You cannot mix cloud proxy, and analytics nodes in a collector group.

If the node running the adapter fails, the adapter is automatically moved to another node in the collector group.

Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see Chapter 8 Adapter and Management Packs Considerations.

Continuous Availability Considerations

6

Continuous Availability (CA) separates the VMware Aria Operations cluster into two fault domains and protects the analytics cluster against the loss of a fault domain.

Cluster Management

Clusters consist of a primary node, a primary replica node, a witness node, and data nodes.

Activating Continuous Availability within VMware Aria Operations is not a disaster recovery solution.

When you activate Continuous Availability, information is stored (duplicated) in two different analytics nodes within the cluster but stretched across fault domains. Due to sizing requirements, continuous availability requires doubling the system's compute and capacity requirements.

If either the primary node or primary replica node is permanently lost, then you must replace the lost node, which will become the new primary replica node. If it is necessary to have the new primary replica node as the primary node, then you can take the current primary node offline and wait until the primary replica node is promoted to the new primary node. Then bring the former primary node back online and it will be the new primary replica node.

Fault Domains

Fault domains consist of analytics nodes, separated into two zones.

A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. When configured, two fault domains allow VMware Aria Operations to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

Witness Node

Witness node is a member of the cluster but not part of the analytics nodes.

To activate CA within VMware Aria Operations, deploy the witness node in the cluster. The witness node does not collect nor store data.

The witness node serves as a tiebreaker when a decision must be made regarding availability of VMware Aria Operations when the network connection between the two fault domains is lost.

Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you activate continuous availability, you protect VMware Aria Operations from data loss if an entire fault domain is lost. If node pairs are lost across fault domains, there may be permanent data loss.

Deploy analytics nodes, within each fault domain, to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the VMware Aria Operations nodes remain on separate hosts.

Collector Group

In VMware Aria Operations, you can create a collector group. A collector group is a collection of nodes (Cloud Proxy, and analytics nodes). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

Note A collector group must contain the same type of nodes. You cannot mix Cloud Proxy, and analytics nodes in a collector group.

When activating continuous availability, collector groups can be created to collect data from adapters within each fault domain.

Collector groups do not have any correlation with fault domains. The functionality of a collector group is to collect data and provide it to the analytics nodes, which then VMware Aria Operations decides how to keep the data.

If the node running the adapter collection fails, the adapter is automatically moved to another node in the collector group.

Theoretically, you can install collectors in any place, provided the networking requirements are being met. However, from a failover perspective, it is not recommended to put all the collectors within a single fault domain. If all the collectors are directed to a single fault domain, VMware Aria Operations stops receiving data if a network outage occurs affecting that fault domain.

Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see Chapter 8 Adapter and Management Packs Considerations.

Continuous Availability FAQs

7

With the introduction of continuous availability in VMware Aria Operations 8, there have been several frequently asked questions. This section is to help increase awareness and knowledge about continuous availability.

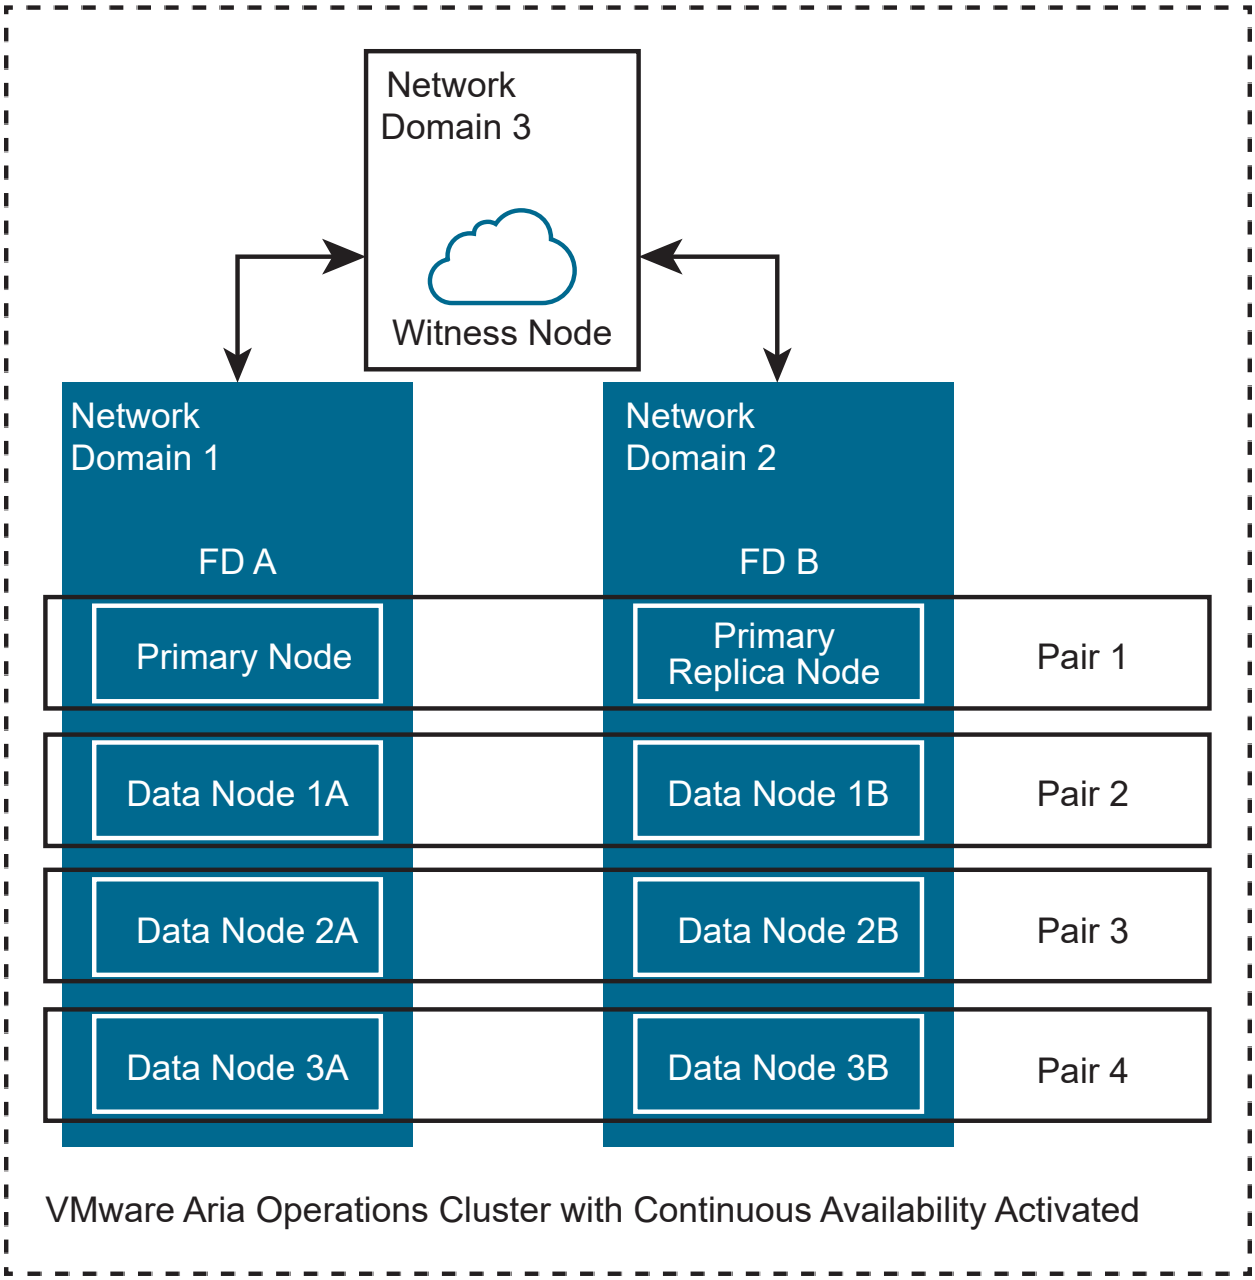

How is the data stored in analytics nodes?

When an object is discovered, VMware Aria Operations determines which node to keep the data, then copies (duplicates) the data to its pair node in the other fault domain. Every object is stored in two analytics nodes (node pairs) across the fault domains and they are always synchronized.

As an example, VMware Aria Operations has eight analytics nodes, CA is activated, and as a result each fault domain has four analytics nodes (see above diagram).

When a new object is discovered, VMware Aria Operations decides to store the data in “Data Node 2B” (primary) and automatically a copy of the data will be saved in “Data Node 2A” (secondary).

If somehow “FD A” becomes unavailable, then “primary” data from “Data Node 2B” will be used.

If somehow “FD B” becomes unavailable, then “secondary” data from “Data Node 2A” will be used.

Which situations break a continuous availability cluster? Simultaneously losing the primary node or primary replica node and data nodes, or two or more data nodes in both fault domains, are not supported.

Each analytics node from fault domain 1 has its node pair in fault domain 2 or vice versa.

Using the previously mentioned example, we will have four node pairs:

Primary + Replica Node

Data Node 1A (FD A) + Data Node 1B (FD B)

Data Node 2A (FD A) + Data Node 2B (FD B)

Data Node 3A (FD A) + Data Node 3B (FD B)

The two nodes of each node pair are always synchronized and storing the same data. Hence, the cluster continues to function without data loss while one node from all node pairs is available.

What happens if one data node from one of the fault domains becomes unavailable?

The cluster will be in a degraded state but continue to operate when one node becomes unavailable in either fault domain. There will be no data loss. The data node must be repaired or replaced so the cluster does not remain in a degraded state.

Will the cluster break if two data nodes in fault domain 1 and the primary replica node in fault domain 2 are lost?

In this example, the cluster will continue to work without data loss. If one analytics node from each node pair is still available, there will be no data loss.

What happens if an entire fault domain becomes unavailable?

The cluster will be in a degraded state but continue to operate when an entire fault domain becomes unavailable. There will be no data loss. The fault domain must be repaired and brought online so the cluster does not remain in a degraded state.

It the fault domain is unrecoverable, it is possible to replace the entire fault domain with newly deployed nodes. From the admin UI, only the primary replica node can be replaced. If the entire fault domain for the primary node is lost, you will need to wait until the primary node failover occurs and the primary replica node has been promoted as the new primary node.

What is the proper process to re-add a failed node to a fault domain? How long will it take to sync up?

The recommended procedure to re-add a failed node is to use the "Replace nodes of cluster" functionality within the admin UI. Once the replacement node has been added, the data will be synced. The sync time strongly depends on the object count, historical period of the objects, network bandwidth, and the load on the cluster.

What happens when network latency between fault domains exceeds 20 ms? How long can VMware Aria Operations tolerate extended latency?

Adhering to latency requirements is necessary to achieve optimal performance. The latency between fault domains should be < 10 ms, with peaks up to 20 ms during 20 sec intervals. For more information about network latency guidelines, see the Sizing Guidelines (KB 2093783).

When network latency between fault domains goes above “20 ms during 20 sec intervals” for some period, but then recovers back to under 10 ms, how long does it take to resynchronize?

High latency does not mean that synchronization has stopped. When an object is discovered, VMware Aria Operations will decide which node needs to keep the data (primary), then a second copy of the data will go to its node pair (secondary). Every object is stored in two analytics nodes (pairs) across both fault domains. Synchronization is an ongoing process where the secondary node is periodically syncs with the primary node. Synchronization is performed based on last synced timestamps of the primary and secondary nodes. Hence, there is no synchronization data queue in VMware Aria Operations.

What is the actual witness node tolerance to missed polls?

Witness node operations are not poll based. The witness node interacts only when one of the nodes is not able to communicate (after various checks) with nodes from the other fault domain.

At what point in time will the primary node and primary replica node failover?

The failover occurs only when the primary node is no longer accessible or not alive.

When is the primary replica node promoted to the primary node?

The primary replica node is promoted to the primary node in only two cases:

- When the existing primary node is down.

- The associated fault domain is down/offline.

When the original primary node returns online, does it resume primary control? How does the data get synchronized?

When operations return to normal, with both primary node and primary replica node online, the newly promoted primary node (formerly primary replica node) remains the new primary node and the new primary replica (formerly primary node) gets synced with the new primary node.

What happens if connectivity between fault domains is completely interrupted, but then recovers?

If communications between the fault domains is completely interrupted for several minutes, then one of the fault domains will automatically go offline. After the network interruption is restored, admin user needs to manually bring the fault domain online which will begin the data synchronization.

What happens to the fault domains when the witness node becomes unavailable?

While both fault domains are healthy and communicating with each other, the unavailability of the witness node will have no effect on the cluster; VMware Aria Operations will continue to function. If there is a communication issue between the fault domains, three situations could occur:

- Witness node is accessible from both fault domains – witness will bring one fault domain offline based on site health.
- Witness node is accessible from one fault domain only – The other fault domain will go offline automatically.
- Witness node is not accessible from both fault domains – Both fault domains will go offline.

When the offline fault domain becomes available again, will the fault domains synchronize all data collected during the communication outage?

The collected data is synchronized immediately once connectivity to the fault domain is restored and synchronized to capture all missed data.

What happens when an analytics node is not able to communicate to analytics nodes in the other fault domain?

If an analytics node is not able to communicate with all nodes from the other fault domain nor the witness node, it will go offline automatically. All nodes or entire fault domain that were taken offline automatically should be brought back online by the Admin user manually after ensuring that all communication issues have been resolved.

If the maximum number of nodes in a standard cluster is 10 extra-large nodes, which supports 440,000 objects, why is the maximum number of nodes in continuous availability more with 12 extra-large nodes, which supports 264,000 objects?

The 12 extra-large nodes are supported only in a continuous availability cluster and references a maximum of six extra large nodes across two separate fault domains. This permits an increase to the number of nodes over a standard cluster and allows for collection for a greater number of objects.

A possible design is six-large nodes in fault domain 1, and six extra-large nodes in fault domain 2, with a witness node in a third site. The latency requirements must be met such that latency between fault domain 1 and fault domain 2 is <10 ms. Details about latency, packet loss and bandwidth are listed in the Sizing Guidelines (KB 2093783).

Is a load balancer supported with Continuous Availability?

Yes, for more information about load balancer configuration, see VMware Aria Operations Load Balancing Configuration guide available under Resources in the the VMware Aria Operations documentation.

When the primary node is connected to the network again after a failover, what is the recommended procedure to return the original primary node to the primary role?

It is not necessary to roll back the primary replica node to the primary node role or vice versa. If you still want to restore the old primary node to the primary role, then use “Take Node Offline/Online” on the new primary node or its fault domain (where the original primary node resides)

Anytime a node goes offline or gets rebooted, is it necessary to bring the corresponding fault domain offline and then online to bring the node back online?

All nodes, after reboot or bringing it offline/online, will automatically continue to work. No additional steps are necessary.

Adapter and Management Packs Considerations

8

Adapters and management packs have specific configuration considerations.

Normal Adapters

Normal adapters require a one-way communication to the monitored endpoint. Deploy normal adapters into collector groups, which are sized to handle a failover.

Following is a sample list of adapters provided by VMware for VMware Aria Operations. Additional adapters can be found on the VMware Solutions Exchange website.

- VMware vSphere
- Management Pack for NSX for vSphere
- Management Pack for VMware Integrated OpenStack
- Management Pack for Storage Devices
- Management Pack for Log Insight

Hybrid Adapters

Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

You must deploy hybrid adapters to a dedicated remote collector. Configure only one hybrid adapter type for each remote collector. You cannot configure hybrid adapters as part of a collector group. For example, two vRealize Operations for Published Applications adapters can exist on the same node, and two vRealize Operations for Horizon adapters can exist on the same node, but a vRealize Operations for Published Applications adapter and a vRealize Operations for Horizon adapter cannot exist on the same node.

Several hybrid adapters are available for VMware Aria Operations.

- vRealize Operations for Horizon adapter
- vRealize Operations for Published Applications adapter
- Management Pack for vRealize Hyperic

Hardware Requirements for Analytics Nodes, Witness Nodes, and Cloud Proxy

9

Analytics nodes, witness nodes, and cloud proxies have various hardware requirements for virtual machines and physical machines.

For information about the components to install on each server profile in your deployment, and the required hardware specifications, see the Sizing Guidelines (KB:2093783).

CPU requirements are 2.0 GHz minimum. 2.4 GHz is recommended. Storage requirements are based on the maximum supported resources for each node.

VMware Aria Operations has a high CPU requirement. In general, the more physical CPU that you assign to the analytics cluster, the better the performance. The cluster will perform better if the nodes stay within a single socket.

Port Requirements for VMware Aria Operations

10

The most up-to-date technical information about ports can be found on Ports and Protocol.

Small Deployment Profile for VMware Aria Operations

11

The small deployment profile is intended for systems that manage up to 20,000 resources.

Virtual Appliance Name

The small deployment profile contains a single large analytics node, analytics-1.ra.local.

Deployment Profile Support

For the configuration which the small deployment profile supports, see the sizing guidelines at KB2093783.

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics-1.ra.local*

This is an example of a small deployment profile.

Table 11-1. Adapter Properties

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytics-1	A	2,000
DEFAULT	analytics-1	B	4,000
DEFAULT	analytics-1	C	2,000
DEFAULT	analytics-1	D	3,000

VMware Aria Operations Small Deployment Profile Architecture

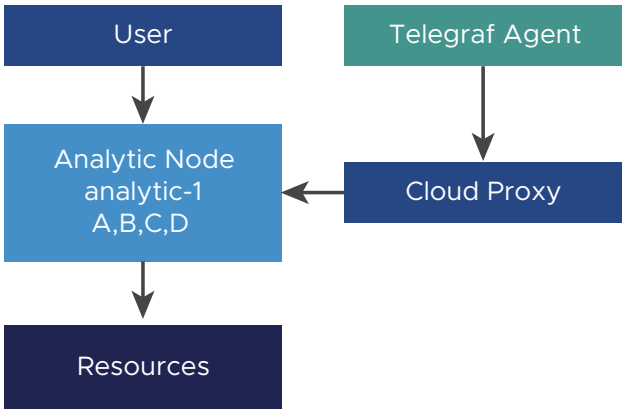

Medium Deployment Profile for VMware Aria Operations

12

The medium deployment profile is intended for systems that manage 68,000 resources, 34,000 of which are activated for High Availability. In the medium deployment profile, adapters are deployed on the analytics nodes by default. If you experience problems with data ingestion, move these adapters to cloud proxies.

Virtual Appliance Names

The medium deployment profile contains eight medium analytics nodes.

- analytics-1.ra.lcoal
- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal
- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

Deployment Profile Support

For the configuration which the medium deployment profile supports, see the sizing guidelines at KB2093783.

Load Balanced Addresses

- analytics.ra.local

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics-1.ra.local*

This is an example of a medium deployment profile.

Table 12-1. Adapter Properties

Collector Group	Collector	Adaptor	Resources
DEFAULT	analytics-1	A	2,000
DEFAULT	analytics-2	B	4,000
DEFAULT	analytics-3	C	2,000
DEFAULT	analytics-4	D	3,000
DEFAULT	analytics-5	E	1,000
DEFAULT	analytics-6	F	2,000
DEFAULT	analytics-7	G	1,500
DEFAULT	analytics-8	H	4,500

VMware Aria Operations Medium Deployment Profile Architecture

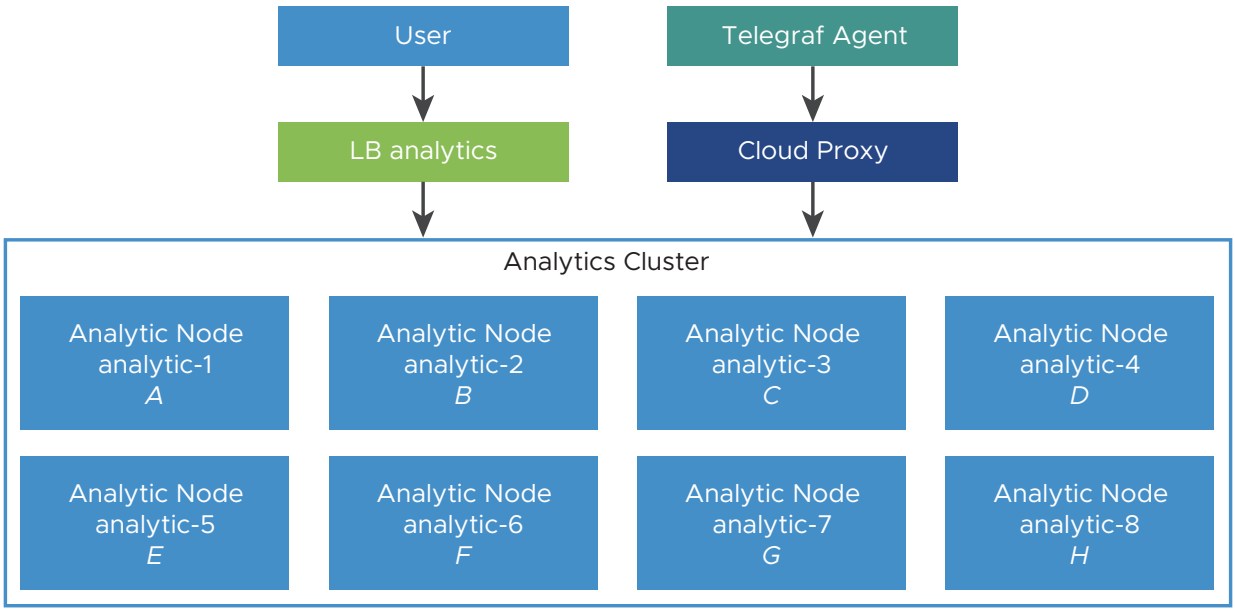

Large Deployment Profile for VMware Aria Operations

13

The large deployment profile is intended for systems that manage 128,000 resources, 64,000 of which are available with High Availability. All adapters are deployed to remote controllers in large deployment profiles to offload CPU usage from the analytics cluster.

Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, and large cloud proxies for adapters and Telegraf agents.

- analytics-1.ra.lcoal
- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal
- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

Deployment Profile Support

For the configuration which the large deployment profile supports, see see the Sizing Guidelines (KB 2093783).

Load Balanced Addresses

- `analytics.ra.local`

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics.refarch.local*
- DNS Name = *analytics-1.ra.local* to DNS Name = *analytics-8.ra.local*
- DNS Name = *remote-1.ra.local* to DNS Name = *remote-N.ra.local*

This is an example of a large deployment profile.

Table 13-1. Adapter Properties

Collector Group	Cloud Proxy	Adapter	Resources
1	CP-1	A	5,000
1	CP-2	B	5,000
		Total	10,000
2	CP-3	C	10,000
2	CP-4	D	5,000
2	CP-5	E	5,000
		Total	20,000

If a cloud proxy is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit for each cloud proxy.

VMware Aria Operations Large Deployment Profile Architecture

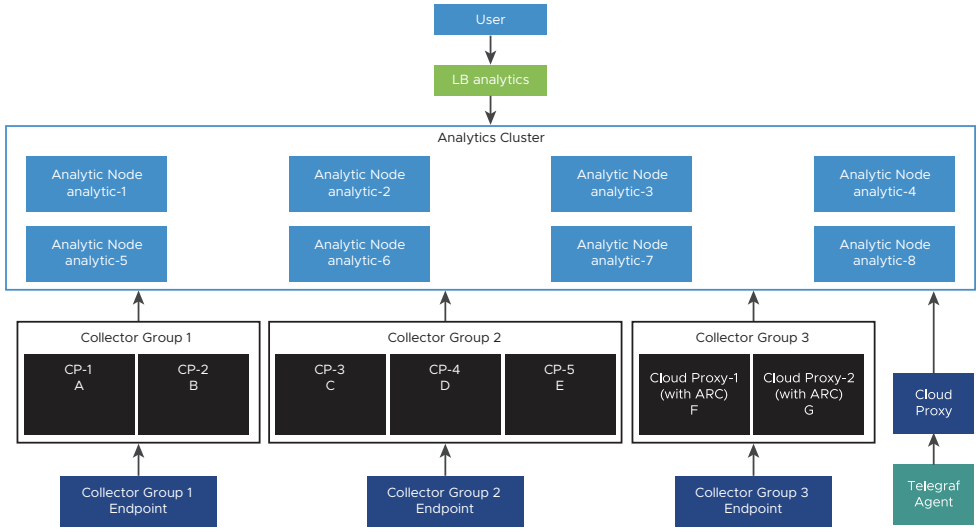

Extra Large Deployment Profile for VMware Aria Operations

14

The extra-large deployment profile is intended for systems that manage 240,000 resources, 120,000 of which are activated for Continuous Availability. This deployment is divided into two data centers and is the maximum supported analytics cluster deployment.

Virtual Appliance Names

The extra-large deployment profile contains six extra-large analytics nodes. Large cloud proxies for adapters and witness node for continuous availability.

- `analytics-1.ra.local`
- `analytics-2.ra.local`
- `analytics-3.ra.local`
- `analytics-4.ra.local`
- `analytics-5.ra.local`
- `analytics-6.ra.local`
- `witness-1.ra.local`

Deployment Profile Support

For the configuration which the extra large deployment profile supports, see the sizing guidelines at KB2093783.

Load Balanced Addresses

- `analytics.ra.local`

Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics.refarch.local*

- DNS Name = *analytics-1.ra.local* to *analytics-16.ra.local*
- DNS Name = *remote-1.ra.local* to *remote-N.ra.local*
- DNS Name = *witness-1.ra.local*

This is an example of an extra-large deployment profile. The adapter in the example provides N-1 redundancy, meaning, if two adapters support 20,000 resources, then a third adapter is added to attain a supported configuration that allows for a single failure.

Table 14-1. Adapter Properties

Collector Group	Data Center	Cloud Proxy	Adapter	Resources
1	A	cp-1	A	5,000
1	A	cp-2	B	5,000
			Total	10,000
2	A	cp-3	C	2,000
2	A	cp-3	D	2,000
2	A	cp-3	E	1,000
2	A	cp-4	F	7,000
2	A	cp-5	G	8,000
2	A	cp-6	H	5,000
2	A	cp-7	I	6,000
			Total	31,000
3	B	cp-8	J	10,000
3	B	cp-9	K	5,000
3	B	cp-10	L	5,000
			Total	20,000

If a cloud proxy is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit for each cloud proxy.

VMware Aria Operations Extra Large Deployment Profile Architecture