

Reference Architecture

12 MAY 2021

vRealize Operations 8.3



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

Copyright © 2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1** Reference Architecture Overview 4
- 2** Best Practices for Deploying vRealize Operations Manager 5
- 3** Initial Considerations for Deploying vRealize Operations Manager 8
- 4** Scalability Considerations 12
- 5** High Availability Considerations 15
- 6** Continuous Availability Considerations 16
- 7** Continuous Availability FAQs 18
- 8** Adapter and Management Packs Considerations 23
- 9** Hardware Requirements for Analytics Nodes, Witness Nodes, and Remote Collectors 25
- 10** Port Requirements for vRealize Operations Manager 26
- 11** Small Deployment Profile for vRealize Operations Manager 27
- 12** Medium Deployment Profile for vRealize Operations Manager 29
- 13** Large Deployment Profile for vRealize Operations Manager 32
- 14** Extra Large Deployment Profile for vRealize Operations Manager 35

Reference Architecture Overview

1

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

For information about software requirements, installation, and supported platforms see [vRealize Operations Manager Documentation](#).

Best Practices for Deploying vRealize Operations Manager

2

Implement all the best practices when you deploy a production instance of vRealize Operations Manager.

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 enabling Continuous Availability.
- Deploy analytics nodes with the same disk size on storage of the same type.
- When enabling 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 vRealize Operations Manager 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 enable Continuous Availability. If the current configuration is an odd number of analytics nodes, deploy an extra analytics node to create an even pairing.

Remote Collector Nodes

Remote collector nodes are additional cluster nodes that allow vRealize Operations Manager to gather more objects into its inventory for monitoring.

- Deploy remote collector nodes when the cluster is online.
- Deploy remote collector nodes one at a time. Adding multiple remote collectors in parallel can cause the cluster to crash.

Witness Nodes

A witness node is required when continuous availability is enabled to manage the analytics nodes in the fault domains.

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

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 remote collector groups to separate data collection into fault domains when continuous availability is enabled.

vRealize Application Remote Collector and Telegraf Agents

- Deploy vRealize Application Remote Collector in the same vCenter Server as the end point VMs on which you want to deploy the Telegraf agents.
- Ensure that your operating system platform is supported by vRealize Application Remote Collector, and the most recent versions of Windows and Linux OS are supported.
- System times must be synchronized between vRealize Application Remote Collector, end point VMs, the vCenter Server, ESX host, and vRealize Operations Manager. To ensure synchronized time, use Network Time Protocol (NTP).
- Disable UAC on Endpoint VMs before installing the Telegraf agent. If you cannot do this due to security restrictions, see [KB article 70780](#) for a work around script.
- Ensure that the latest version 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 *vRealize Operations Manager Configuration Guide*.

Deployment Formats

Deploy vRealize Operations Manager with the same vRealize Operations Manager vApp version for the following node types:

- Primary
- Primary Replica
- Data
- Remote Collector
- Witness

See the *vRealize Operations Manager vApp Deployment and Configuration Guide* for more information.

Initial Considerations for Deploying vRealize Operations Manager

3

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

Sizing

vRealize Operations Manager supports up to 320,000 monitored resources spread across eight extra-large analytics nodes.

Size your vRealize Operations Manager instance to ensure performance and support. For more information about sizing, refer to the KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783) .

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 enabled, 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 enabled.

When continuous availability is enabled, analytics nodes in fault domains should be located on the same VLAN and IP subnet, and communication between fault domains must be available. The witness node might be located in a separate VLAN and IP subnet but must be able to communicate with all analytics nodes.

Latency between analytics nodes cannot exceed 5 milliseconds, except when continuous availability is enabled, 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. A full analytics node, for example a large analytics node that monitors 20,000 resources, requires one virtual CPU to physical CPU. 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 [Troubleshooting a virtual machine that has stopped responding: VMM and Guest CPU usage comparison \(1017926\)](#).

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

Multiple Data Centers

vRealize Operations Manager can be stretched across data centers only when continuous availability is enabled. 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 datacenter 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 datacenter is located in Santa Clara so the latency between the two datacenters is greater than 5 milliseconds but less than 10 milliseconds. Refer to the KB article, [vRealize Operations Manager Sizing Guidelines \(KB 2093783\)](#) for network requirements.

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

It is recommended that you monitor collections to validate that they are completing in less than five minutes. Check the KB article, [vRealize Operations Manager 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 vRealize Operations Manager.

Configure a Certificate Authority signed certificate against the system before you configure End Point Operations Management agents.

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

You can configure End Point Operations Management agents to trust the root or intermediate certificate to avoid having to reconfigure all agents if the certificate on the analytics nodes and remote collectors is modified. For more information about root and intermediate certificates, see *Specify the End Point Operations Management Agent Setup Properties* in the *VMware vRealize Operations Manager Configuration Guide*.

Adapters

It is recommended that you configure adapters to remote collectors in the same data center as the analytics cluster for large and extra-large deployment profiles. Configuring adapters to remote collectors improves performance by reducing load on the analytics node. As an example, you might decide to configure an adapter to remote collectors if the total resources on a given analytics node begin to degrade the node's performance. You might configure the adapter to a large remote collector with the appropriate capacity.

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

vRealize Application Remote Collector

For the production instance of vRealize Application Remote Collector and Telegraf agents to function optimally, your environment must adhere to certain configurations. You must review these configurations before you start deploying vRealize Application Remote Collector and Telegraf agents.

Option	Configurations
Sizing	The vRealize Application Remote Collector supports up to a maximum of 10,000 Telegraf agents using a large vRealize Application Remote Collector. Size your vRealize Application Remote Collector instance to ensure optimal performance and support. For more information about sizing, see the KB article, vRealize Operations Manager Sizing Guidelines (KB 2093783) .
Environment	Deploy the vRealize Application Remote Collector in the same vCenter Server as the end point VMs where you want to deploy Telegraf agents. Latency between the vRealize Application Remote Collector and a vRealize Operations Manager node cannot exceed 10 milliseconds.

Authentication

You can use the Platform Services Controller for user authentication in vRealize Operations Manager. 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 *vRealize Operations Manager Load Balancing Guide*.

Scalability Considerations

4

Configure your initial deployment of vRealize Operations Manager based on the anticipated use.

Analytics Nodes

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

For enterprise deployments of vRealize Operations Manager, 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. vRealize Operations Manager supports various node sizes.

Table 4-1. Analytics Nodes Deployment Sizes

Node Size	vCPU	Memory
Extra small	2	8 GB
Small	4	16 GB
Medium	8	32 GB
Large	16	48 GB
Extra large	24	128 GB

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 vRealize Operations Manager vApp Node*. 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)

vRealize Operations Manager supports up to eight extra-large analytics nodes in a cluster, or up to 10 extra-large nodes in a cluster when continuous availability is enabled.

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

Witness Node

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

Table 4-2. Witness Node Deployment Sizes

Node Size	vCPU	Memory
Witness	2	8 GB

Remote Collectors

vRealize Operations Manager 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 vRealize Operations Manager monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed. Determine the areas of the environment in which the latency is greater than 20 milliseconds and install a remote collector in those areas.

Table 4-3. Supported Remote Collector Sizes

Collector Size	Resources	End Point Operations Management Agents
Standard	6000	250
Large	32,000	2,000

For more information about sizing, see the KB article [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783).

vRealize Application Remote Collector

vRealize Operations Manager supports three sizes for application remote collectors; small, medium, and large. The number of Telegraf agents you want to deploy determines the size of vRealize Application Remote Collector you deploy.

Currently, vRealize Application Remote Collector can collect data on 20 different application sources.

If you have more than 6,000 Telegraf agents installed, increase vCPU and memory of the large configurations so that you can monitor up to 10,000 Telegraf agents.

An increase of memory usage depends on the number of services and their configurations on the VMs that are being monitored. When you monitor 1000 operating system objects, memory usage increases by around 1-1.5 GB.

Table 4-4. Supported vRealize Application Remote Collector Sizes

vRealize Application Remote Collector Size	Maximum number of Telegraf Agents Supported
Small	500
Medium	3000
Large	6000

High Availability Considerations

5

High availability creates a replica for the vRealize Operations Manager 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, data nodes, and remote collector nodes.

Enabling High Availability within vRealize Operations Manager is not a disaster recovery solution. When you enable 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 disable, and then re-enable 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 enable High Availability, you protect vRealize Operations Manager 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 vRealize Operations Manager nodes remain on separate hosts.

Remote Collectors

In vRealize Operations Manager, you can create a collector group. A collector group is a collection of nodes (analytics nodes and remote collectors). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

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 vRealize Operations Manager 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, data nodes, and remote collector nodes.

Enabling Continuous Availability within vRealize Operations Manager is not a disaster recovery solution.

When you enable 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 enable vRealize Operations Manager 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 enable CA within vRealize Operations Manager, 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 vRealize Operations Manager 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 enable continuous availability, you protect vRealize Operations Manager 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 vRealize Operations Manager nodes remain on separate hosts.

Remote Collectors

In vRealize Operations Manager, you can create a collector group. A collector group is a collection of nodes (analytics nodes and remote collectors). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

When enabling 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 vRealize Operations Manager 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, vRealize Operations Manager stops receiving data if a network outage occurs affecting that fault domain.

The recommendation is to keep remote collectors outside of fault domains or keep half of the remote collectors in fault domain 1 and the remaining remote collectors in fault domain 2.

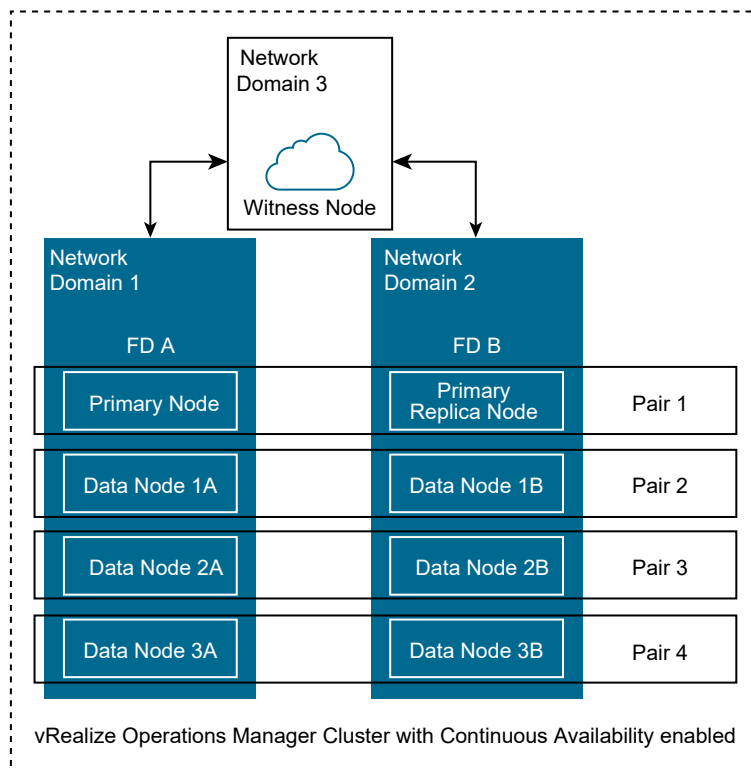
Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see [Chapter 8 Adapter and Management Packs Considerations](#).

Continuous Availability FAQs

7

With the introduction of continuous availability in vRealize Operations Manager 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, vRealize Operations Manager 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, vRealize Operations Manager has eight analytics nodes, CA is enabled, and as a result each fault domain has four analytics nodes (see above diagram).

When a new object is discovered, vRealize Operations Manager 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.

If 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 vRealize Operations Manager 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 KB article [vRealize Operations Manager 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, vRealize Operations Manager 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 vRealize Operations Manager.

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; vRealize Operations Manager 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 8 extra-large nodes, which supports 320,000 objects, why is the maximum number of nodes in continuous availability more with 10 extra-large nodes, which supports 200,000 objects?

The 10 extra-large nodes are supported only in a continuous availability cluster and references a maximum of five 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 five extra-large nodes in fault domain 1, and 5 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 KB article, [vRealize Operations Manager Sizing Guidelines](#) (KB 2093783).

Is a load balancer supported with Continuous Availability?

Yes, for more information about load balancer configuration, see vRealize Operations Manager Load Balancing Configuration guide available under Resources in the [vRealize Operations Manager Documentation](#) page.

The documentation states, “When CA is enabled, the replica node can take over all functions that the primary node provides, in case of a primary node failure. The failover to the replica is automatic and requires only two to three minutes of vRealize Operations Manager downtime to resume operations and restart data collection.”

During testing, by disconnecting the network interface on the primary node, the switchover to the new primary worked within 5 minutes, you get kicked out of the product UI or get strange errors.

The stated two or three minutes are approximate medium values, so 5 minutes is acceptable.

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



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 vRealize Operations Manager. 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 vRealize Operations Manager.

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

End Point Operations Management Adapter

By default, End Point Operations Management adapters are installed on all data nodes. Large and extra-large analytics nodes can support 2,500 endpoint agents and large remote

collectors can support 2,000 per node. To reduce the ingestion load on the cluster, you can point End Point Operations Management adapters at remote collectors. Assign the dedicated remote collectors to their own collector group, which helps the End Point Operations Management adapter maintain the state of End Point Operations Management resources if a node in the collector group fails.

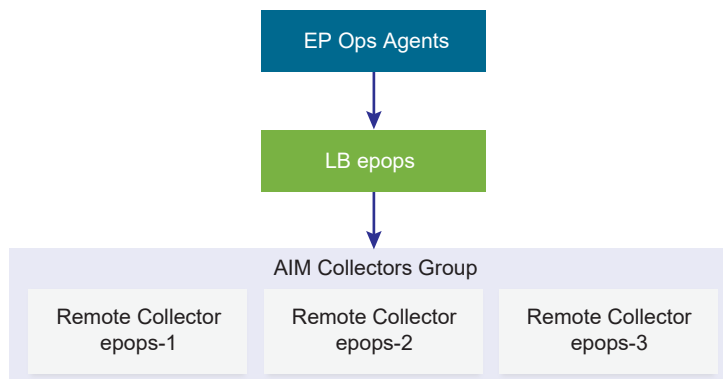
To reduce the cost of reconfiguring the system, it is recommended that you install End Point Operations Management agents against a DNS entry specific to End Point Operations Management agents if you plan to scale the system beyond a single node.

vRealize Application Management Pack

When you activate application monitoring in vRealize Operations Manager, as part of the process, you have to either download the vRealize Application Remote Collector OVA from within vRealize Operations Manager or externally from My VMware.

This is a dedicated virtual appliance which acts as a proxy between vRealize Operations Manager, the target vCenter Server, and the end point VMs where Telegraf agents are deployed.

Remote Collectors Behind a Load Balancer for End Point Operations Management Agents



Hardware Requirements for Analytics Nodes, Witness Nodes, and Remote Collectors

9

Analytics nodes, witness nodes, and remote collectors have various hardware requirements for virtual machines and physical machines.

The following table specifies the components to install on each server profile in your deployment, and the required hardware specifications.

Table 9-1. Hardware Requirements for System Components

Server Roles	Virtual CPU	Memory	Storage Requirements
Small Analytics Node	4 vCPU	16 GB	1276 IOPS
Medium Analytics Node	8 vCPU	32 GB	1875 IOPS
Large Analytics Node	16 vCPU	48 GB	3750 IOPS
Extra Large Analytics Node	24 vCPU	128 GB	12758 IOPS
Standard Remote Collector	2 vCPU	4 GB	N/A
Large Remote Collector	4 vCPU	16 GB	N/A
Witness Node	2 vCPU	8 GB	N/A
Small vRealize Application Remote Collector	4 vCPU	8 GB	N/A
Medium vRealize Application Remote Collector	8 vCPU	16 GB	N/A
Large vRealize Application Remote Collector	16 vCPU	24 GB	N/A

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

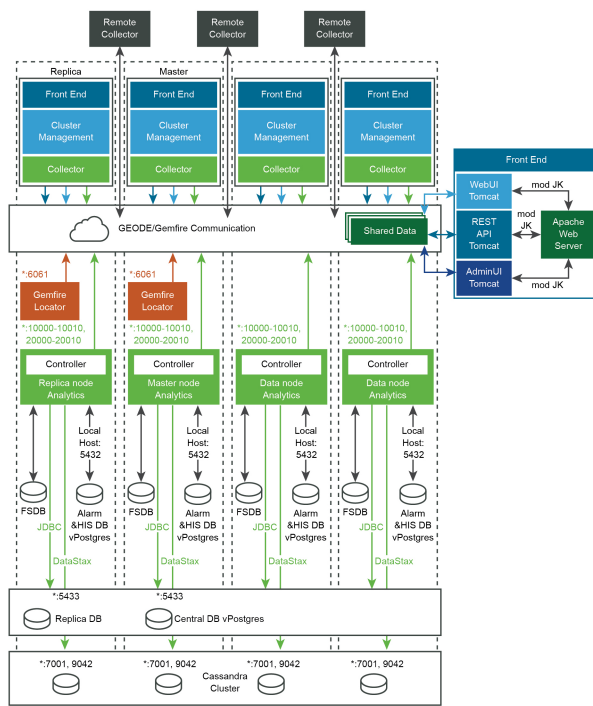
vRealize Operations Manager has a high CPU requirement. In general, the more physical CPU that you assign to the analytics cluster, the better the performance. The cluster will perform better if the nodes stay within a single socket.

Port Requirements for vRealize Operations Manager

10

vRealize Operations Manager has certain port requirements for its components. All ports specified are default ports.

Port Requirements for vRealize Operations Manager



Ports Information for vRealize Operations Manager

Ports information for vRealize Operations Manager is available on [Ports and Protocol](#).

Small Deployment Profile for vRealize Operations Manager

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

The small deployment profile supports the following configuration.

- 20,000 resources
- 2,500 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 months

Additional DNS Entries

You can add additional DNS entries for your organization's future requirements. If you do not expect your planned deployment to exceed a single node, you can configure End Point Operations Management agents against the analytics nodes.

`epops.ra.local -> analytics-1.ra.local`

Certificate

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

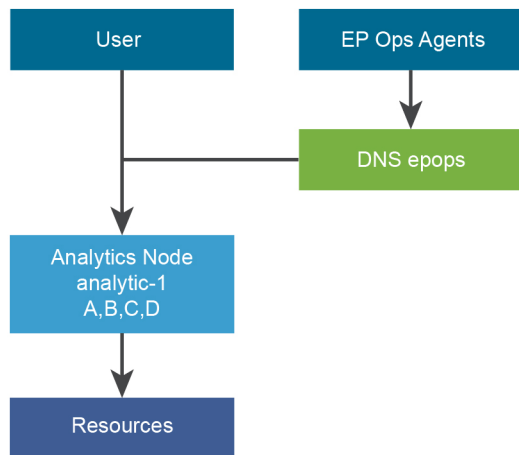
- DNS Name = *epops.refarch.local*
- DNS Name = *analytics-1.ra.local*

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

vRealize Operations Manager Small Deployment Profile Architecture



Medium Deployment Profile for vRealize Operations Manager

12

The medium deployment profile is intended for systems that manage 68,000 resources, 34,000 of which are enabled 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 remote controllers.

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

The medium deployment profile supports the following configuration.

- 68,000 total resources, 34,000 enabled for HA
- 9,600 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 months

Load Balanced Addresses

- analytics.ra.local
- epops.ra.local

Certificate

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

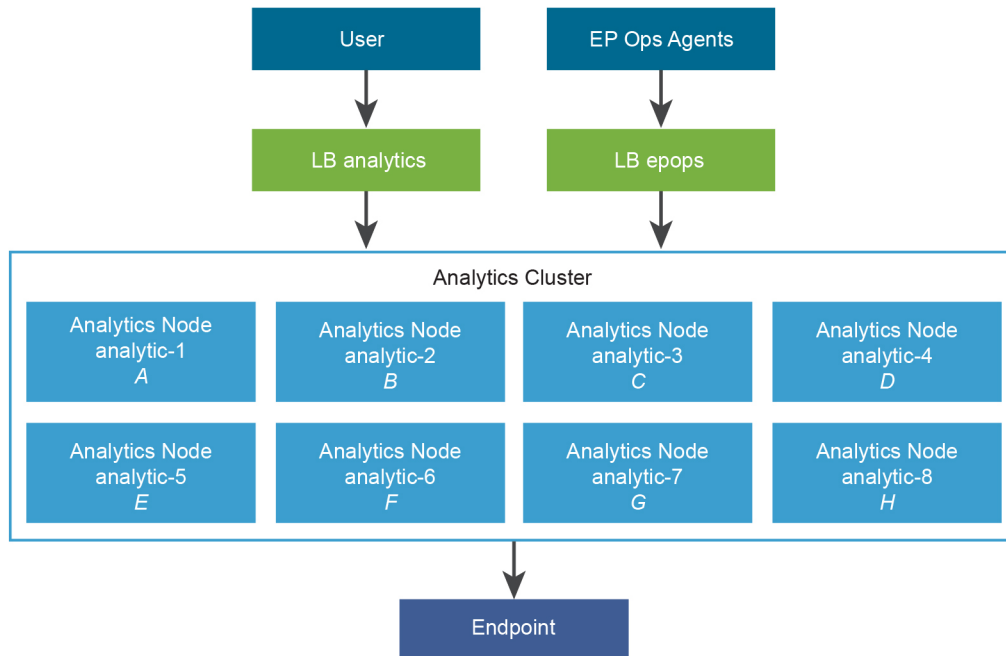
- DNS Name = *epops.refarch.local*
- DNS Name = *analytics-1.ra.local*

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

vRealize Operations Manager Medium Deployment Profile Architecture



Large Deployment Profile for vRealize Operations Manager

13

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

In addition, vRealize Application Remote Collector can be deployed to collect application level data for up to 6,000 end point VMs using Telegraf agents.

Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, large remote collectors for adapters, and large remote collectors for 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

The large deployment profile supports the following configuration.

- 128,000 total resources, 64,000 enabled for HA
- 6,000 Telegraf agents
- 20,000 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 months

Load Balanced Addresses

- analytics.ra.local
- epops.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 = *epops.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*
- DNS Name = *epops-1.ra.local* to DNS Name = *epops-N.ra.local*

This is an example of a large deployment profile.

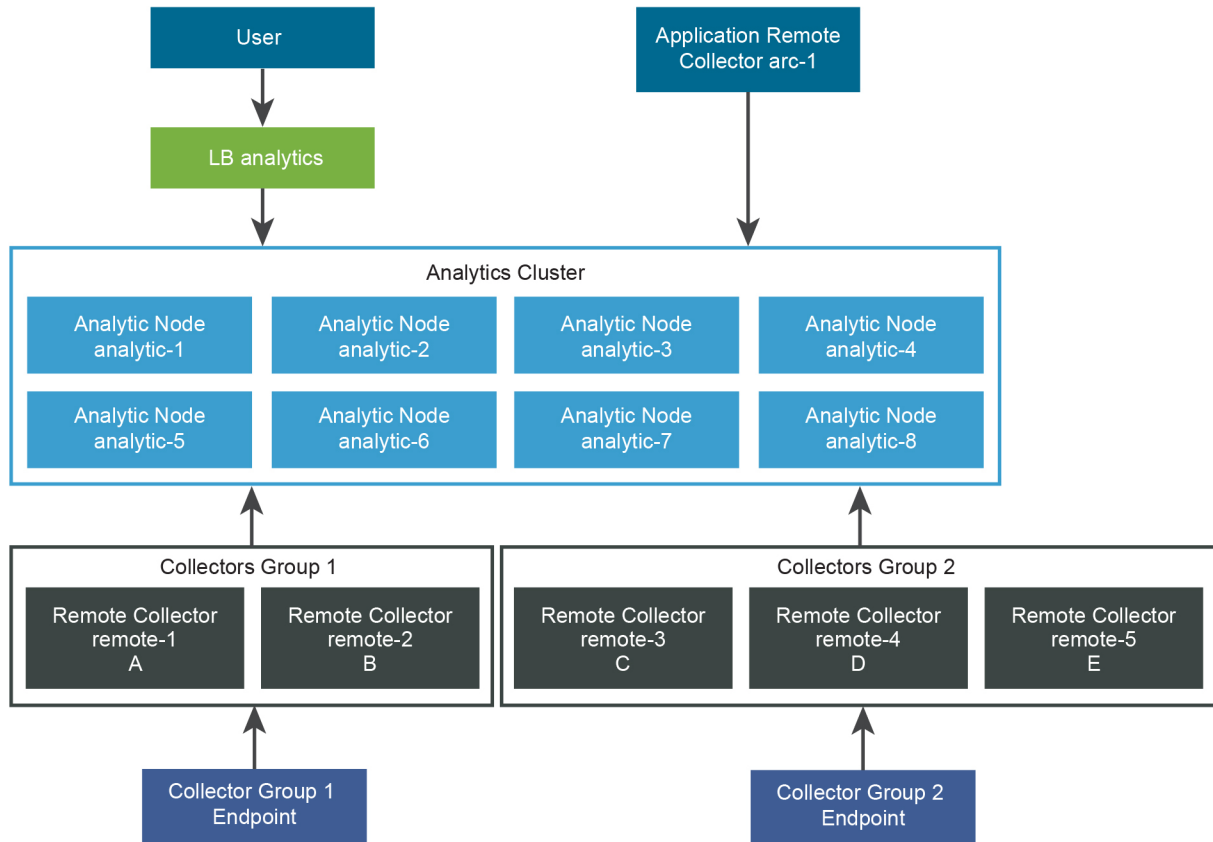
Table 13-1. Adapter Properties

Collector Group	Remote Collector	Adapter	Resources	End Point Operations Management Agents
1	remote-1	A	5,000	N/A
1	remote-2	B	5,000	N/A
		Total	10,000	N/A
2	remote-3	C	10,000	N/A
2	remote-4	D	5,000	N/A
2	remote-5	E	5,000	N/A
		Total	20,000	N/A
AIM	epops-1	epops	4,800	800
	epops-2	epops	4,800	800
		Total	9,600	1,600

If a remote collector is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit of 32,000 resource for each remote collector.

The estimate of 9,600 resources uses six resources for each End Point Operations Management agent.

vRealize Operations Manager Large Deployment Profile Architecture



Extra Large Deployment Profile for vRealize Operations Manager

14

The extra-large deployment profile is intended for systems that manage 240,000 resources, 120,000 of which are enabled 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 remote collectors for adapters, large remote collectors for End Point Operations Management agents, 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

- 240,000 total resources, 120,000 enabled for CA
- 20,000 End Point Operations Management agents
- Data retention for six months
- Additional Time Series Retention for 36 months

Load Balanced Addresses

- `analytics.ra.local`
- `epops-a.ra.local`
- `epops-b.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 = *epops-a.refarch.local*
- DNS Name = *epops-b.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 = *epops-1.ra.local* to *epops-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	Remote Collector	Adapter	Resources	End Point Operations Management agents
1	A	remote-1	A	5,000	N/A
1	A	remote-2	B	5,000	N/A
Total				10,000	
2	A	remote-3	C	2,000	N/A
2	A	remote-3	D	2,000	N/A
2	A	remote-3	E	1,000	N/A
2	A	remote-4	F	7,000	N/A
2	A	remote-5	G	8,000	N/A
2	A	remote-6	H	5,000	N/A
2	A	remote-7	I	6,000	N/A
Total				31,000	
3	B	remote-8	J	10,000	N/A
3	B	remote-9	K	5,000	N/A
3	B	remote-10	L	5,000	N/A
Total				20,000	

Table 14-1. Adapter Properties (continued)

Collector Group	Data Center	Remote Collector	Adapter	Resources	End Point Operations Management agents
AIM-1	A	epops-1	epops	8,004	1,334
AIM-1	A	epops-2	epops	7,998	1,333
	A	epops-3	epops	7,998	1,333
			Total	24,000	4,000
AIM-2	B	epops-4	epops	8,004	1,334
AIM-2	B	epops-5	epops	7,998	1,333
AIM-2	B	epops-6	epops	7,998	1,333
			Total	24,000	4,000

If a remote collector is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit of 32,000 resource for each remote collector.

The estimate of 24,000 resources for AIM-1 and AIM-2 collector groups uses six resources for each End Point Operations Management agent.

vRealize Operations Manager Extra Large Deployment Profile Architecture

