

vRealize Suite 2017 Backup and Restore by Using EMC Avamar

vRealize Suite 2017



vmware®

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

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

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

docfeedback@vmware.com

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

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

Contents

vRealize Suite 2017 Backup and Restore by Using EMC Avamar	4
1 Backup and Restore Introduction	5
2 Preparing to Back Up vRealize Components	6
vRealize Components Backup Order	6
vRealize Business Preparations for Backing Up	7
vRealize Log Insight Preparations for Backing Up	8
vRealize Operations Manager Preparations for Backing Up	10
vRealize Orchestrator Preparations for Backing Up	13
vRealize Automation Preparations for Backing Up	13
3 Backing Up vRealize Suite by Using EMC Avamar	19
Perform an On-Demand Backup of a Client	19
Perform an On-Demand Group Backup	21
Schedule Backups	21
4 Restoring, Powering On, and Validating vRealize Suite	23
vRealize Suite Startup Order	23
vRealize Automation System Recovery	24
vRealize Orchestrator Restore Process	33
vRealize Operations Manager Restore Process	34
vRealize Log Insight Restore Process	39
vRealize Business Restore Process	45
5 Restore vRealize Suite by Using EMC Avamar	46
6 Restart the Cluster Analytic Service	49
7 Manual Backup Procedure Appears to Stall	50

vRealize Suite 2017 Backup and Restore by Using EMC Avamar

vRealize Suite 2017 Backup and Restore by Using EMC Avamar provides information about how to perform full, differential, incremental, and synthetic full backups leveraging VMware Changed Block Tracking to back up and restore VMware vRealize™ Suite by using EMC Avamar™.

Intended Audience

This information is intended for anyone who wants to back up and restore vRealize Suite 2017 components by using Avamar. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Backup and Restore Introduction

1

You can back up and restore the vRealize Suite components.

You can back up and restore the following versions of vRealize Suite components with EMC Avamar 7.4.

- vRealize Automation 7.3.0, 7.2.0, 7.1.0, and 7.0.1
- vRealize Orchestrator 7.3.0, 7.2.0, 7.1.0, and 7.0.1
- vRealize Operations Manager 6.6.0, 6.5.0, 6.4.0, 6.3.0, and 6.2.0
- vRealize Log Insight 4.5.0, 4.3.0 , 4.0.0, 3.6.0, 3.3.x
- vRealize Business 7.3.0, 7.2.0, 7.1.0, 7.0.1

When to Back Up Components

You should back up vRealize Suite components for the following reasons:

- To prepare for major maintenance of any of the system components
- To implement scheduled maintenance backup
- To prepare for updating certificates
- To protect certificates after updating them

Preparing to Back Up vRealize Components

2

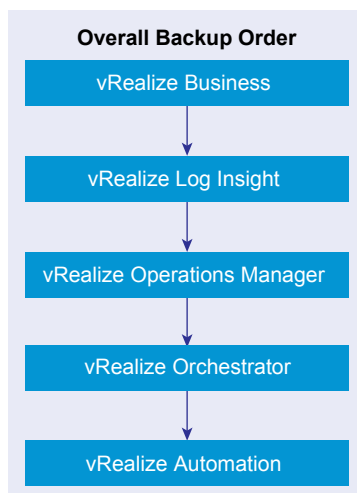
You can successfully prepare the vRealize product components for backup by planning and observing some basic guidelines.

This section includes the following topics:

- [vRealize Components Backup Order](#)
- [vRealize Business Preparations for Backing Up](#)
- [vRealize Log Insight Preparations for Backing Up](#)
- [vRealize Operations Manager Preparations for Backing Up](#)
- [vRealize Orchestrator Preparations for Backing Up](#)
- [vRealize Automation Preparations for Backing Up](#)

vRealize Components Backup Order

You should back up the VMs for vRealize product components in a specific order.



Depending on the vRealize product components that you have configured and your requirements, schedule backups for your vRealize product components in the following order. If you do not have a particular component, you can move to the next component in the specified order. The order of the VMs can be also defined in the backup tools.

- 1 vRealize Business. Back up the VM for vRealize Business.
- 2 vRealize Log Insight. If the resources are not an issue, it is recommended to enable concurrent cluster node backups to speed up the backup process. Back up all the nodes at the same time.
- 3 vRealize Operations Manager. All nodes are backed up at the same time.
- 4 vRealize Orchestrator. You can take backups of the vRealize Orchestrator VMs, individually and in cluster mode, in no particular order.
 - If vRealize Orchestrator is a standalone component, back it up before vRealize Automation components in no particular order.
 - If vRealize Orchestrator is embedded with in the vRealize Automation deployment, back it up as part of the vRealize Appliance.
- 5 vRealize Automation. Back up the vRealize Automation components in the following order:
 - a Proxy Agents
 - b DEM Workers
 - c DEM Orchestrator
 - d Manager Services
 - e Web sites
 - f vRealize Automation Appliances
 - g PostgreSQL, if applicable
 - h MS SQL

If you have multiple components on a VM, select the order considering the latter component on the VM from the list.

vRealize Business Preparations for Backing Up

To minimize system downtime and data loss when failures occur, administrators back up the vRealize Business Standard installation on a regular basis. If your system fails, you can recover it by restoring the last known working backup. You back up vRealize Business by exporting or cloning the virtual appliance and use backups to restore the virtual appliance.

vRealize Business Backup Order

vRealize Business Appliance

(Only one Node)

Guidelines for Backing Up


Use the following guidelines to plan backups:

- Verify that vRealize Business is up and running and vRealize Automation is registered with it.
- Verify that you can view the **Business Management** tab in your vRealize Automation deployment.
- Verify that vRealize Business is calculating the correct cost of the VMs.
- Verify that the VMs provisioned for vRealize Automation and vRealize Orchestrator are visible in vRealize Business and that vRealize Business can calculate the cost for the VMs.

vRealize Log Insight Preparations for Backing Up

You can perform full, differential, and incremental backups and restores of vRealize Log Insight VMs.

If resources are not a problem back up all the nodes at the same time, to speed up the backup process. Verify that you are increasing the number of concurrent backups from one, which is the default. Linear backup is also supported but it slows down the restore operation.



Guidelines for Planning Backups

You can use the following information for backing up vRealize Log Insight 3.3 clusters in a new environment.

- Ensure that you have no configuration problems on source and target sites before performing the backup and restore operations.
- During the backup operation, the memory usage can increase due to the vRealize Log Insight cluster usage. In some cases, the worker nodes might be disconnected for 1 to 3 minutes due to high memory usage. To reduce the memory throttling on vRealize Log Insight nodes, follow these guidelines:
 - Allocate additional memory over the vRealize Log Insight recommended configuration.
 - Schedule the recurring backups during off-peak hours.
- Disable quiesced snapshots, because vRealize Log Insight does not support them.
- The vRealize Log Insight server supports Linux and Windows agents.
 - If the agent configuration is created on the server side, a separate backup of the agent node is not required.


- If you use the agent nodes for more than installing the agent software and if these nodes need a full backup, follow the same backup procedure as for any VM.
- If the agent configuration is done on the client side, on the agents, and if the agent nodes are used only to install vRealize Log Insight agent software, scheduling a backup of the agent configuration file is sufficient. Back up the `liagent.ini` file and replace the file on the recovered agent or Linux or Windows machine with the backup file.
- If concurrent backup is not possible, ensure that the vRealize Log Insight master node is backed up first before the worker nodes. Worker nodes can be backed up in any order.
- Ensure that the backup frequency and backup types are selected based on the available resources and customer-specific requirements.
- Use the following guidelines when scheduling recurring backups.
 - For a reasonable loaded cluster setup, it might take a while for the first backup to finish irrespective of the tool.
 - The first backup is usually a full backup. Successive backups can be incremental or full backups, Successive backups finish relatively fast, compared to the first backup operation.
- Use static IP addresses for all nodes in a vRealize Log Insight cluster.
 - Using static IP addresses eliminates the need to update the IP addresses of vRealize Log Insight cluster nodes each time the IP address of a vRealize Log Insight node changes.
 - vRealize Log Insight includes all node IP addresses in each cluster node configuration file at `/storage/core/loginsight/config/loginsight-config.xml#<n>` where `<n>` is the largest number.
 - Some products that integrate with vRealize Log Insight to feed their logs, use a fully qualified domain name (FQDN) or IP address as the syslog target. For example, vSphere ESXi, vSphere, and vRealize Operations Manager use the nodes of the cluster master's or the load balancer's (if configured) FQDN or IP address as the syslog target.
- Use an FQDN for all nodes in the vRealize Log Insight cluster.
 - For the master node, when you use a load balancer, a fully resolvable FQDN is required. Otherwise, the ESXi hosts fail to feed the syslog messages to vRealize Log Insight or to any remote target.
 - Using an FQDN saves time on post-restore and recovery configuration changes, assuming that the same FQDN can be resolved on the recovery site.
 - For system alerts, vRealize Log Insight uses FQDN host names if available instead of IP addresses.
 - Assuming that only underlying IP addresses change post-backup and recovery or disaster recovery operations, using FQDN eliminates the need to change the syslog target address (master node FQDN or internal load balancer FQDN) on all the external devices feeding logs to the vRealize Log Insight cluster.

- With vRealize Log Insight 2.5, you must update the configuration file, located at `/storage/core/loginsight/config/loginsight-config.xml#<n>` where `<n>` is the largest number. This configuration file replaces the worker node IP address with the new IP address used for the restored nodes because the FQDN is not used for worker node addresses in the configuration file. You need to make this change only on the master node to synchronize the changes with all the worker nodes.
- Join requests from a vRealize Log Insight worker node should use the FQDN of the vRealize Log Insight master node.
 - Beginning in vRealize Log Insight 2.5, the master node host value in the configuration file on each of the nodes, located at `/storage/core/loginsight/config/loginsight-config.xml#<n>`, is based on the value used by the first worker node sending a join request. Using the FQDN of the master node for the join request prevents making any manual changes to the master node host value post-disaster recovery. Otherwise, the worker nodes cannot rejoin the master node until the master node host name is updated in the configuration files on all restored cluster nodes.
- Provide static IP addresses as well as optional virtual IP addresses for the load balancer.
 - When configuring an integrated load balancer, provide the optional FQDN for the virtual IP address. This optional FQDN enables vRealize Log Insight to revert to the FQDN when an IP address is not reachable for any reason.

vRealize Operations Manager Preparations for Backing Up

To minimize vRealize Operations Manager downtime and data loss if a failure occurs, back up on a regular basis. If your system fails, you can recover it by restoring to the last full or incremental backup.

You can backup and restore vRealize Operations Manager single-node or multi-node clusters by using backup tools. You can perform full or incremental backups and restores of virtual machines.



Note All nodes are backed up and restored at the same time. You cannot back up and restore individual nodes.

Guidelines for Planning Backups

Verify that the following prerequisites are met before you back up vRealize Operations Manager systems by using any tool.

- Do not quiesce the file system.
- Use a resolvable host name and a static IP address for all nodes.
- All nodes must be powered on and accessible during backup.
- Back up the entire VM. You must back up all VMDK files that are part of the virtual appliance.
- Do not stop the cluster while performing the backup.

Verify that the following prerequisites are met for the vCenter server where vRealize Operations Manager is deployed.

- The vSphere Data Protection appliance is installed on the vCenter Server instance where the vRealize Operations Manager cluster is deployed.
- You have sufficient disk space available for your vSphere Data Protection instance. This depends on the size of the multi-node cluster that you want to back up
- Use the vSphere Web Client to log in as an administrator to the vCenter Server instance that manages your environment.
- In the vSphere Web Client verify that the virtual machines have the latest version of VMware Tools installed.

Note Do not perform a backup while dynamic threshold (DT) calculations are running because this backup might lead to performance problems or loss of nodes. DT calculations run at 2 a.m. by default. The initial backup might take longer to complete depending on the cluster size and number of nodes, so you should turn off the DT. Schedule the differential or incremental backups so that they end before the DT calculations begin.

If you are using backup tools, such as NetBackup and vSphere Data Protection, and have manually created vSphere snapshots of any of the VMs to be backed up, be aware that the tools delete all existing snapshots at the time of the backup or restore.

Common Backup Scenarios

The common backup scenarios for vRealize Operations Manager systems include a full backup of a single node virtual appliance system and a full backup of a multiple node virtual appliance cluster.

Single-Node Virtual Appliance

This scenario backs up a single-node system on the same host.

- 1 Assign a static IP address.
- 2 Ensure that the power is on for the entire backup process.

- 3 If the system is a Linux or Windows installation, you need to prepare the system before you start the backup. For more information, see *Preparing for vRealize Operations Manager Installation* in the VMware vRealize Operations Manager Documentation Center.

Multiple-Node Virtual Appliance Clusters

This scenario backs up a multiple-node virtual appliance clusters.

- 1 Assign a static IP address for each node.
- 2 Ensure that all nodes that are configured with high availability are accessible and are enabled for high availability.
- 3 Ensure that the power remains on during the entire backup process.
- 4 If the system is a Linux or Windows installation, you must prepare the system before you start the backup. For more information, see *Preparing for vRealize Operations Manager Installation* in the VMware vRealize Operations Manager Documentation Center.

Disable Quiescing

Before you backup your vRealize Operations Manager multi-node cluster by using vSphere Data Protection, disable quiescing of the file system.

Procedure

- 1 Log in to the ESXi host with an SSH session, and power off all nodes.
- 2 Navigate to the `/vmfs/volumes/virtual_machine_datastore/ virtual_machine/` directory, and open the `virtual_machine.vmx` file for editing.
- 3 Set the `disk.EnableUUID` parameter to `false`.
You may have to add the `disk.EnableUUID` parameter to the `virtual_machine.vmx` file.
- 4 Save and close the file.
- 5 Power on all nodes.
- 6 Open a console session to the virtual machine, and log in to each node.
- 7 Navigate to the `/etc/vmware-tools` directory, and open the `tools.conf` file for editing.
If you cannot locate the `tools.conf` file, run the `vi tools.conf` command to create a file.
- 8 Add these parameters to the file.

```
[vmbackup]
enableSyncDriver = false
```

This runs a synchronization operation before the snapshot, and does not run a freeze on the file system.

- 9 Save and close the file.

What to do next


Backup your vRealize Operations Manager multi-node cluster by using vSphere Data Protection.

vRealize Orchestrator Preparations for Backing Up

You can backup your vRealize Orchestrator VMs in no particular order. You can also back up the vRealize Orchestrator elements that you modified.

If vRealize Orchestrator is embedded within the vRealize Automation appliance, perform its backup along with Manager Services for vRealize Automation.

If vRealize Orchestrator is a standalone component, perform its backup before backing up vRealize Automation components in no particular order.



Each vRealize Orchestrator server instance has unique certificates, and each vCenter Server plug-in instance has a unique ID. The certificates and the unique ID identify the vRealize Orchestrator server and the vCenter Server plug-in. If you do not back up the vRealize Orchestrator elements or export the vRealize Orchestrator configuration for backup purposes, make sure that you change these identifiers.

All components of the vRealize Orchestrator must be backed up together and at the same time including the database components. You must back up the vRealize Orchestrator database and VMs (custom workflows and packages).

vRealize Orchestrator Database

You can take full database backups of the database in your environment before a full VM backup. The main purpose is to ensure consistency of the data when you have to restore.

Follow your in-house procedures to back up the vRealize Orchestrator database outside of the vRealize Suite framework.

vRealize Automation Preparations for Backing Up

A system administrator backs up the full vRealize Automation installation on a regular basis. Plan the backup around efficiencies and periods of low activity.

You can use several strategies, singly or in combination, to back up vRealize Automation system components. For virtual machines (VMs), you can use the Snapshot function to create snapshot images of critical components or use tools like Symantec NetBackup and vSphere Data Protection. If a system failure occurs, you can use these images to restore components to their state when the images were created. Alternatively, and for non-virtual machine components, you can create copies of critical configuration files for system components, which can be used to restore these components to a customer configured state following reinstallation.

A complete backup includes the following components:

- IaaS components
 - Proxy Agents
 - DEM Workers
 - DEM Orchestrator
 - Manager Services
 - Web sites
- vRealize Automation appliance
- PostgreSQL database. Applicable only for legacy installations that do not use an embedded appliance database
- Infrastructure MSSQL database
- (Optional) Software load balancers
- (Optional) Load balancers that support your distributed deployment. Consult the vendor documentation for your load balancer for information about backup considerations

Guidelines for Planning Backups


Use the following guidelines to plan backups:

- When you back up a complete system, back up all instances of the vRealize Automation appliance and databases as near simultaneously as possible, preferably within seconds.
- Minimize the number of active transactions before you begin a backup. Schedule your regular backup to when your system is least active.
- Back up all databases at the same time.
- When you back up the vRealize Automation appliance and the IaaS components, disable in-memory snapshots and quiesced snapshots.
- Create a backup of instances of the vRealize Automation appliance and the IaaS components when you update certificates.

vRealize Automation Backup Order

You must back up the VMs for vRealize Automation components in a specific order.

When you use backup tools to schedule backups for your vRealize Automation deployment, you must back up the components in the following order:



If you have multiple components on a VM, select the order considering the latter component on the VM from the list.

If vRealize Orchestrator is a standalone component, then it should be backed up in no particular order before vRealize Automation components.

Backing Up vRealize Automation Certificates

A system administrator backs up certificates and certificate chains at installation time or when a certificate is replaced.

Back up the following certificates:

- vRealize Automation appliance certificates and the entire corresponding certificate chain.
- IaaS certificates and the entire corresponding certificate chain.

Backing Up vRealize Automation Databases

The database administrator backs up the Infrastructure MSSQL Server and Appliance Database.

As a best practice, back up the Infrastructure MSSQL and Appliance Database or legacy PostgreSQL databases as nearly simultaneously as possible to prevent or minimize data loss. Also, when applicable, back up databases with Point-in-Time enabled. By using Point-in-Time recovery, you ensure that the two databases are consistent with each other.

Note If only one database fails, you must restore the running database to the most recent backup so that the databases are consistent.

Infrastructure MSSQL Database

Back up the Infrastructure MSSQL and Appliance Database or legacy PostgreSQL databases as nearly simultaneously as possible to prevent or minimize data loss. Also, when applicable, back up databases with Point-in-Time enabled. By using point-in-time recovery, you ensure that the two databases are consistent with each other. If only one database fails, you must restore the running database to the most recent backup so that the databases are consistent.

Follow your in-house procedures to back up the Infrastructure MSSQL database outside of the vRealize Suite framework.

Use the following guidelines when creating a backup:

- If possible, check that all IaaS workflows are complete and that all IaaS services are stopped or that activity is minimized.
- Back up with Point-in-Time enabled.
- Back up the MSSQL database at the same time that you back up the other components.
- Back up the passphrase for your database.

Note Your database is protected by a passphrase. Have the passphrase available when you restore the database. Typically, you record the passphrase in a safe and accessible location at install time.

Appliance Database or Legacy PostgreSQL Database

If you are using an Appliance Database or a legacy PostgreSQL database embedded in a vRealize Automation appliance, you can back up the database by backing up the entire appliance.

If you are using a standalone legacy PostgreSQL appliance, you must back up the appliance.

If you are using a legacy PostgreSQL database, you can also backup the database separately. For more information, see the VMware Knowledge Base article *Migrating from external vPostgres appliance to vPostgres instance located in the vCAC appliance (2083562)* at <http://kb.vmware.com/kb/2083562> for more information.

Backing Up Load Balancers

Load balancers distribute work among servers in high-availability deployments. The system administrator backs up the load balancers on a regular basis at the same time as other components.

Follow your site policy for backing up load balancers, keeping in mind the preservation of network topology and vCloud Automation Center backup planning.

Backing Up IaaS Components

The system administrator backs up the IaaS components in a specific order.

When you back up IaaS components, disable in-memory snapshots and quiesced snapshots.

You can back up IaaS components by taking a snapshot of the VMs in the following order:

- Proxy Agents and DEMs
- Manager Service
- Web sites

Backing Up vRealize Automation Appliance

You must back up secondary vRealize Automation appliance nodes, followed by the master node.

When you back up the vRealize Automation appliance, disable in-memory snapshots and quiesced snapshots.

The primary vRealize Automation appliance instance contains the writeable Appliance Database, if applicable, and is the last appliance that you shut down in an ordered shutdown procedure.

(Optional) Shut Down vRealize Automation

Shutting down vRealize Automation is optional. If you decide to shut down your vRealize Automation system, use the specified order.

If you are using vCenter Server to manage your virtual machines, use the guest shutdown command to shut down vRealize Automation.

Procedure

- 1 Shut down the DEM Orchestrator and Workers and all vRealize Automation agents in any order and wait for all components to finish shutting down.
- 2 Shut down the VMs that are running the Manager Service and wait for the shutdown to finish.
- 3 (Optional) For distributed deployments, shut down all secondary Web nodes and wait for the shutdown to finish.
- 4 Shut down the primary Web node, and wait for the shutdown to finish.
- 5 (Optional) For distributed deployments, shut down all secondary vRealize Appliance instances and wait for the shutdown to finish.

- 6 Shut down the primary vRealize Appliance and wait for the shutdown to finish.

If applicable, the primary vRealize Appliance is the one that contains the master, or writeable, Appliance Database. Make a note of the name of the primary vRealize Appliance. You use this information when you restart vRealize Automation.

- 7 If you are using a legacy standalone PostgreSQL database, also shut down that machine.

You have shut down your vRealize Automation deployment.

Backing Up vRealize Suite by Using EMC Avamar

3

You can back up vRealize Suite components by using EMC Avamar to perform on-demand backups and schedule backups.

You can take on-demand backups and schedule Avamar to perform scheduled backups. Scheduled backups run automatically to ensure that backups occur on an ongoing basis. You can schedule backups to run daily, weekly, or monthly. For information on scheduling backups in Avamar, see the [EMC Avamar for VMware User Guide](#).

This section includes the following topics:

- [Perform an On-Demand Backup of a Client](#)
- [Perform and On-Demand Group Backup](#)
- [Schedule Backups](#)

Perform an On-Demand Backup of a Client

You can perform an on-demand backup of an individual client. If you configure scheduled backups for a group of clients, then you can also perform an on-demand backup of a group or an on-demand back of a single client by using group policy settings.

Prerequisites

- Deploy and configure the Avamar appliance. See the [EMC Avamar for VMware User Guide](#).
- Verify that all nodes are powered on and are accessible during backup.
- Add all vRealize Suite components to backup as clients. See the [EMC Avamar for VMware User Guide](#).
- Disable quiesced snapshots.

Procedure

- 1 In Avamar Administrator, click **Backup & Restore**.
- 2 Click the **Backup** tab.

The top left pane contains a list of domains.

- 3 Select a domain in the upper tree, and select a virtual machine client, VMware folder, resource pool, or vApp in the lower tree.

- 4 In the **Browse for File, Folders, or Directories** pane, select the data to back up.

Option	Description
Back up the entire image	Right-click the top (root) folder, and select Select All .
Back up specific virtual disks	Select one or more disks to back up only the selected virtual disks.

- 5 In the **Actions** menu, select **Backup Now**.

The **On Demand Backup Options** dialog box appears.

- 6 Select the backup retention setting.

Option	Description
Retention period	Automatically delete this backup from the Avamar server after a specific amount of time. Specify the number of days, weeks, months, or years for the retention period.
End date	Automatically delete this backup from the Avamar server on a specific calendar date. Browse to the date to delete this backup on the calendar.
No end date	Keep this backup while this client remains active in the Avamar server.

- 7 From the Avamar encryption method list, select the encryption method to use for data transfer between the client and the Avamar server during the backup.

- 8 (Optional) Select a proxy to perform backup.

The default setting is automatic, which enables the Avamar server to select the best proxy for this operation.

- 9 Click **More Options**.

The **Backup Command Line Options** dialog box appears.

- a Select the plug-in type.

In most cases, the system default settings are the optimum settings for on-demand image backups.

- b (Optional) To increase performance, select **Use Changed Block Tracking (CBT)**.

- c To disable the quiesced snapshots, click **More**.

- d Enter the attribute and attribute value, and click **+** to add the flag.

Option	Value
Attribute	[avvcbimage]quiesce_fs
Attribute Value	false

- 10 Click **OK**.

The **On Demand Backup Options** dialog box closes and the following status message appears:
Backup initiated.

- 11 Click **OK**.

What to do next

For information on monitoring and cancelling backup procedures, see the [EMC Avamar for VMware User Guide](#).

Perform and On-Demand Group Backup

You can use EMC Avamar to back up an entire group of clients.

Prerequisites

- Deploy and configure the Avamar appliance. See the [EMC Avamar for VMware User Guide](#).
- Verify that all nodes are powered on and are accessible during backup.
- Add all vRealize Suite components to backup as clients. See the [EMC Avamar for VMware User Guide](#).
- Disable quiesced snapshots.

Procedure

- 1 In Avamar Administrator, click **Policy**.
The **Policy** window appears.
- 2 Click the **Policy Management** tab.
- 3 Select the group to back up by clicking the **Groups** tab and selecting the group from the list.
- 4 Click **Back Up** and click **OK**.

What to do next

For information on monitoring and cancelling backup procedures, see the [EMC Avamar for VMware User Guide](#).

Schedule Backups

Scheduled backups run automatically to ensure that backups occur on an ongoing basis.

A scheduled backup can include multiple clients or a single server.

Prerequisites

- Deploy and configure the Avamar appliance. See the [EMC Avamar for VMware User Guide](#).
- Verify that all nodes are powered on and are accessible during backup.
- Add all vRealize Suite components to backup as clients. See the [EMC Avamar for VMware User Guide](#).

Procedure

- 1 Create a dataset to define the data that is included in the backups.

- 2 Disable quiesced snapshots.
 - a Select the plug-in type.
 - b For each selected plug-in type, click **More**.
 - c Enter the attribute `[avvcbimage]quiesce_fs` and attribute value `false`, and click **+** to add the flag.
- 3 Create a schedule for when the backups should occur.
- 4 Create a retention policy to define how long to keep the backups in the system.
- 5 Create a group for the backups.
 - a Assign the new dataset to the new group.
 - b Assign a schedule to the new group.
 - c Assign a retention policy to the new group.
 - d Add one or more clients to the new group.
 - e Assign the proxies to the new group.
- 6 Enable scheduling for the group.

What to do next

For information on monitoring and cancelling backup procedures, see the [EMC Avamar for VMware User Guide](#).

Restoring, Powering On, and Validating vRealize Suite

4

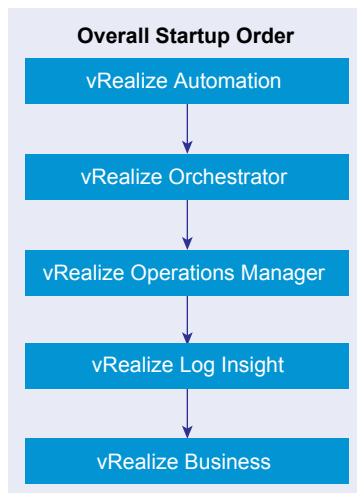
You can restore, power on, and validate that the vRealize Suite components are restored correctly by using the following information.

This section includes the following topics:

- [vRealize Suite Startup Order](#)
- [vRealize Automation System Recovery](#)
- [vRealize Orchestrator Restore Process](#)
- [vRealize Operations Manager Restore Process](#)
- [vRealize Log Insight Restore Process](#)
- [vRealize Business Restore Process](#)

vRealize Suite Startup Order

You should start up the VMs for vRealize Suite components in a specific order.



Depending on the vRealize Suite components that you have backed up, restore your vRealize Suite components in the specified order. If you have taken any snapshots, you must remove the snapshots before you restore.


- 1 vRealize Automation. Start up the vRealize Automation components in the following order:
 - a MS SQL
 - b PostgreSQL, if applicable
 - c vRealize Automation Appliances
 - d Websites
 - e Manager Services
 - f DEM vRealize Orchestrator
 - g DEM Workers
 - h Proxy Agents
- 2 vRealize Orchestrator. You can take start up the vRealize Orchestrator VMs in no particular order.
- 3 vRealize Operations Manager. All nodes can start at the same time.
- 4 vRealize Log Insight. Start up the master node, followed by VMs with the worker nodes in any order.
- 5 vRealize Business. Start up the VM for vRealize Business.

vRealize Automation System Recovery

A system administrator uses backups to restore vRealize Automation to a functional state after a system failure. If IaaS components such as Manager Service machines fail, you must reinstall them.

If you restore from a backup, machines that were provisioned after the backup still exist, but are not managed by vRealize Automation. For example, they do not appear in the items list for the owner. Use the Infrastructure Organizer to import virtual machines and bring them back under management.

You must start up your vRealize Automation components in a specified order.



Restoring vRealize Automation Databases

A system administrator restores the IaaS MSSQL database and the PostgreSQL database.

Recover a database in the following situations:

- If both databases fail, restore them from the last known time when both databases were backed up.
- If one database fails, restore it and revert the functional database to the version that was in use when the backup used to restore the failed database was created.

The backup time for each database can differ. The greater the gap between the last working time of the databases, the greater the potential for data loss.

You should back up full VMs of databases, instead of backing up PostgreSQL database directly. For information about how to restore a PostgreSQL database, see the VMware Knowledge Base article [Migrating from external vPostgres appliance to a vPostgres instance located in the vCAC appliance \(2083562\)](#).

Database Passphrases

IaaS MSSQL database security requires a security passphrase to generate an encryption key that protects the data. You specify this passphrase when you install vRealize Automation.

If you lose the passphrase, or want to change the passphrase, consult VMware technical support for more information.

Restoring the vRealize Automation Appliance and Load Balancer

If a failure occurs, a system administrator restores the vRealize Automation appliance. If a load balancer is used, the administrator restores the load balancer and the virtual appliances that it manages. For vRealize Automation 7.0, you cannot change the host names during restoration.

You might need to restore a failed virtual appliance in the following circumstances:

- You are running a minimal deployment and your only vRealize Automation appliance fails or becomes corrupted.
- You are running a distributed deployment and some, but not all, virtual appliances fail.
- You are running a distributed deployment and all virtual appliances fail.

How you restore a vRealize Automation appliance or virtual appliance load balancer depends on your deployment type and on which appliances failed.

- If you are using a single virtual appliance whose name is unchanged, restore the virtual appliance, or redeploy it and restore a set of backed up files. No further steps are required.
- If you are running a distributed deployment that uses a load balancer, and you change the name of the virtual appliance or the virtual IP address of the load balancer, you must redeploy the appliance and restore its backed up VMs or files. Also, you must regenerate and copy certificates for your deployment.

If you are redeploying, reconfiguring, or adding virtual appliances to a cluster, see *Installation and Configuration* documentation for vRealize Automation appliance in the [vRealize Automation product documentation](#) for more information.

Restoring the IaaS Website, Manager Services, and Their Load Balancers

A system administrator restores the IaaS Website and Manager Service and their associated load balancers.

Restoring the DEM Orchestrator and the DEM Workers

If a failure occurs, a system administrator restores all DEMs from the available backups.

Restoring the IaaS Agents

The system administrator restores all the IaaS agents from the available backups.

If you reinstall vSphere agents, use the same endpoint name used at backup.

Restore vRealize Automation in an Isolated Environment

A system administrator can use backups to restore vRealize Automation in an environment that is isolated from other networks.

Prerequisites

- Verify that the environment has a DNS server with identical zones and records to the production server and that the vRealize Automation virtual appliance and Windows and IaaS servers are configured to work with it.

Procedure

- 1 Configure the vRealize Automation virtual appliance.
 - a Check the vRealize Automation virtual appliance PostgreSQL status, and reset secondary node if needed.
 - b If the primary node is unable to start, reset RabbitMQ on the primary node.
 - c Shut down the secondary node.
 - d Restart the primary node.
 - e Wait 30 seconds, and restart the secondary node.
 - f Verify that the primary server has all services except IaaS as **Registered**.
 - g If needed, perform a join cluster operation on the secondary node.
 - h Verify that both servers have all services except IaaS as **Registered**.

If the environment uses Active Directory, you can skip all following steps.

2 Configure IaaS machine authentication.

Option	Description
Using a local administrative account	<p>This option is the most secure and simulates a production environment with an Active Directory and Integrated Authentication to the SQL server, so tests can be closer to reality.</p> <p>Create a user with administrative privileges and with the same user name and password on all servers, including the SQL server. Grant this user sysadmin rights on the SQL Server and owner rights on the IaaS database. You can also modify the connection string to reflect this new user or configure it to use SQL authentication.</p>
Using the local system account	<p>This option is less secure and might cause problems with the manager service authenticating to the Web repository.</p> <p>Start all services with the local system account and grant the local system accounts on all servers sysadmin rights on the SQL Server and owner rights on the IaaS database.</p>

3 Configure the Web IaaS machine.

- a Check connection string in the **web.config** file of the Model Manager Web installation folder and modify it according to the account you used for IaaS machine authentication.

If you are using SQL authentication, it should look similar to the following:

```
Max Pool Size=200;MultipleActiveResultSets=True;Connect Timeout=200;User
Id=sa;Password=password;" />
```

- b Change the IIS IaaS application pools to use the account you used for IaaS machine authentication.
- c Restart the IIS services by running `iisreset` in command prompt on the IaaS Web servers.

4 Configure the Manager Service IaaS machine.

- a If the machine running the SQL server is part of a failover cluster, disable the cluster service and reboot the machine.
- b If the SQL server is part of a failover cluster, disable the cluster service and reboot the SQL server.

If the machine running the SQL server is part of a failover cluster, the cluster service does not work because there is no Active Directory. In a failover cluster, the MSDTC service is controlled by the cluster service. When the Active Directory domain is missing, the cluster service is not able to bring up its resources and keeps the DTC in a locked state. When you disable the cluster service and reboot the machine running the SQL server, this unlocks the local DTC, and the Manager Service can successfully initialize.

- c Enable **NoAuthenticationRequired** on the DTCs of both the SQL server and the manager server from their respective component services console.

5 On the DEM Orchestrator IaaS machine, change the service account to the account you used for IaaS machine authentication.

- 6 Change the management agent service account on all IaaS machines to the account you used for IaaS machine authentication.

(Optional) Change the IP Addresses of vRealize Automation Appliances After Restore

You should perform these steps, only if you have changed the IP addresses of the vRealize Automation appliances after the restore of vRealize Automation has started. This is a workaround in order to get the vRealize Automation up and running after the restore, if you have changed the IP addresses during the restore process.

If you have changed the IP addresses of the vRealize Automation appliances during the restore process, you must perform the following steps in order to have the vRealize Automation operational after the restore.

Procedure

- 1 Change the DNS server so that the FQDNs of all the vRealize Automation machines points to the new IP addresses.
- 2 Add the new IP addresses in the load balancer pools, such as virtual appliances, web, and manager.
- 3 Update the `/etc/hosts` file for each of the vRealize Automation VA nodes with the new IP addresses.
- 4 Update the `/etc/sysconfig/elasticsearch` for each of the vRealize Automation VA nodes with the new IP addresses.
- 5 Restart all the vRealize Automation VA Nodes.

You should now be able to startup vRealize Automation.

Start Up vRealize Automation

When you start vRealize Automation from the beginning, such as after a power outage, a controlled shutdown or after recovery, you must start its components in a specified order.

Prerequisites

Verify that the load balancers that your deployment uses are running.

Procedure

- 1 Start the MS SQL database machine. If you are using a legacy PostgreSQL standalone database, start that machine as well.
- 2 (Optional) If you are running a deployment that uses load balancers with health checks, disable the health check before you start the vRealize Automation appliance. Only ping health check should be enabled.

- 3 Start all instances of vRealize Automation appliance at the same time and wait for approximately 15 minutes for the appliances to startup. Verify that the vRealize Automation appliance services are up and running.

If you have more than one node and you start only one node, you may have to wait for extra 35 minutes. However, the extra wait time would be canceled out as soon as you start the second node.

- 4 Start the primary Web node and wait for the startup to finish.
- 5 (Optional) If you are running a distributed deployment, start all secondary Web nodes and wait 5 minutes.
- 6 Start the primary Manager Service machine and wait for 2 to 5 minutes, depending on your site configuration.
- 7 (Optional) If you are running a distributed deployment, start secondary Manager Service machines and wait 2 to 5 minutes.

On secondary machines, do not start or run the Windows service unless you are configured for automatic Manager Service failover.

- 8 Start the Distributed Execution Manager Orchestrator and Workers and all vRealize Automation proxy agents.

You can start these components in any order and you do not need to wait for one startup to finish before you start another.

- 9 If you disabled health checks for your load balancers, reenable them.

- 10 Verify that the startup succeeded.

- a Open a Web browser to the vRealize Automation appliance management interface URL.
- b Click the **Services** tab.
- c Click the **Refresh** tab to monitor the progress of service startup.

When all services are listed as registered, the system is ready to use.

Validate vRealize Automation

After the recovery is complete, you must validate that the environment and application functionality is restored.

Procedure

- 1 Verify that the vRealize Automation services are registered on both virtual appliances.
 - a Log in to the vRealize Automation Console on <https://<IP of VM>:5480/> by using root credentials.
 - b Click the **Services** tab to verify that the vRealize Automation services have the status of Registered and the IaaS service is running.

- 2 Verify that the PostgreSQL replication is connected.
 - a Log in to the vRealize Automation Console on `https://<IP of VM>:5480/` by using root credentials.
 - b Click the **vRA Settings** tab and click **Database**.
 - c Verify that the **Connection Status** is **CONNECTED**.
- 3 Verify that RabbitMQ is registered correctly on all virtual appliances.
 - a Click vRealize Automation **Settings > Messaging** tab.
 - b Verify that all virtual appliances have a **Running** status of Yes.
 - c Verify that the **Connection Status** is **CONNECTED**.If RabbitMQ fails to register correctly, see [KB 2106969](#) to reset RabbitMQ.
- 4 Verify virtual appliance and IaaS service last connected status.
 - a Click **vRealize Automation Settings** and click **Cluster**.
 - b Verify that each virtual appliance node has a **Last Connected** time of 9 minutes or less.
 - c Verify that each IaaS node has a **Last Connected** time of three minutes or less.
 - a Verify that all services are in sync.
- 5 On the node where you installed Web IaaS, verify that there are no errors displayed in **Event Viewer > Administrative Events**.
- 6 Verify that the Web IaaS service is initialized.
 - a In your browser, navigate to `https://localhost/wapi/api/status`.
 - b Verify that the displayed file contains the string `"serviceInitializationStatus":"REGISTERED"`.
- 7 Verify that the Model Manager Web log repository contains no exceptions.
 - a Navigate to `C:\Program Files (x86)\VMware\VCAC\Server\Model Manager Web\Logs\Repository`.
 - b Verify that the file contains no exceptions.
 - c Verify that the file contains the following strings:

```
Model TrackingModel service has been initialized.
Model MetaModel service has been initialized.
Model CoreModel service has been initialized.
```
- 8 On the node where you installed the manager service, verify that the VMware vCloud Automation Center service is running.
- 9 Verify that the manager service log contains no errors.
 - a Navigate to `C:\Program Files (x86)\VMware\VCAC\Server\Logs\All.log`.
 - b Verify that there are no exception errors after the string `Manager service started for processed workflows`.

- 10 In your browser, navigate to `https://localhost/VMPS2Proxy`, and verify that it contains an XML file similar to the following:

```
- <wsdl:service name="ProxyAgentService">
  - <wsdl:port name="BasicHttpBinding_VMPSProxyAgent"
binding="tns:BasicHttpBinding_VMPSProxyAgent">
  <soap:address location="https://us06-2-ims1.vmware.com/VMPS2proxy"/>
```

- 11 On the node where you installed DEM Orchestrator, verify that the DEM Orchestrator services are running.
- 12 Navigate to `C:\Program Files (x86)\VMware\VCAC\Distributed Execution Manager\DEMOName-hostname DEO\Logs\hostname DEO_All.log`, and verify all processed and scheduled workflows.
- 13 On the node where you installed DEM Worker, verify that the DEM Worker services are running.
- 14 Navigate to `C:\Program Files (x86)\VMware\VCAC\Distributed Execution Manager\DEMNAME\Logs\DEMNAME_All.log`, and verify all executed workflows.
- 15 On the node where you installed the Proxy Agent, verify that the Proxy Agent services are running.
- 16 Navigate to `C:\Program Files (x86)\VMware\VCAC\Agents\AgentName\logs\vsphereAgent.log`, and verify that the agent can successfully connect to the manager service and that it is attempting to connect to the endpoint.
- 17 Verify that you can log in to all tenants with external directory credentials if available or `vsphere.local` credentials.
- 18 In the vRealize Automation Console, verify that the **Infrastructure** tab loads successfully.
- 19 In the vRealize Automation Console, click the **Monitoring** tab and verify that the DEM services have a green status.
- 20 If you are validating vRealize Automation version 7.2, click the **Containers** tab in the vRealize Automation Console, and verify that the tab load successfully and displays all connected servers.
- 21 . Click **Administration > Health**, and configure and run system and tenant tests using vRealize Automation health service.
- IaaS administrators can configure the health service to run tests that assess the health of a selected vRealize Automation virtual machine. These tests tell you if the components are registered and all the necessary resources are available. See [Monitoring vRealize Automation Health](#).
- 22 Click **Infrastructure > Managed Machines**, and verify that all the managed virtual machines are present.
- 23 Click **Infrastructure > Compute Resources**, select each endpoint, and click **Data Collection**, **Request now**, and **Refresh** to verify that the endpoints are working.
- 24 Click **Administration > Catalog Management**, and verify the contents listed under **Services**, **Catalog Items**, **Actions**, and **Entitlements**.


- 25 Click **Design > Blueprints**, and verify the elements of each blueprint.
- 26 Click **Design > XaaS**, and verify the contents listed under **Custom Resources, Resource Mappings, XaaS Blueprints, and Resource Actions**.
- 27 Verify the details for the provisioned virtual machines.
 - a Click **Items > Deployment**.
 - b Select a provisioned, powered-off virtual machine, and select **Actions > Power On**.
 - c Click **Submit**, and click **OK**.
 - d Verify that the virtual machine powers on correctly.
 - e Repeat these steps for all provisioned virtual machines.
- 28 Request a new catalog item.
 - a Click **Catalog** and request a new catalog item.
 - b On the **General** tab, enter the request information.
 - c Click the machine icon and accept the default settings.
 - d Click **Submit**, and click **OK**.
 - e Verify that the request completes successfully.
- 29 Recreate or update the contents created or updated after backup.

Virtual machines provisioned after backup still exist on vCenter, but they lose their management. You cannot use bulk import to bring these virtual machines under management. They must be deleted manually in vCenter.

vRealize Orchestrator Restore Process

A system administrator uses backups to restore vRealize Orchestrator to a functional state after a system failure.

After you restore the images from the backup, you must power the nodes on in any order.



The vRealize Orchestrator server status appears as `Service is starting`. The first boot can take 5 to 10 minutes because the server is installing the vRealize Orchestrator plug-in's content in the database tables.

A message states that the service has started successfully.

Validate vRealize Orchestrator

After the recovery is complete, you must validate that the environment and application functionality is restored.

Procedure

- 1 Log in to the vRealize Orchestrator node 1 appliance configuration Web interface at `https://vRO_node1:8283`.
- 2 Restart the vRealize Orchestrator service from the Startup options page and wait for the service to start.
- 3 Log in to the vRealize Orchestrator node 2 appliance configuration Web interface at `https://vRO_node2:8283`.
- 4 Restart vRealize Orchestrator service from the Startup options page and wait for the service to start.
- 5 Ensure that your MSSQL database host server is running.
- 6 Ensure that the two vRealize Orchestrator nodes are working in a cluster mode.
- 7 Verify that every tab in the configuration is green.
- 8 Log in to the vRealize Orchestrator server and verify that all the running tasks, policies, and workflows appear and continue running on the home page.

vRealize Operations Manager Restore Process

A system administrator uses backups to restore vRealize Operations Manager to a functional state after a system failure.

You must start the vRealize Operations Manager components in a specified order after the restore is complete.

Common Restore Scenarios

The common restore scenarios for vRealize Operations Manager systems include a full restore of a single-node virtual appliance system and restore of multiple-node virtual appliance clusters.

Satisfy the following prerequisites for the vRealize Operations Manager cluster.

- Power off the virtual machines in the multi-node cluster that you want to restore.
- Before restoring to a different host, power off virtual machines at the original location, and then bring up the environment on the new host to avoid hostname or IP conflict. Verify that the datastore on the new host has sufficient capacity for the new cluster.
- Verify that all VMDK files have been assigned to the same datastore.
- Back up the vRealize Operations Manager multi-node cluster.

Single-Node Virtual Appliance

This scenario restores a single node system on the same host.

- 1 After the restore is complete, power on the VM.
- 2 Verify that you have set up a static IP address for the node and that the IP address is restored.
- 3 Log in to the node to verify that all your data is preserved and that all vmware-vcops services are running.
- 4 If your root password was reset, you must change it now.

For incremental backups, a user must create a cumulative differential policy. Every backup after a full backup is an incremental backup.

Note If you restore to another host, you must power off the environment at the original location and start the environment on the new host.

Multiple-Node Virtual Appliance Clusters

This scenario restores multiple-node virtual appliance clusters.

- 1 After the restore is complete, you must power on the nodes in the following order:
 - a master
 - b replica
 - c data
 - d remote collectorEnsure that each node is online before attempting to start the next component.
- 2 Verify that you have set up a static IP address for the node and that the IP address is restored.
- 3 Log in to the administrator interface of the master node and verify that all the nodes are online and running. Log in to each node to verify that all your data is preserved and all of the vmware-vcops services are up and running by using the **vmware-vcops status** command.
- 4 Log in to the administrator interface of the master node.
 - a Verify that high availability is enabled. If the **Enable High Availability** button appears, high availability was disabled during the backup and restore process. Enable high availability.
 - b Verify that all nodes are collecting metrics.
- 5 If your root password was reset, you must change it now.
- 6 (Optional) If a node in a non-high availability cluster does restart, you must take it offline with the administrator interface.
- 7 (Optional) If the replica node in a high availability cluster does not restart, you must assign a new master-replica node.

- 8 (Optional) You can execute an incremental backup only if a full backup already exists. After an incremental backup, you can choose to restore to either the full or incremental backup.

Note If you restore to another host, you must power off the environment at the original location before starting the environment on the new host.

Verify the Restore of vRealize Operations Manager Systems

After the restore operation of the VM is finished, verify that vRealize Operations Manager is in a functional state.

Note Do not power on any vRealize Operations Manager nodes during the restore operation. Wait until the entire cluster restore has finished before you power on any node.

- 1 For a non-HA cluster, power on the master node followed by the data nodes. For an HA environment, power on the master node followed by the replica node, data nodes, and remote collectors.
- 2 Use SSH to log in to the vRealize Operations Manager master node to verify the vRealize Operations Manager service status.
 - a Use SSH to switch to the vRealize Operations Manager master node and enter the `service vmware-vcops status` command.

```
# service vmware-vcops status
Slice Online=true
admin Role Enabled=true
    vRealize Operations vPostgres Replication Database is running (31810).
    vRealize Operations Gemfire Locator is running (31893).
data Role Enabled=true
    vRealize Operations vPostgres Database is running (32013).
    vRealize Operations Cassandra Distributed Database is running (21062).
    vRealize Operations Analytics is running (32142).
    vRealize Operations Collector is running (32225).
    vRealize Operations API is running (32331).
ui Role Enabled=true
remote collector Role Enabled=false
```

- b Confirm that the admin, data, and UI roles are running.
- 3 Log in to the administration UI of the master node and verify that all nodes in the cluster are up and collecting data.
 - a Open a browser and go to the administration UI of the master node: `https://<Master_Node_IP>/admin/login.action`.
 - b Log in as administrator.
 - c Verify that each node is in the Online status.
 - d Click each node and verify that adapter instances are in the Data receiving status.

- 4 If the High Availability state indicates **Enabled, degraded** after the restore, one or more nodes are inaccessible, and you must power on the nodes or you must restart the cluster.
 - a Open a browser and go to the administration UI of the master node:
https://<Master_Node_IP>/admin/login.action.
 - b Verify that all nodes are in Running state and Online .
 - 1 If a node is not in Running state and Online, power on the node and start it online.
 - 2 If all nodes are in Running state and Online, but HA is still **Enabled, degraded**, restart the cluster.

Checking the Restore of vRealize Operations Manager Systems

After you have restored a vRealize Operations Manager system, verify that the system nodes are up and running.

Procedure

- 1 Power on the master node for a simple cluster, and the master node and replica node for HA clusters.
- 2 Use SSH to log into the vRealize Operations Manager master node to check the vRealize Operations Manager service status, and run `service vmware-vcops status`.

```
# service vmware-vcops status
Slice Online=true
admin Role Enabled=true
    vRealize Operations vPostgres Replication Database is running (31810).
    vRealize Operations Gemfire Locator is running (31893).
data Role Enabled=true
    vRealize Operations vPostgres Database is running (32013).
    vRealize Operations Cassandra Distributed Database is running (21062).
    vRealize Operations Analytics is running (32142).
    vRealize Operations Collector is running (32225).
    vRealize Operations API is running (32331).
ui Role Enabled=true
remote collector Role Enabled=false
```

- 3 Confirm that the admin, data, and ui roles are running.
- 4 Verify that all the nodes in the cluster are up and collecting data. If you have an HA-enabled cluster, verify that HA mode is enabled.
 - a In a Web browser, navigate to the vRealize Operations Manager administration interface at https://<Master_Node_IP>/admin/login.action.
 - b Log in with the admin username and password.
 - c Verify that each node is online.
 - d Click each node, and verify that the status of adapter instances is Data receiving.
 - e Verify that HA mode is enabled. If the cluster is running in degraded mode, restart the cluster.

Change the IP Address of Nodes After Restoring a Cluster on a Remote Host

After you have restored a vRealize Operations Manager cluster to a remote host, change the IP address of the master nodes and data nodes to point to the new host.

Prerequisites

- Verify that the restore job has completed successfully.
- Verify that the datastore on the new host has sufficient capacity for the new cluster.

Procedure

- 1 Shut down the vRealize Operations Manager cluster at the original location.
- 2 In the Virtual Appliance Management Interface (VAMI), access the machine from the vCenter console and run the `/opt/vmware/share/vami/vami_set_network eth0 STATICV4 new IP netmask gateway` to change the IP address for each node in the cluster.

For example:

```
/opt/vmware/share/vami/vami_set_network
eth0 STATICV4 10.145.152.170 255.255.252.0 10.145.155.253
```


- 3 After the command runs successfully, restart the network, reboot each node, and power on the remote collector node.
- 4 Use SSH to access the master, data, and remote collector nodes, and run the `$VMWARE_PYTHON_BIN /usr/lib/vmware-vcopsuite/utilities/sliceConfiguration/bin/vcopsConfigureRoles.py --action=bringSliceOffline --offlineReason=restore cluster` command to take the cluster offline.
- 5 Update the CaSA database with the new IP address first on the master nodes, and then on the data nodes.
 - a Run the `vmware-casa stop` command to stop the CaSA service.
 - b Open the `/storage/db/casa/webapp/hsqldb/casa.db.script` file for editing, and replace all instances of the old IP address and with the new IP address.
 - c Run the `vmware-casa start` command to start the CaSA service.
- 6 In the following configuration files, use a text editor to replace all instances of the old IP address with the new IP address.
 - `/usr/lib/vmware-vcopsuite/utilities/sliceConfiguration/data/roleState.properties.`
 - `/usr/lib/vmware-vcops/user/conf/gemfire.properties.`

- `/usr/lib/vmware-vcops/user/conf/gemfire.locator.properties`. This configuration file only runs on the master node. Edit the `locator` parameter.
 - `/usr/lib/vmware-vcops/user/conf/gemfire.native.properties`.
 - `/usr/lib/vmware-vcops/user/conf/persistence/persistence.properties`.
- 7 Navigate to the `/usr/lib/vmware-vcops/user/conf/cassandra/` directory, and edit the `cassandra.yaml` file so that the `seeds` parameter points to the new IP address of the master node, and the `listen_address` and `broadcast_rpc_address` point to the IP addresses of the data nodes.
 - 8 Log in to the vRealize Operations Manager administration interface, and bring the cluster online.

vRealize Log Insight Restore Process

A system administrator uses backups to restore vRealize Log Insight to a functional state after a system failure.

You must start the master node for the vRealize Log Insight, followed by worker nodes in any order after the restore is complete.



Restore Guidelines

Use the following information for restoring and validating after a restore operation.

- Verify that restored nodes are in a powered-off state. Restore the nodes in a specific order and apply manual configuration changes where applicable.
- Verify that the vRealize Log Insight master node is restored first before restoring worker nodes. Worker nodes can be restored in any order.
- You can restore the VMs to the same host, to a different host on the same data center, or to a different host on a target remote data center, depending on the backup tool used.
 - Unless the vRealize Log Insight cluster is completely shut down and inaccessible, the cluster instances are powered off before you restore the cluster to a new site.
 - Verify that split-brain behavior does not occur when the same IP addresses and FQDNs are used on the recovery site. Verify that no one is unintentionally using a partially working cluster on the primary site.
- During an outage, recover the vRealize Log Insight cluster as soon as possible.
- When a successful restoration is finished, perform a quick spot check of the cluster that was restored.

Post-Recovery Configuration Change Guidelines

Depending on the recovery target and IP customizations applied during the backup configuration, manual configuration changes are required to one or more vRealize Log Insight nodes before the restored site can become fully functional.

Recovering to the Same Host

You can restore vRealize Log Insight cluster to the same host by using any back up tool.

- All network, IP, and FQDN settings that are used for the production environment should be preserved in the restored site.
- The original copy of the cluster is overwritten with the restored version unless a new name is provided to the virtual machine, during the restore process.
- If the same IP addresses and FQDNs are used for the restored cluster nodes as per the default settings, power down the existing cluster before beginning the restore.
- After a successful restoration and validation, delete the old copy to conserve resources and to prevent potential issues.

Recovering to a Different Host

You must perform manual configuration on vRealize Log Insight, if you are restoring to a different host cluster. For information about changes that are specific to vRealize Log Insight 3.3.0 versions, see [Restore to a Different Host](#). It is assumed that the restored vRealize Log Insight nodes have been assigned different IP addresses and FQDNs than their source counterparts from which a backup was taken.

Recovering vRealize Log Insight Forwarders

The manual instructions for recovering vRealize Log Insight forwarders are the same as that of the vRealize Log Insight server as described above.

Recovering vRealize Log Insight Agents

If the complete agent OS is backed up, follow the tool-specific workflow to recover the agent OS.

- If agent configuration is made on the client side, that is on agent OS, replace the agent.ini using the backup copy.
- If configuration changes are made on the server side, that is vRealize Log Insight master node, no backup and recovery is required for the agent virtual machines.

Confirming the Restoration

You must confirm that all restored vRealize Log Insight clusters are fully functional.

- Verify that you can access the vRealize Log Insight user interface using the Internal Load Balancer (ILB) IP address or FQDN (if configured) as well as access all individual cluster nodes using respective IP addresses or FQDNs.

- From the vRealize Log Insight Administration page:
 - Verify the status of cluster nodes from the cluster page and make sure the ILB, if configured, is also in an active state.
 - Verify the vSphere integration. If required, reconfigure the integration. This occurs when the ILB and/or the master node IP address or FQDN is changed post-recovery.
 - Verify the vRealize Operations Manager integration and reconfigure again if needed.
 - Verify that all content packs and UI features are functioning correctly.
 - Verify that vRealize Log Insight forwarders and agents, if configured, are functioning correctly.
- Verify that other key features of vRealize Log Insight are functioning as expected.

Restore to a Different Host

When you restore your system to a different host, you should make some configuration changes on the vRealize Log Insight cluster.

The configuration changes listed are specific to vRealize Log Insight 3.3.0. It is assumed that the restored vRealize Log Insight nodes are assigned different IP addresses and FQDNs than their source counterparts from which the backup was taken.

Procedure

- 1 List all new IP addresses and FQDNs that were assigned to each vRealize Log Insight node.
- 2 Perform the following configuration changes on the master node:
 - a Power on the master node, if it is not ON.

Note Steps b through e are applicable for vRealize Log Insight 2.5. You can not make changes to the configuration files directly from the appliance console for vRealize Log Insight 3.0 and higher. To make changes to the internal configuration options by using the web UI interface for vRealize Log Insight 3.0 and higher, refer to the Knowledge Base article [KB 2123058](#).

- b Use SSH to connect as a root user to the node's virtual appliance.
- c If the vRealize Log Insight service is running, stop the service first by running this command `service loginsight stop`.
- d Run `cd /storage/core/loginsight/config`.
- e Run `cp loginsight-config.xml#<n> backup-loginsight-config.xml` where <n> represents the largest number that is automatically suffixed to `loginsight-config.xml` during configuration changes.

- f Open the copied version of the configuration file in your favorite editor or in the vRealize Log Insight 3.0 web UI and look for lines that resemble the following lines. This configuration change is applicable to both vRealize Log Insight 2.5 and 3.0.

```
<distributed overwrite-children="true">
  <daemon host="prod-es-vrli1.domain.com" port="16520" token="c4c4c6a7-f85c-4f28-
a48f-43aeea27cd0e">
    <service-group name="standalone" />
  </daemon>
  <daemon host="192.168.1.73" port="16520" token="a5c65b52-aff5-43ea-8a6d-38807ebc6167">
    <service-group name="workernode" />
  </daemon>
  <daemon host="192.168.1.74" port="16520" token="a2b57cb5-a6ac-48ee-8e10-17134e1e462e">
    <service-group name="workernode" />
  </daemon>
</distributed>
```

In this code snippet, there are three nodes. The first one is the master node which shows `<service-group name=standalone>` and the remaining two nodes are worker nodes and show `<service-group name="workernode">`.

- g For the master node, in the newly recovered environment, verify if the DNS entry that was used in the pre-recovery environment can be reused.
- If the DNS entry can be reused, you only need to update the DNS entry to point to the new IP address of the master node.
 - If the DNS entry cannot be reused, replace the master node entry with the new DNS name, pointing to the new IP address.
 - If the DNS name cannot be assigned, as a last option, update the configuration entry with the new IP address.
- h Update the worker node IP addresses to reflect the new IP addresses.

- i In the same configuration file, look for entries that represent NTP, SMTP and database, and appenders sections.

This applies to vRealize Log Insight 2.5 and 3.0.

Note The `<logging><appenders>...</appenders></logging>` section is applicable only to the vRealize Log Insight 2.5 and is not available for vRealize Log Insight 3.0.

```
<ntp>
  <ntp-servers value="ntp1.domain.com, ntp2.domain.com" />
</ntp>

<smtp>
  <server value="smtp.domain.com" />
  <default-sender value="source.domain.com@domain.com" />
</smtp>

<database>
  <password value="xserttt" />
  <host value="vrli-node1.domain.com" />
  <port value="12543" />
</database>

<logging>
  <appenders>
    <appender name="REMOTE" class="com.vmware.loginsight.commons.logging.ThriftSocketAppender">
      <param name="RemoteHost" value="vdl-node1.domain.com" />
    </appender>
  </appenders>
</logging>
```

- If the configured NTP server values are not valid in the new environment, update these in the `<ntp>...</ntp>` section.
- If the configured SMTP server values are not valid in the new environment, update these in the `<smtp>...</smtp>` section.
- Optionally, change the `default-sender` value in the SMTP section. The value can be any value, but as a good practice, you should represent the source from where the email was sent.
- In the `<database>...</database>` section, change the host value to point to the master node FQDN or IP address.
- In the `<logging><appenders>...</appenders></logging>` section, change the parameter value for `RemoteHost` to reflect the new master node FQDN or IP address.

- j In the same configuration file, update the vRealize Log Insight ILB configuration section

For a vRealize Log Insight appliance,

```
<load-balancer>
<leadership-lease-renewal-secs value="5" />
<high-availability-enabled value="true" />
<high-availability-ip value="10.158.128.165" />
<high-availability-fqdn value="LB-FQDN.eng.vmware.com" />
<layer4-enabled value="true" />
<ui-balancing-enabled value="true" />
</load-balancer>
```

For a vRealize Log Insight 2.5 appliance,

```
<load-balancer>
  <leadership-lease-renewal-secs value="5" />
  <high-availability-enabled value="true" />
  <high-availability-ip value="192.168.1.75" />
  <layer4-enabled value="true" />
</load-balancer>
```

- k Under the `<load-balancer>...</load-balancer>` section, update the `high-availability-ip` value if it is different from the current setting.
- l In the vRealize Log Insight, make sure to also update the FQDN of the load balancer.
- m Rename the updated configuration file to finish the changes.

Note This step is applicable for vRealize Log Insight 2.5 only. In vRealize Log Insight 3.0 the changes are made through web UI.

Run : `mv backup-loginsight-config.xml loginsight-config.xml#<n+1>` where n represents the current maximum number suffixed to the `loginsight-config.xml` files.

- n For vRealize Log Insight 2.5, restart the vRealize Log Insight service and run : `service loginsight start`.

Note For vRealize Log Insight 3.0, this can be achieved from the web UI by going to the Cluster tab on the Administration page. For each node listed, select its hostname or IP address to open the details panel and click **Restart Log Insight**. The configuration changes are automatically applied to all cluster nodes.

- o Wait for two minutes after the vRealize Log Insight service starts in order to give enough time for Cassandra services to come up before bringing other worker nodes online.

Note You can skip steps 3 to 9 for vRealize Log Insight 3.0. These steps are only applicable for vRealize Log Insight 2.5.

- 3 SSH onto the first worker node using root credentials.

- 4 Stop the vRealize Log Insight service and run : `service loginsight stop.`
- 5 Copy the latest `loginsight-config.xml` file from the master node to the worker node.
- 6 On the worker node, run : `scp root@[master-node-ip]:/storage/core/loginsight/config/loginsight-config.xml#<n> /storage/core/loginsight/config/.`
- 7 Run : `service loginsight start.`
- 8 Wait for 2 minutes after the vRealize Log Insight service starts in order to give enough time for Cassandra service to start completely.
- 9 Repeat the steps for each worker node.

vRealize Business Restore Process

You can restore vRealize Business to the last full or incremental backup.

If a failure occurs, a system administrator must restore vRealize Business Standard to a functional state.

Verify that the vRealize Automation system is running before you start the vRealize Business VM.

vRealize Business Startup Order

vRealize Business Appliance

(Only one Node)

If you change the certificate of vRealize Automation, you must reregister vRealize Business with vRealize Automation.

Validate vRealize Business

- 1 Log in to `https://<VRB IP>:5480/` by using root credentials and ensure that no configuration has changed after the restore.
- 2 Verify that vRealize Automation is in the Registered status in vRealize Business.
- 3 Log in to vRealize Automation and verify that the cost profiles that are set to Automatic continue to pull the rates from vRealize Business.
- 4 Log in to vRealize Automation and verify that the **Business Management** tab appears.
- 5 Log in to vRealize Automation and verify that vRealize Business is collecting data after the restore and that the cost of the VMs is correctly calculated. This cost should be the same as before the backup.
- 6 Provision any VM by using vRealize Orchestrator or vRealize Automation and validate that vRealize Business is able to calculate the VM costs .
- 7 If you are adding more endpoints to vRealize Automation, vRealize Business should be able to calculate these endpoints' managed VM cost.

Restore vRealize Suite by Using EMC Avamar

5

You can restore the backed-up data for vRealize Suite components by using EMC Avamar to restore the full image to the original virtual machine.

If a failure occurs, you can perform an image restore. Avamar offers the following choices for restore destinations:

- Original virtual machine
- Different (existing) virtual machine
- New virtual machine

For information on restoring to a destination other than the original virtual machine, see the [EMC Avamar for VMware User Guide](#).

Prerequisites

Verify that a backup of the vRealize Suite components is available.

Procedure

- 1 In the vSphere Client or vSphere Web Client, ensure that the target virtual machine is powered off.
- 2 In Avamar Administrator, click **Backup & Restore**.
- 3 Click the **Restore** tab.

The upper left pane contains a list of domains.

- 4 Select a virtual machine client or VMware container.
 - a Select the domain that contains the virtual machine client or VMware container.

You cannot view clients outside the domain for the login account. To view all clients, log in to the root domain. A list of Avamar clients appears in the pane under the domains list.
 - b From the list of clients, select the virtual machine client or VMware container.

5 Select a backup.

- a Click the **By Date** tab.
- b Select the backup date from the calendar.

Valid backups occurred on dates with a yellow highlight. A list of backups on that date appears in the Backups table next to the calendar.

- c Select a backup from the **Backups** table.

6 Click **Browse for Image Restore** directly above the contents pane.

7 Select **All virtual disks** to restore the entire image.

8 From the **Actions** menu, select **Restore Now**.

The **Restore Option** dialog box appears.

9 Select the restore destination.

If you are not restoring to the original virtual machine, configure the destination. When restoring an image backup to the original virtual machine, the **Configure Destination** button is disabled.

10 (Optional) Select **Restore virtual machine configuration** to restore VMware configuration files.

11 (Optional) Select a proxy to perform restore.

The default setting is automatic, which enables the Avamar server to select the best proxy for this operation.

12 Click **More Options**.

The **Restore Command Line Options** dialog box appears.

13 (Optional) Select **Use Changed Block Tracking (CBT)** to increase performance.

14 From the **Encryption method from Data Domain system** list, select the encryption method to use for data transfer between the Data Domain system and the client during the restore.

15 Select **Do not power on VM after restore** from the **Select Post Restore Options** list.

16 (Optional) Click **More**.

To include additional plug-in options with this restore, type **Enter Attribute** and **Enter Attribute Value** settings, and click **+**.

17 Click **OK** on the **Restore Command Line Options** dialog box.

18 Click **OK** on the **Restore Options** dialog box. The following status message appears: Restore initiated.

19 Click **OK**.

20 If the restore target virtual machine will be using changed block tracking for future backups, enable changed block tracking by performing any of the following actions on that virtual machine: reboot, power on, resume after suspend, or migrate.

What to do next

- If you restored using the Restore to virtual machine option, verify that the original MAC addresses are configured on the NICs. If you cannot recover the original addresses, open a command prompt on the restored appliances. Delete the `/etc/udev/rules.d/70-persistent-net.rules` file, and then restart the virtual machines that run each appliance.

You delete the file and restart each appliance because the operating systems used by vRealize Operations Manager and vRealize Automation preserve the MAC address in this file. You must perform this action on each of your restored virtual machines or the applications will not start. When you delete the file and restart the virtual machines, the operating system regenerates the file and orders the interfaces based on their place on the PCI bus. The first NIC is `eth0`, the second `eth1`, and so on. This is the order expected by vRealize Operations Manager and vRealize Automation.

- After the restore is finished, start the cluster and verify that the system is working correctly. After you restore your vRealize Suite components, refer to the validation process for each component in [Chapter 4 Restoring, Powering On, and Validating vRealize Suite](#) to verify that your system is working correctly.
- If any cluster displays the status **waiting for analytics** in vRealize Operations Manager after the restore is finished, restart the analytics service. See [Chapter 6 Restart the Cluster Analytic Service](#).
- For information on monitoring and cancelling backup procedures, see the [EMC Avamar for VMware User Guide](#).

Restart the Cluster Analytic Service

6

Restart the analytic service on a cluster if the service does not automatically restart after you restore vRealize Suite.

Problem

If files are corrupted during the restore process, the cluster analytic service fails to restart after you restore vRealize Suite. When this occurs, vRealize Operations Manager displays a **waiting for analytics** status for the cluster.

Solution

- 1 Search all restored files for the text string @^.

Files that contain this text string are corrupt.

- 2 Replace all corrupt files with files copied from a node that still has uncorrupted files.
- 3 Take the cluster offline and back online.

If the analytic service successfully restarts when you bring the cluster online, you can skip the following steps.

- 4 Run the following command on all nodes.

```
$VMWARE_PYTHON_BIN $ALIVE_BASE/./vmware-vcopssuite/utilities/sliceConfiguration/bin/vcopsConfigureRoles.py --action bringSliceOffline --offlineReason "test"
```

- 5 Stop the casa service by running the following command.

```
service vmware-casa stop
```

- 6 In the `/data/db/casa/webapp/hsqldb/casa.db.script` script file, set the role of all nodes to offline.

- 7 Start the casa service by running the following command.

```
service vmware-casa start
```

- 8 Take the cluster offline and back online.
- 9 Verify that all nodes successfully come back online.

Manual Backup Procedure Appears to Stall



When you run a backup job manually by using vSphere Data Protection, the progress of the job might reach 92% and stall. It appears as though the job has stopped running.

Problem

The task details in the **Running** tab of the Recent Tasks pane might show that the job has stopped running when it reaches 92%. Often, the job might still be running in the background. The status of the backup job can be verified in the vSphere Data Protection Appliance.

Solution

- 1 Use SSH to log in to vSphere Data Protection appliance.
- 2 Run `mccli activity show` to view a list of backup jobs and their status.
- 3 In the Client column, search for the ID of the backup job, and the corresponding virtual machines.
- 4 In the Status column, verify that the job is still running.