

vRealize Operations Manager Reference Architecture

vRealize Operations Manager 6.3

vmware[®]

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

<https://docs.vmware.com/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

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

Contents

1	VMware vRealize Operations Manager Reference Architecture Overview	5
2	Best Practices for Deploying vRealize Operations Manager	7
3	Initial Considerations for Deploying vRealize Operations Manager	9
4	Scalability Considerations	11
5	High Availability Considerations	13
6	Adapter and Management Packs Considerations	15
7	Hardware Requirements for Analytic Nodes and Remote Collectors	17
8	Port Requirements for vRealize Operations Manager	19
9	Small Deployment Profile for vRealize Operations Manager	23
10	Medium Deployment Profile for vRealize Operations Manager	25
11	Large Deployment Profile for vRealize Operations Manager	27
12	Extra Large Deployment Profile for vRealize Operations Manager	31
	Index	37

VMware vRealize Operations Manager Reference Architecture Overview

1

The *VMware vRealize Operations Manager Reference Architecture Guide* provides recommendations for deployment topology, hardware requirements, and interoperability, and scalability for VMware vRealize Operations Manager 6.3.

For information about software requirements, installation, and supported platforms see [VMware vRealize Operations Manager 6.3 Documentation](#).

Best Practices for Deploying vRealize Operations Manager

2

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

Analytics Nodes

Analytics nodes consist of master nodes, replica nodes, and data nodes.

- Deploy analytics nodes in the same vSphere Cluster.
- Deploy analytics nodes on storage of the same type.
- 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 resource reservation to ensure optimal performance. Ensure that the virtual CPU to physical CPU ratio is not negatively impacting the performance of analytic 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.

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.

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 120,000 monitored resources spread across 16 analytic nodes.

Size your vRealize Operations Manager instance to ensure performance and support. For more information about sizing see the following KB article: https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2093783.

Environment

Deploy analytic nodes in the same vSphere cluster and use identical or similar hosts and storage. If you cannot deploy analytic nodes in the same vSphere cluster, you must deploy them in the same geographical location. vRealize Operations Manager does not support deploying analytics nodes in multiple geographical locations.

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

- vMotion
- Storage vMotion
- HA
- DRS

Due to a high level of traffic between analytics nodes, all analytics nodes must be Layer 2 Adjacent. Layer 2 Adjacent means that each node is located on the same VLAN and IP subnet, and that VLAN is not stretched between data centers. Latency between analytics nodes cannot exceed 5 milliseconds, and the bandwidth must be equal to or higher than 1 GB per second. It is recommended that bandwidth be 10 GB per second.

If you deploy analytics nodes in to a highly consolidated vSphere cluster, configure resource reservations. A full analytics node, for example a large analytics node that monitors 10,000 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 ration 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 behind a firewall. You cannot use NAT between remote collectors and analytics nodes.

Multiple Data Centers

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 center. 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 latency between sites is less than 200ms. When latency exceeds 200ms, it is recommended that you monitor collections to validate that they are completing in less than five minutes. If collections are not completed in this 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 Endpoint Operations Management agents.

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

You can configure Endpoint 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 are modified. For more information about root and intermediate certificates, see [Specify the End Point Operations Management Agent Setup Properties](#).

Adapters

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

You should also deploy adapters to remote collectors 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 vRealize Operations Manager. For more information about deploying a highly-available Platform Services Controller instance, see [VMware vCenter Server 6.0 Deployment Guide](#).

Load Balancer

For more information about load balancer configuration, see the vRealize Operations Manager documentation.

Scalability Considerations

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

Analytics Nodes

Analytics nodes consist of master nodes, master replica nodes, and data nodes.

For enterprise deployments of vRealize Operations Manager, deploy all nodes as medium or large deployments, depending on 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. 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

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 [Add Disk Space to a vApp Node](#) or [Add Disk Space to Linux or Windows Node](#). You cannot modify the disks of virtual machines that have a snapshot. You must remove all snapshots before you increase disk size.

Scaling Horizontally (Adding nodes)

vRealize Operations Manager 6.2 supports up to 16 analytic nodes in a cluster.

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

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 4-2. Supported Remote Collector Sizes

Collector Size	Resources	Endpoint Operations Management Agents
Standard	1,500	250
Large	12,000	2,500

High Availability Considerations

HA creates a replica for the vRealize Operations Manager master node and protects the analytics cluster against the loss of a node.

Cluster Management

Clusters consist of master nodes and master replica nodes.

When you enable High Availability, information is stored on the master nodes and master replica nodes.

If the master nodes or master replica nodes are permanently lost, then you must disable and re-enable high availability to reassign the master roles or master replica roles. This process, which includes a hidden cluster rebalance, can take a long time.

Analytics Nodes

Analytics nodes consist of master nodes, master replica nodes, and data nodes.

Enabling High Availability within vRealize Operations Manager is not a disaster recovery solution. Enabling High Availability duplicates data in the system, and doubles the system's compute and capacity requirements. When you enable high availability, you protect vRealize Operations Manager from data loss in the event that a single node is lost. If two or more nodes are lost, the data loss is permanent.

Deploy all analytics nodes to separate hosts to reduce the chance of data loss in the event that a host fails. You can use DRS anti-affinity rules to ensure that VMs remain on separate hosts.

Adapters

In vRealize Operations Manager 6.1 and later, you can create a collector group. A collector group is a collection of nodes (analytic nodes and remote collectors). You can assign adapters to the 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. Do not deploy hybrid adapters in collector groups. For more information about adapters, see the documentation for the specific adapters.

Adapter and Management Packs Considerations

6

Adapters and management packs have specific configuration considerations.

Normal Adapters

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

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

- vSphere adapter
- Management Pack for NSX for vSphere
- Management Pack for OpenStack
- Management Pack for Storage Devices
- Management Pack for Log Insight

Hybrid Adapters

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

You must deploy hybrid adapters to a dedicated remote controller. You should configure only one hybrid adapter type for each remote controller. 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

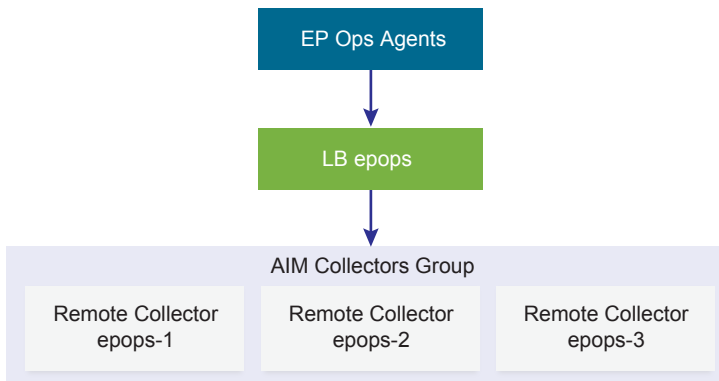
Endpoint Operations Management Adapter

By default, Endpoint Operations Management adapters are installed on all data nodes. Large analytic nodes can support 2,500 agents and large remote collectors can support 2,000 to 10,000 agents for a single cluster. To reduce ingestion load on the cluster, you can point Endpoint Operations Management adapters at remote collectors. You should

assign the dedicated remote collectors to their own collector group, which helps the Endpoint Operations Management adapter maintain the state of Endpoint 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 Endpoint Operations Management agents against a DNS entry specific to Endpoint Operations Management agents if you plan to scale the system beyond a single node.

Remote Collectors Behind a Load Balancer for Endpoint Operations Management Agents



Hardware Requirements for Analytic Nodes and Remote Collectors

7

Analytics 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 7-1. Hardware Requirements for System Components

Server Roles	Virtual CPU	Memory	CPU Requirements	Storage Requirements
Medium analytic node	8 vCPU	32 GB	2.0 Ghz minimum, 2.4 Ghz recommended	1875 IOPS
Large analytic node	16 vCPU	48 GB	2.0 Ghz minimum, 2.4 Ghz recommended	3750 IOPS
Standard remote collector	2 vCPU	4 GB	2.0 Ghz minimum, 2.4 Ghz recommended	N/A
Large remote collector	4 vCPU	16 GB	2.0 Ghz minimum, 2.4 Ghz recommended	N/A

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. You must use a minimum of eight physical CPU dual socket hosts, but it is recommended that you use 12 or more physical CPU dual socket hosts.

Port Requirements for vRealize Operations Manager

8

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

Internal Communications

The following components require internal communication.

Table 8-1. Communication Between Master Node and Replica Node

Component	Protocol	Port
XDB	TCP	1235
Postgres Replica Database	TCP	5433

The XDB ports are required only when you upgrade to vRealize Operations Manager 6.1 or later and are not required for after the upgrade.

Table 8-2. Communication Between Analytics Nodes

Component	Protocol	Port
HTTPS	TCP	443
Gemfire Locator	TCP	6061
Gemfire	TCP	10000
Gemfire	TCP	20000:20010
Cassandra (inter-node)	TCP	7001
Cassandra client	TCP	9042

Table 8-3. Communication From Remote Collector to Analytics Node

Component	Protocol	Port
HTTPS	TCP	443
Gemfire Locator	TCP	6061,
Gemfire	TCP	10000

Table 8-4. Communication Between Remote Collector and Analytics Node

Component	Protocol	Port
HTTPS (Casa)	TCP	443

Table 8-5. Communication Between Remote Collector and Master and Data Nodes

Component	Protocol	Port
HTTP	TCP	80
HTTPS	TCP	443
Gemfire Locator	TCP	6061
Gemfire	TCP and UDP	10000:10010
Gemfire	TCP and UDP	20000:20010
NTP	UDP	123

Table 8-6. Communication From Endpoint Operations Management Agent to Analytics Node

Component	Protocol	Port
HTTPS	TCP	443

Table 8-7. Communication From Endpoint Operations Management Agent to Remote Collector

Component	Protocol	Port
HTTPS	TCP	443

External Communications

The following components require external communications.

Table 8-8. Communication from Analytics Nodes and Remote Collectors to External Resources

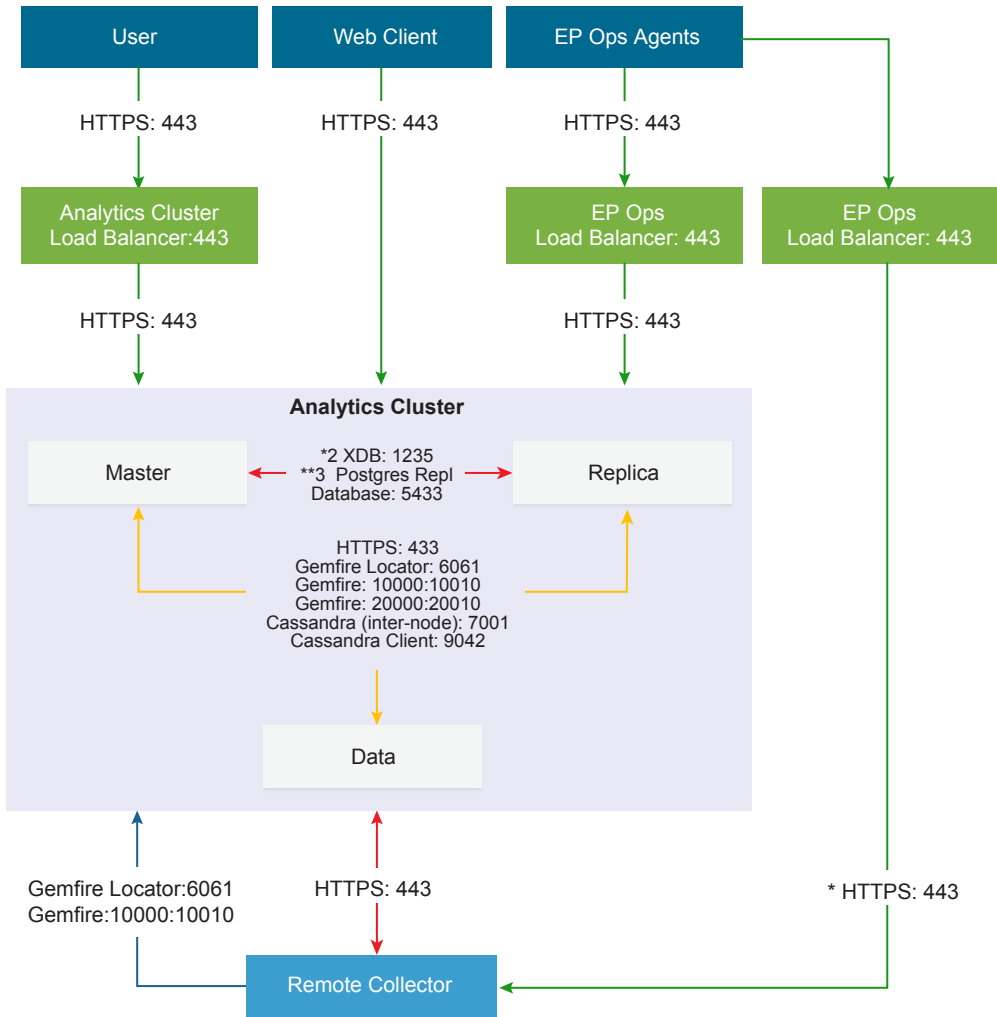
Component	Protocol	Port
Platform Services Controller	TCP	443
DNS	TCP, UDP	53
LDAP	TCP	389
LDAPS	TCP	636
GC TCP	TCP	3268, 3269
NTP	UDP	123
SMTP	TCP	25
SNMP	UDP	161
Adapters	TCP	**

** Ports required for adapters to communicate with external devices vary based upon the requirements of the device. Consult adapter documentation for required ports.

NOTE vROPS requires a TCP connection over HTTP via Port 10433 to connect to vSphere 5.x when retrieving inventory tag information.

NOTE The user interface and administrative interface to vROPS Operations Manager are through Port 443 with a TCP connection. See additional vROPS port information in the VMware vRealize Operations Manager 6.3 Information Center. Search on "How vRealize Operations Manager Uses Network Ports."

Port Requirements for vRealize Operations Manager



Protocols are not in the diagram.

* Required for upgrading from vRealize Operations Manager 6.0 to 6.1. The ports are closed after the upgrade.

** Required only for High Availability.

Small Deployment Profile for vRealize Operations Manager

9

The small deployment profile is intended for systems that manage up to 12,000 resources.

Virtual Appliance Name

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

Deployment Profile Support

The small deployment profile supports the following configuration.

- 12,000 resources
- 1,000 Endpoint Operations Management agents
- Data retention for six 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 Endpoint Operations Management agents against the analytics nodes.

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

Certificate

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

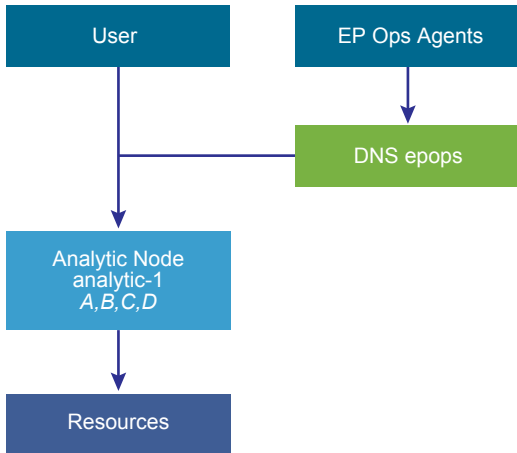
- DNS Name = `epops.refarch.local`
- DNS Name = `analytic-1.ra.local`

This is an example of a small deployment profile.

Table 9-1. Adapter Properties

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

vRealize Operations Manager Small Deployment Profile Architecture



Medium Deployment Profile for vRealize Operations Manager

10

The medium deployment profile is intended for systems that manage 40,000 resources, 20,000 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.

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

Deployment Profile Support

The medium deployment profile supports the following configuration.

- 40,000 total resources, 20,000 enabled for HA
- 6,000 Endpoint Operations Management agents
- Data retention for six 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*

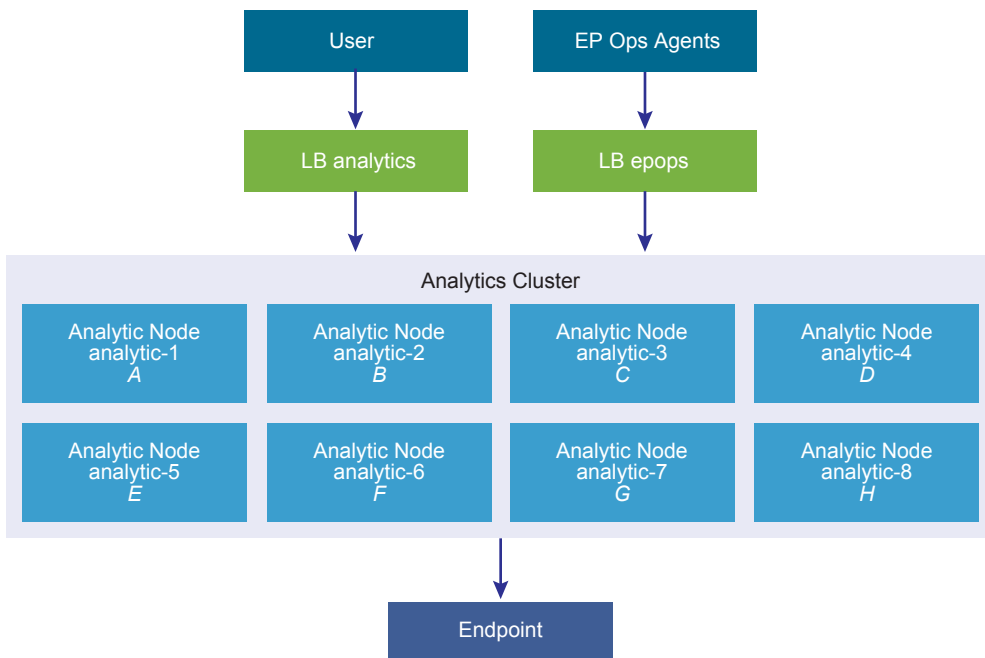
- DNS Name = *analytic-1.ra.local*

This is an example of a medium deployment profile.

Table 10-1. Adapter Properties

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

vRealize Operations Manager Medium Deployment Profile Architecture



Large Deployment Profile for vRealize Operations Manager

11

The large deployment profile is intended for systems that manage 80,000 resources, 40,000 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.

Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, large remote collectors for adapters, and large remote collectors for Endpoint Operations Management agents.

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

Deployment Profile Support

The large deployment profile supports the following configuration.

- 80,000 total resources, 40,000 enabled for HA
- 10,000 Endpoint Operations Management agents
- Data retention for six 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 = *analytic.refarch.local*

- DNS Name = *epops.refarch.local*
- DNS Name = *analytic-1.ra.local* to DNS Name = *analytic-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.

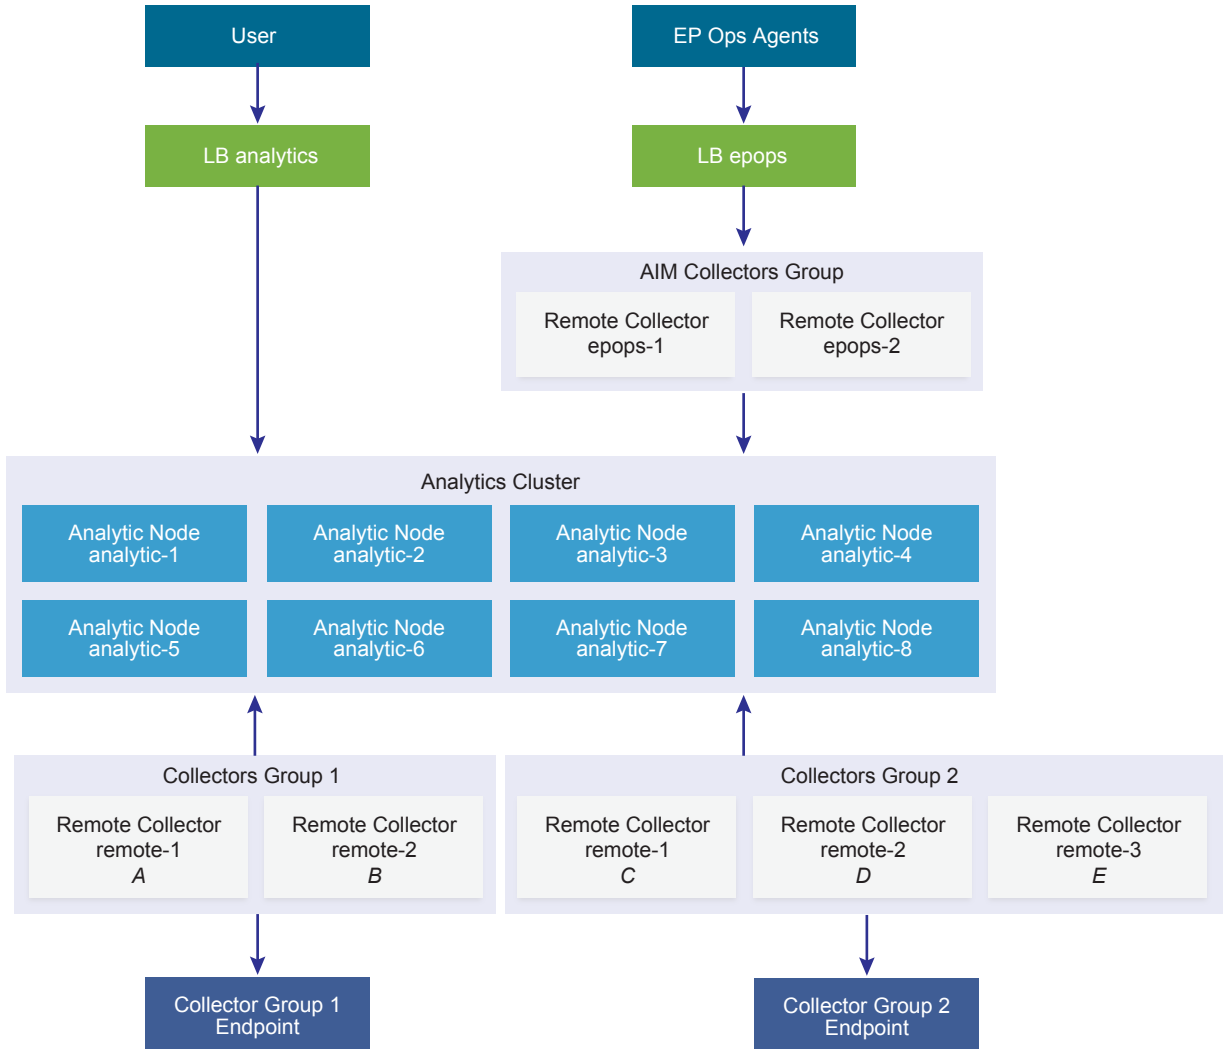
Table 11-1. Adapter Properties

Collector Group	Remote Collector	Adapter	Resources	Endpoint 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 10,000 resource for each remote collector.

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

vRealize Operations Manager Large Deployment Profile Architecture



Extra Large Deployment Profile for vRealize Operations Manager

12

The extra large deployment profile is intended for systems that manage 120,000 resources, 60,000 of which are enabled for High 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 16 large analytics nodes, X large remote collectors for adapters, and Y large remote collectors for Endpoint Operations Management agents.

- `analytic-1.ra.local`
- `analytic-2.ra.local`
- `analytic-3.ra.local`
- `analytic-4.ra.local`
- `analytic-5.ra.local`
- `analytic-6.ra.local`
- `analytic-7.ra.local`
- `analytic-8.ra.local`
- `analytic-9.ra.local`
- `analytic-10.ra.local`
- `analytic-11.ra.local`
- `analytic-12.ra.local`
- `analytic-13.ra.local`
- `analytic-14.ra.local`
- `analytic-15.ra.local`
- `analytic-16.ra.local`

Deployment Profile Support

- 120,000 total resources, 60,000 enabled for HA
- 10,000 Endpoint Operations Management agents
- Data retention for six 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 = *analytic.refarch.local*
- DNS Name = *epops-a.refarch.local*
- DNS Name = *epops-b.refarch.local*
- DNS Name = *analytic-1.ra.local* to *analytic-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*

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 12-1. Adapter Properties

Collector Group	Data Center	Remote Collector	Adapter	Resources	Endpoint 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
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

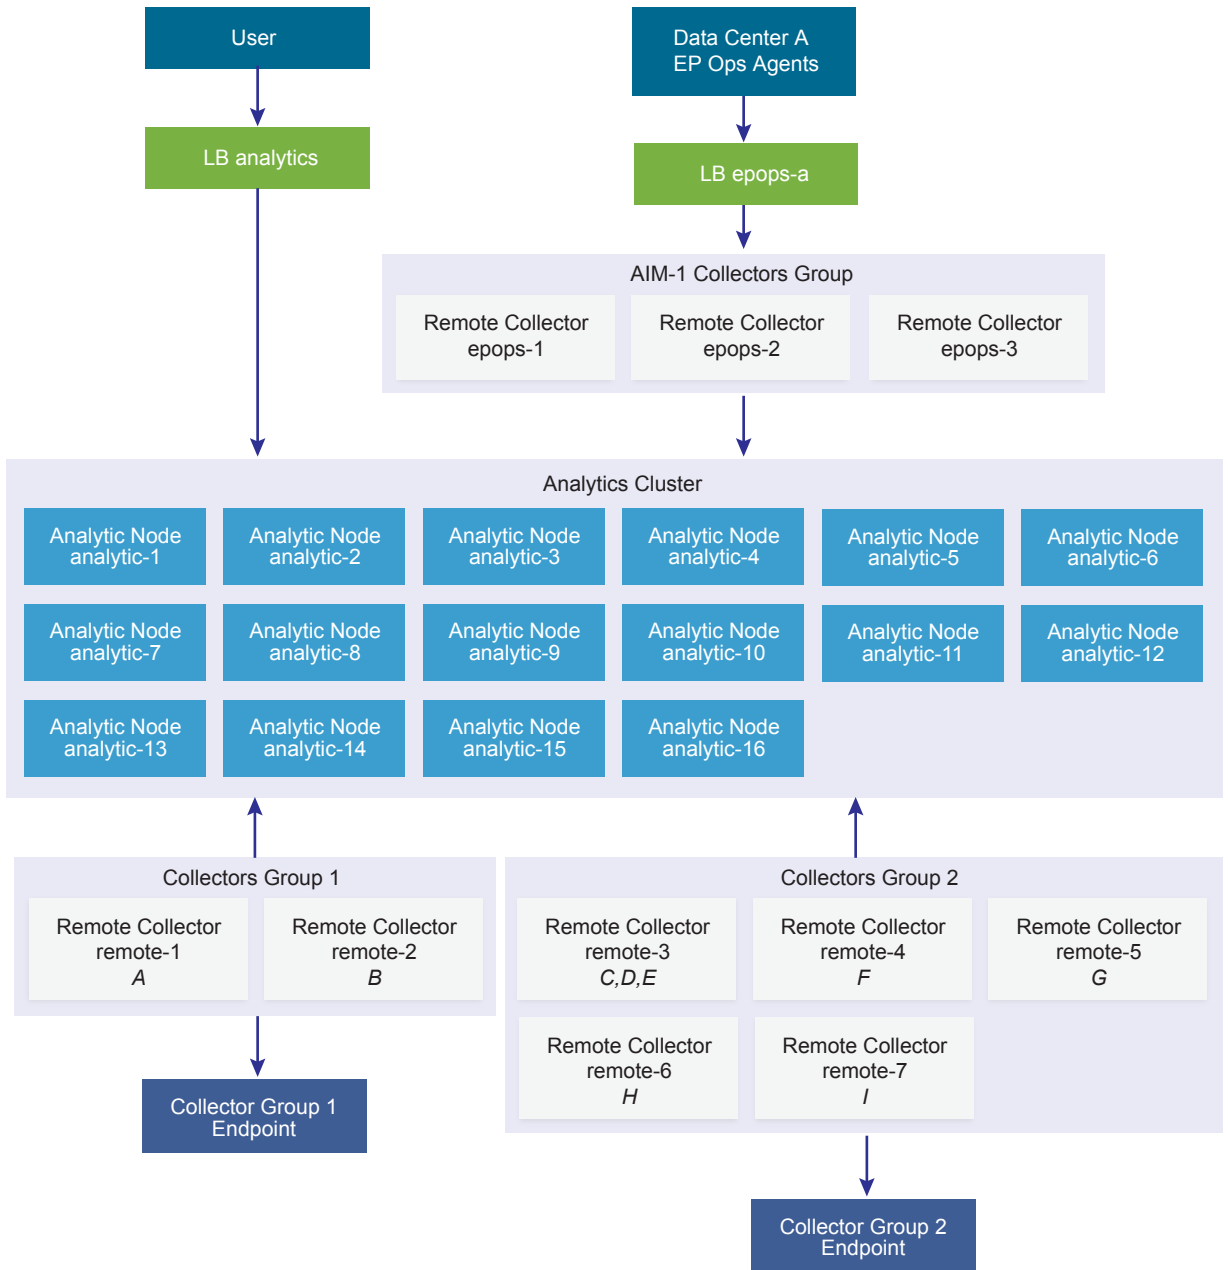
Table 12-1. Adapter Properties (Continued)

Collector Group	Data Center	Remote Collector	Adapter	Resources	Endpoint Operations Management agents
			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

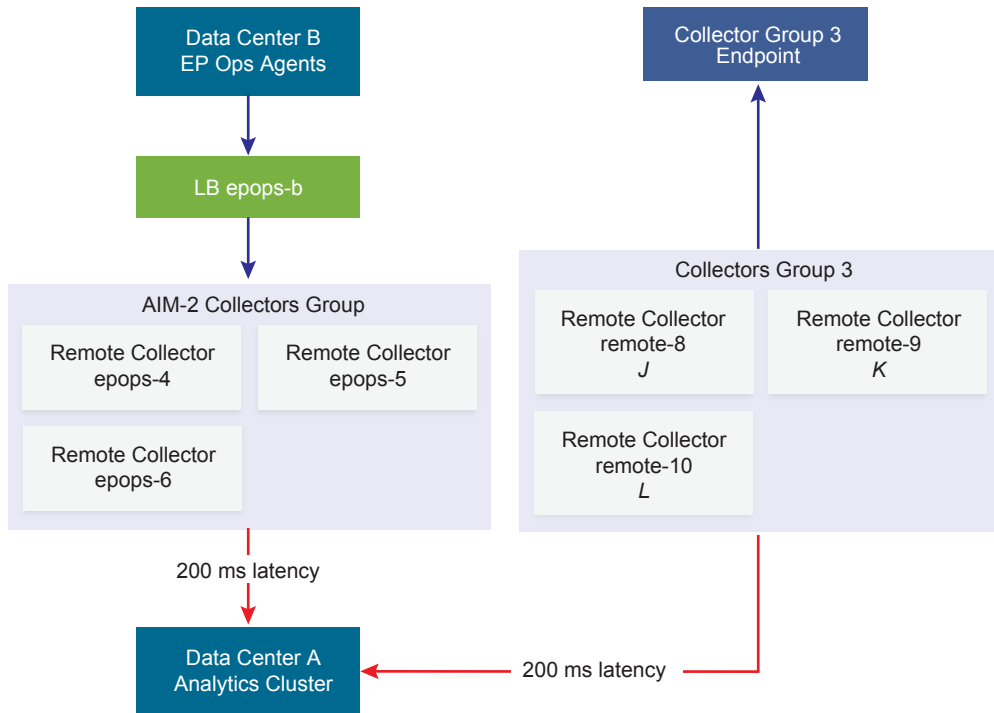
If a remote collector is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit of 10,000 resource for each remote collector.

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

vRealize Operations Manager Extra Large Deployment Profile Architecture - Data Center A



vRealize Operations Manager Extra Large Deployment Profile Architecture - Data Center B



Index

A

- adapters
 - best practices 7
 - considerations 9, 15
- analytics nodes
 - best practices 7
 - hardware requirements 17
 - scalability considerations 11
- authentication, considerations 9

B

- best practices
 - adapters 7
 - analytics nodes 7
 - management packs 7

C

- certificates, considerations 9
- considerations
 - adapters 15
 - High Availability 13
 - management packs 15

D

- data centers, considerations 9
- deployment profiles
 - extra large 31
 - large 27
 - medium 25
 - small 23

E

- environment, considerations 9
- extra large deployment profile 31

H

- hardware requirements
 - analytic nodes 17
 - remote collectors 17
- High Availability, considerations 13

L

- large deployment profile 27
- load balancers, considerations 9

M

- management packs
 - best practices 7
 - considerations 15
- medium deployment profile 25

O

- overview
 - architecture 5
 - deployment 5
 - scalability 5

P

- ports, requirements 19

R

- remote collectors
 - hardware requirements 17
 - scalability considerations 11

S

- scalability considerations
 - analytics nodes 11
 - remote collectors 11
 - storage 11
- sizing, considerations 9
- small deployment profile 23
- storage, scalability considerations 11

