Migrating vRealize Automation to 7.3

vRealize Automation 7.3

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-002425-01

vmware®

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

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

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

Contents

Updated Information

This *Migrating vRealize Automation to 7.3* is updated with each release of the product or when necessary.

This table provides the update history of the *Migrating vRealize Automation to 7.3* documentation.

Revision	Description
002425-01	■ Made minor editorial updates. ■ Changed title and added information to "Run Test Connection and Verify Migrated Endpoints," on page 30. ■ Added vRealize Orchestrator migration topics. ■ "Migrating an External vRealize Orchestrator Server to vRealize Automation 7.3," on page 32 ■ "Migrate an External vRealize Orchestrator 6.x on Windows to vRealize Automation 7.3," on page 33 ■ "Migrate an External vRealize Orchestrator 6.x Virtual Appliance to vRealize Automation 7.3," on page 35 ■ "Migrate an External vRealize Orchestrator 7.x to vRealize Automation 7.3," on page 37 ■ "Configure the Built-In vRealize Orchestrator Server," on page 38 ■ "Control Center Differences Between External and Embedded Orchestrator," on page 32
002425-00	Initial release.

1 Migrating vRealize Automation

You can perform a side-by-side upgrade of your current vRealize Automation environment using migration.

Migration moves all data, except for tenants and identity stores, from your current vRealize Automation source environment to a target deployment of the latest version of vRealize Automation.

Migration does not change your source environment except to stop vRealize Automation services for the time required to collect and copy the data safely to your target environment. Depending on the size of the source vRealize Automation database, migration can take from a few minutes to hours.

You can migrate your source environment to a minimal deployment or a high-availability deployment.

If you plan to put your target environment into production, do not put your source environment back into service after migration. Changes to your source environment after migration are not synchronized with your target environment.

If your source environment is integrated with vCloud Air or vCloud Director or has physical endpoints, you must use migration to perform an upgrade. Migration removes these endpoints and everything associated with them from the target environment. Migration also removes a 6.x
VMware vRealize Application Services integration from the target environment.

Migration Prerequisites 2

The migration prerequisites differ depending on your target environment.

You can migrate to a minimal environment or to a high-availability environment.

This chapter includes the following topics:

Prerequisites for Migration to a Minimal Environment

Ensure a successful migration to a minimal environment by reviewing these prerequisites.

Prerequisites

- Verify that you have a new target environment of vRealize Automation.
- Install relevant proxy agents on the target environment according to these requirements.
  - Target proxy agent name must match the source proxy agent name for vSphere, Hyper-V, Citrix XenServer, and Test proxy agents.

    Note Finish these steps to obtain an agent name.

    1 Go to the agent installation directory on the IaaS node.

    2 Open the `VRMAgent.exe.config` file.

    3 Under the serviceConfiguration tag, look for the value of the agentName attribute.
  - Target proxy agent endpoint name must match the source proxy agent endpoint name for vSphere, Hyper-V, Citrix XenServer, and Test proxy agents.
  - Do not create an endpoint for vSphere, Hyper-V, Citrix XenServer, or Test proxy agents on the target environment.
- Review the version numbers of vRealize Automation components.

  a In your target vRealize Automation 7.3 environment, start a browser. Go to the vRealize Automation appliance management console at https://*vra-va-hostname.domain.name*:5480.

  b Log in with the user name **root** and the password you entered when you deployed the appliance.

  c Select **vRA Settings > Cluster**.

  d Expand the Host / Node Name records by clicking the triangle.

  Verify that the version numbers of the vRealize Automation IaaS components match.

- Verify that the target Microsoft SQL Server version for the vRealize Automation target IaaS database is 2012, 2014, or 2016.
- Verify that port 22 is open between the source and target vRealize Automation environments. Port 22 is required to establish Secure Shell (SSH) connections between source and target virtual appliances.
- Verify that the IaaS server node in the target environment has at least Java SE Runtime Environment (JRE) 8, update 111 (64 bit) installed. After you install the JRE, make sure the JAVA_HOME system variable points to the Java version you installed on each IaaS node. Revise the path if necessary.
- Verify that each IaaS node has PowerShell 3.0 or later installed.
- Verify that the source and target vRealize Automation environments are running.
- Verify that no user and provisioning activities are happening on the source vRealize Automation environment.
- Security software must not interact with the operating system and its components running on IaaS nodes in the target vRealize Automation environment during migration. If you have any antivirus or security software installed, verify that the software is correctly configured or disabled for migration.

What to do next

Chapter 3, "Pre-Migration Tasks," on page 13.

Prerequisites for Migration to a High-Availability Environment

Ensure a successful migration to a high-availability environment by reviewing these prerequisites.

Prerequisites

- Verify that you have a new target installation of vRealize Automation with a master and replica virtual appliance configured for high availability. See *vRealize Automation High Availability Configuration Considerations* in *Reference Architecture*.
- Verify that all vRealize Automation virtual appliances use the same password for root user.
- Install relevant proxy agents on the target environment according to these requirements.
  - Target proxy agent name must match the source proxy agent name for vSphere, Hyper-V, Citrix XenServer, and Test proxy agents.

    Note Finish these steps to obtain an agent name.

    1 Go to the agent installation directory on the IaaS node.

    2 Open the `VRMAgent.exe.config` file.

    3 Under the serviceConfiguration tag, look for the value of the agentName attribute.

  - Target proxy agent endpoint name must match the source proxy agent endpoint name for vSphere, Hyper-V, Citrix XenServer, and Test proxy agents.
  - Do not create an endpoint for vSphere, Hyper-V, Citrix XenServer, or Test proxy agents on the target environment.
- Check the version numbers of vRealize Automation components.

  a In your target vRealize Automation 7.3 environment, start a browser and go to the vRealize Automation appliance management console at https://*vra-va-hostname.domain.name*:5480.

  b Log in with the user name **`root`** and the password you entered when you deployed the appliance.

  c Select **vRA Settings > Cluster**.

  d To expand the Host / Node Name records so you can see the components, click the expand button.

Verify that the version numbers of vRealize Automation components match across all virtual appliance nodes.

Verify that the version numbers of vRealize Automation IaaS components match across all IaaS nodes.

- Perform these steps to direct traffic to only the master node.
  a. Disable all the redundant nodes.
  b. Remove the health monitors for these items according to your load balancer documentation:
     - vRealize Automation virtual appliance
     - IaaS Website
     - IaaS Manager Service
- Verify that the vRealize Automation appliance master node connects to the PostgreSQL database in MASTER mode.
  a. In your target vRealize Automation 7.3 environment, start a browser and go to the master vRealize Automation appliance management console at https://*vra-va-hostname.domain.name*:5480.
  b. Log in with the user name **root** and the password you entered when you deployed the appliance.
  c. Select **vRA Settings > Database**.
  d. Verify that the database node host mode is MASTER.
- Verify that the target Microsoft SQL Server version for the vRealize Automation target IaaS database is 2012, 2014, or 2016.
- Verify that port 22 is open between the source and target vRealize Automation environments. Port 22 is required to establish Secure Shell (SSH) connections between source and target virtual appliances.
- Verify that the IaaS Web Service and Model Manager Server nodes in the target environment have the right Java Runtime Environment. You must have Java SE Runtime Environment (JRE) 8, update 111 (64 bit) or later installed. Make sure the JAVA_HOME system variable points to the Java version you installed on each IaaS node. Revise the path if necessary.
- Verify that each IaaS node has at least PowerShell 3.0 or later installed.
- Verify that the source and target vRealize Automation environments are running.
- Verify that no user and provisioning activities are happening on the source vRealize Automation environment.
- Verify that any antivirus or security software that might interact with the operating system and its components running on IaaS nodes in the target vRealize Automation environment is correctly configured or disabled.
- Security software must not interact with the operating system and its components running on IaaS nodes in the target vRealize Automation environment during migration. If you have any antivirus or security software installed, verify that it is correctly configured or disabled for migration.

What to do next

Chapter 3, “Pre-Migration Tasks,” on page 13.

Pre-Migration Tasks 3

Before you migrate, you must perform several pre-migration tasks.

The pre-migration tasks you perform before you migrate your source vRealize Automation environment to the target vRealize Automation 7.3 environment vary depending on your source environment.

This chapter includes the following topics:

Gather Information Required for Migration

Use these tables to record the information that you need for migration from your source and target environments.

Prerequisites

Finish verifying the prerequisites for your situation.

- "Prerequisites for Migration to a Minimal Environment," on page 9.
- "Prerequisites for Migration to a High-Availability Environment," on page 10.

Table 3-1. Source vRealize Automation Appliance

Option	Description	Value
Host name	Log in to your source vRealize Automation appliance management console. Find the host name on the **System** tab. The host name must be a fully qualified domain name (FQDN).	
Root username	root	
Root password	The root password that you entered when you deployed your source vRealize Automation appliance.	

Table 3-2. Target vRealize Automation Appliance

Option	Description	Value
Root username	root	
Root password	The root password that you entered when you deployed your target vRealize Automation appliance.	
Default tenant	The default tenant you created when you configured single sign-on in the vRealize Automation Installation wizard, usually vsphere.local.	
Administrator username	Default tenant administrator user name that you entered when you deployed the target vRealize Automation environment, usually administrator.	
Administrator password	Password for the default tenant administrator user that you entered when you deployed the target vRealize Automation environment.	

Table 3-3. Target IaaS Database

Option	Description	Value
Database server	The location of the Microsoft SQL Server where the restored vRealize Automation IaaS Microsoft SQL database resides. If a named instance and non-default port is used, enter in SERVER,PORT\INSTANCE-NAME format.	
Cloned database name	Name of the source vRealize Automation 6.2.x or 7.x IaaS Microsoft SQL database that you backed up on the source and restored on the target environment.	
Login name	Login name of a user with db_owner role for the cloned IaaS Microsoft SQL database in the target environment. For Windows Authentication, the Windows account for the vCloud Automation Center Management Agent service must be db_owner for the cloned IaaS SQL database.	
Password	Password for the SQL Server user who has the db_owner role for the cloned IaaS Microsoft SQL database.	

Table 3-3. Target IaaS Database (Continued)

Option	Description	Value
Original encryption key	Original encryption key that you retrieve from the source environment. See "Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.	
New passphrase	A series of words used to generate a new encryption key. You use this passphrase each time you install a new IaaS component in the target vRealize Automation environment.	

What to do next

"Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.

Obtain the Encryption Key from the Source vRealize Automation Environment

You must enter the encryption key from the source vRealize Automation environment as part of the migration procedure.

Prerequisites

Verify that you have administrator privileges on the active Manager Service host virtual machine in your source environment.

Procedure

1 Open a command prompt as an administrator on the virtual machine that hosts the active Manager Service in your source environment and run this command.

```
"C:\Program Files
(x86)\VMware\vCAC\Server\ConfigTool\EncryptionKeyTool\DynamicOps.Tools.EncryptionKeyTool.exe"
key-read -c "C:\Program Files (x86)\VMware\vCAC\Server\ManagerService.exe.config" -v
```

If your installation directory is not in the default location, `C:\Program Files (x86)\VMware\vCAC`, edit the path to show your actual installation directory.

2 Save the key that appears after you run the command.

The key is a long string of characters that looks similar to this example:

```
NRH+f/BlnCB6yvasLS3sxespgdkcFWAEuyV0g4lfryg=.
```

What to do next

- If you are migrating from a vRealize Automation 6.2.x environment: "Add Each Tenant from the Source vRealize Automation Environment to the Target Environment," on page 16.
- If you are migrating from a vRealize Automation 7.x environment: "List Tenant and IaaS Administrators from the Source vRealize Automation 6.2.x Environment," on page 16.

List Tenant and IaaS Administrators from the Source vRealize Automation 6.2.x Environment

Before you migrate a vRealize Automation 6.2.x environment, you must make a list of the tenant and IaaS administrators for each tenant.

Perform the following procedure for each tenant in the source vRealize Automation console.

Note If you migrate from a vRealize Automation 7.x environment, you do not need to perform this procedure.

Prerequisites

Log in to the source vRealize Automation console.

1 Open the vRealize Automation console using the fully qualified domain name of the source virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

For a high-availability environment, open the console using the fully qualified domain name of the source virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

2 Log in with the user name **`administrator@vsphere.local`** and the password that you entered when you deployed the source vRealize Automation.

Procedure

1 Select **Administration > Tenants**.

2 Click a tenant name.

3 Click **Administrators**.

4 Make a list of each tenant and IaaS administrator user name.

5 Click **Cancel**.

What to do next

"Add Each Tenant from the Source vRealize Automation Environment to the Target Environment," on page 16.

Add Each Tenant from the Source vRealize Automation Environment to the Target Environment

You must add tenants in the target environment using the name of each tenant in the source environment.

For successful migration, it is mandatory that each tenant in the source environment is created in the target environment. You must also use a tenant-specific access URL for each tenant that you add using the tenant URL name from the source environment. If there are unused tenants in the source environment that you do not want to migrate, delete them from the source environment before migration.

Perform this procedure for each tenant in your source environment.

- When you migrate from a vRealize Automation 6.2.x environment, you migrate your existing SSO2 tenants and identity stores on the source environment to the VMware Identity Manager on the target environment.
- When you migrate from a vRealize Automation 7.x environment, you migrate your existing VMware Identity Manager tenants and identity stores on the source environment to the VMware Identity Manager on the target environment.

Prerequisites

- "Gather Information Required for Migration," on page 13.
- Log in to the target vRealize Automation console.
  a. Open the vRealize Automation console using the fully qualified domain name of the target virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

     For a high-availability environment, open the console using the fully qualified domain name of the target virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

  b. Log in with the user name **`administrator@vsphere.local`** and the password that you entered when you deployed the target vRealize Automation.

Procedure

1. Select **Administration > Tenants**.
2. Click the **New** icon ( ).
3. In the **Name** text box, enter a tenant name that matches a tenant name in the source environment.

   For example, if the tenant name in the source environment is DEVTenant, enter **`DEVTenant`**.
4. (Optional) Enter a description in the **Description** text box.
5. In the **URL Name** text box, enter a tenant URL name that matches the tenant URL name in the source environment.

   The URL name is used to append a tenant-specific identifier to the vRealize Automation console URL.

   For example, if the URL name for DEVTenant in the source environment is dev, enter **`dev`** to create the URL https://*vra-va-hostname.domain.name*/vcac/org/dev.
6. (Optional) Enter an email address in the **Contact Email** text box.
7. Click **Submit and Next**.

What to do next

"Create an Administrator for Each Added Tenant," on page 17.

Create an Administrator for Each Added Tenant

You must create an administrator for each tenant that you added to the target environment. You create an administrator by creating a local user account and assigning tenant administrator privileges to the local user account.

Perform this procedure for each tenant in your target environment.

Prerequisites

- "Add Each Tenant from the Source vRealize Automation Environment to the Target Environment," on page 16.
- Log in to the target vRealize Automation console.
  a. Open the vRealize Automation console using the fully qualified domain name of the target virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

     For a high-availability environment, open the console using the fully qualified domain name of the target virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

  b. Log in with the user name **`administrator@vsphere.local`** and the password that you entered when you deployed the target vRealize Automation.

Procedure

1 Select **Administration > Tenants**.

2 Click a tenant that you added.

For example, for DEVTenant, click **DEVTenant**.

3 Click **Local users**.

4 Click the **New** icon ( ).

5 In User Details, enter the requested information to create a local user account to assign the tenant administrator role.

The local user name must be unique to the default local directory, vsphere.local.

6 Click **OK**.

7 Click **Administrators**.

8 Enter the local user name in the **Tenant administrators** search box and press Enter.

9 Click the appropriate name in the search returns to add the user to the list of tenant administrators.

10 Click **Finish**.

11 Log out of the console.

What to do next

- For a minimal deployment: "Synchronize Users and Groups for an Active Directory Link Before Migration to a Minimal Environment," on page 18.
- For a high-availability deployment: "Synchronize Users and Groups for an Active Directory Link Before Migration to a High-Availability Environment," on page 20.

Synchronize Users and Groups for an Active Directory Link Before Migration to a Minimal Environment

Before you import your users and groups to a minimal deployment of vRealize Automation, you must connect to your Active Directory link.

Perform this procedure for each tenant. If a tenant has more than one Active Directory, perform this procedure for each Active Directory that the tenant uses.

Prerequisites

- "Create an Administrator for Each Added Tenant," on page 17.
- Verify that you have access privileges to the Active Directory.
- Log in to the tenanted target vRealize Automation console at `https://vra-va-hostname.domain.name/vcac/org/tenant-URL-name` with the tenant administrator user name and password.

Procedure

1 Select **Administration > Directories Management > Directories**.

2 Click **Add Directory** icon ( ) and select **Add Active Directory over LDAP/IWA**.

3 Enter your Active Directory account settings.

- For Non-Native Active Directories

Option	Sample Input
Directory Name	Enter a unique directory name. Select **Active Directory over LDAP** when using Non-Native Active Directory.
This Directory Supports DNS Service Location	Deselect this option.
Base DN	Enter the distinguished name (DN) of the starting point for directory server searches. For example, **cn=users,dc=rainpole,dc=local**.
Bind DN	Enter the full distinguished name (DN), including common name (CN), of an Active Directory user account that has privileges to search for users. For example, **cn=config_admin infra,cn=users,dc=rainpole,dc=local**.
Bind DN Password	Enter the Active Directory password for the account that can search for users and click **Test Connection** to test the connection to the configured directory.

- For Native Active Directories

Option	Sample Input
Directory Name	Enter a unique directory name. Select **Active Directory (Integrated Windows Authentication)** when using Native Active Directory.
Domain Name	Enter the name of the domain to join.
Domain Admin Username	Enter the user name for the domain admin.
Domain Admin Password	Enter the password for the domain admin.
Bind User UPN	Use the email address format to enter the name of the user who can authenticate with the domain.
Bind DN Password	Enter the Active Directory bind account password for the account that can search for users.

4 Click **Save & Next**.

Select the Domains displays a list of domains.

5 Accept the default domain setting and click **Next**.

6 Verify that the attribute names are mapped to the correct Active Directory attributes, and click **Next**.

7 Select the groups and users to synchronize.

a Click the **New** icon ( ).

b Enter the user domain and click **Find Groups**.

For example, enter **`dc=vcac,dc=local`**.

c To select the groups to synchronize, click **Select** and click **Next**.

d On Select Users, select the users to synchronize and click **Next**.

8 Review the users and groups you are syncing to the directory, and click **Sync Directory**.

The directory synchronization takes some time and runs in the background.

What to do next

"Run NSX Network and Security Inventory Data Collection in the Source vRealize Automation Environment," on page 21

Synchronize Users and Groups for an Active Directory Link Before Migration to a High-Availability Environment

Before you import your users and groups to a high-availability vRealize Automation environment, you must connect to your Active Directory link.

- Perform steps 1- 8 for each tenant. If a tenant has more than one Active Directory, perform this procedure for each Active Directory that the tenant uses.
- Repeat steps 9–10 for each identity provider associated with a tenant.

Prerequisites

- "Create an Administrator for Each Added Tenant," on page 17.
- Verify that you have access privileges to the Active Directory.
- Log in to the tenanted target vRealize Automation console at `https://vra-va-lb-hostname.domain.name/vcac/org/tenant-URL-name` with the tenant administrator user name and password.

Procedure

1 Select **Administration > Directories Management > Directories**.

2 Click **Add Directory** icon ( ) and select **Add Active Directory over LDAP/IWA**.

3 Enter your Active Directory account settings.

- For Non-Native Active Directories

Option	Sample Input
Directory Name	Enter a unique directory name. Select **Active Directory over LDAP** when using Non-Native Active Directory.
This Directory Supports DNS Service Location	Deselect this option.
Base DN	Enter the distinguished name (DN) of the starting point for directory server searches. For example, **cn=users,dc=rainpole,dc=local**.
Bind DN	Enter the full distinguished name (DN), including common name (CN), of an Active Directory user account that has privileges to search for users. For example, **cn=config_admin infra,cn=users,dc=rainpole,dc=local**.
Bind DN Password	Enter the Active Directory password for the account that can search for users and click **Test Connection** to test the connection to the configured directory.

- For Native Active Directories

Option	Sample Input
Directory Name	Enter a unique directory name. Select **Active Directory (Integrated Windows Authentication)** when using Native Active Directory.
Domain Name	Enter the name of the domain to join.

Option	Sample Input
Domain Admin Username	Enter the user name for the domain admin.
Domain Admin Password	Enter the password for the domain admin account.
Bind User UPN	Use the email address format to enter the name of the user who can authenticate with the domain.
Bind DN Password	Enter the Active Directory bind account password for the account that can search for users.

4 Click **Save & Next**.

The Select the Domains page displays the list of domains.

5 Accept the default domain setting and click **Next**.

6 Verify that the attribute names are mapped to the correct Active Directory attributes, and click **Next**.

7 Select the groups and users to synchronize.

a Click the **New** icon  .

b Enter the user domain and click **Find Groups**.

For example, enter `dc=vcac,dc=local`.

c To select the groups to synchronize, click **Select** and click **Next**.

d On the Select Users page, select the users to synchronize and click **Next**.

8 Review the users and groups you are syncing to the directory, and click **Sync Directory**.

The directory synchronization takes some time and runs in the background.

9 Select **Administration > Directories Management > Identity Providers**, and click your new identity provider.

For example, **WorkspaceIDP__1**.

10 On the page for the identity provider that you selected, add a connector for each node.

a Follow the instructions for **Add a Connector**.

b Update the value for the **IdP Hostname** property to point to the fully qualified domain name (FQDN) for the vRealize Automation load balancer.

c Click **Save**.

What to do next

"Run NSX Network and Security Inventory Data Collection in the Source vRealize Automation Environment," on page 21.

Run NSX Network and Security Inventory Data Collection in the Source vRealize Automation Environment

Before you migrate, you must run NSX Network and Security Inventory data collection in the source vRealize Automation environment.

This data collection is necessary for the load balancer reconfigure action to work in vRealize Automation 7.3 for 7.1 and 7.2 deployments.

NOTE You do not perform this data collection in a vRealize Automation 6.2.x environment. The load balancer reconfigure action is not supported in a 6.2.x deployment.

Prerequisites

- For a minimal deployment: "Synchronize Users and Groups for an Active Directory Link Before Migration to a Minimal Environment," on page 18.
- For a high-availability deployment: "Synchronize Users and Groups for an Active Directory Link Before Migration to a High-Availability Environment," on page 20.

Procedure

- Run NSX Network and Security Inventory data collection in your source vRealize Automation environment before you migrate to vRealize Automation 7.3. See *Start Endpoint Data Collection Manually* in *Managing vRealize Automation*.

What to do next

"Manually Clone the Source vRealize Automation IaaS Microsoft SQL Database," on page 22.

Manually Clone the Source vRealize Automation IaaS Microsoft SQL Database

Before migration, you must back up your IaaS Microsoft SQL database in the vRealize Automation source environment and restore it to a new blank database created in the vRealize Automation target environment.

Prerequisites

- "Run NSX Network and Security Inventory Data Collection in the Source vRealize Automation Environment," on page 21.
- Obtain information about backing up and restoring an SQL Server database. Find articles on the Microsoft Developer Network about creating a full SQL Server database backup and restoring an SQL Server database to a new location.

Procedure

- Create a full backup of your source vRealize Automation 6.2.x or 7.x IaaS Microsoft SQL database. You use the backup to restore the SQL database to a new blank database created in the target environment.

What to do next

"Snapshot the Target vRealize Automation Environment," on page 22.

Snapshot the Target vRealize Automation Environment

Take a snapshot of each target vRealize Automation virtual machine. If migration is unsuccessful, you can try again using the virtual machine snapshots.

For information, see your vSphere documentation.

Prerequisites

"Manually Clone the Source vRealize Automation IaaS Microsoft SQL Database," on page 22.

What to do next

Perform one of the following procedures:

- "Migrate vRealize Automation Source Data to a vRealize Automation 7.3 Minimal Environment," on page 23.
- "Migrate vRealize Automation Source Data to a vRealize Automation 7.3 High-Availability Environment," on page 25.

Migration Procedures 4

The procedure you perform to migrate your source vRealize Automation environment data depends whether you migrate to a minimal environment or to a high-availability environment.

This chapter includes the following topics:

Migrate vRealize Automation Source Data to a vRealize Automation 7.3 Minimal Environment

You can migrate your current vRealize Automation environment to a new installation of vRealize Automation 7.3.

Prerequisites

- "Gather Information Required for Migration," on page 13.
- "Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.
- "Add Each Tenant from the Source vRealize Automation Environment to the Target Environment," on page 16.
- "Create an Administrator for Each Added Tenant," on page 17.
- "Synchronize Users and Groups for an Active Directory Link Before Migration to a Minimal Environment," on page 18.
- "Manually Clone the Source vRealize Automation IaaS Microsoft SQL Database," on page 22.
- "Snapshot the Target vRealize Automation Environment," on page 22.

Procedure

1 In your target vRealize Automation 7.3 environment, start a browser and go to the vRealize Automation appliance management console at https://*vra-va-hostname.domain.name*:5480.

2 Log in with the user name **`root`** and the password you entered when you deployed the appliance.

3 Select **vRA Settings > Migration**.

4 Enter the information for the source vRealize Automation appliance.

Option	Description
Host name	The host name for the source vRealize Automation appliance.
Root username	`root`
Root password	The root password that you entered when you deployed the vRealize Automation appliance.

5 Enter the information for the target vRealize Automation appliance.

Option	Description
Root username	`root`
Root password	The root password that you entered when you deployed the target vRealize Automation appliance.
Default tenant	The default tenant you created when you configured single sign-on in the Installation wizard, usually vsphere.local.
Administrator username	The tenant administrator user name that you entered when you deployed the target vRealize Automation appliance. Change existing value if necessary.
Administrator password	The password that you entered for the default tenant administrator when you deployed the target vRealize Automation appliance.

6 Enter the information for the target IaaS database server.

Option	Description
Database server	The location of the Microsoft SQL Server where the restored vRealize Automation IaaS Microsoft SQL database resides. If a named instance and a non-default port are used, enter in *SERVER,PORT\INSTANCE-NAME* format.
Cloned database name	Name of the source vRealize Automation 6.2.x or 7.x IaaS Microsoft SQL database that you backed up on the source and restored on the target environment.
Authentication mode	■ **Windows** If you use the Windows authentication mode, the IaaS service user must have the SQL Server db_owner role. The same permissions apply when using SQL Server authentication mode. ■ **SQL Server** **SQL Server** opens the **Login name** and **Password** text boxes.
Login name	Login name of the SQL Server user with the db_owner role for the cloned IaaS Microsoft SQL database.
Password	Password for the SQL Server user with the db_owner role for the cloned IaaS Microsoft SQL database.
Original encryption key	Original encryption key that you retrieve from the source environment. See "Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.
New passphrase	A series of words used to generate a new encryption key. You use this passphrase each time you install a new IaaS component in the target vRealize Automation environment.

7 Click **Validate**.

The page displays the validation progress.

- If all the items validate successfully, go to step 8.

- If an item fails to validate, inspect the error message and the validation log file on the IaaS nodes. For log file locations, see "Migration Log Locations," on page 52. Click **Edit Settings** and edit the problem item. Go to step 7.

8 Click **Migrate**.

The page displays the migration progress.

- If migration is successful, the page displays information about the Software Agent post-migration update.
- If migration is unsuccessful, inspect the migration log files on the virtual appliance and the IaaS nodes. For log file locations, see "Migration Log Locations," on page 52.

Finish these steps before you restart migration.

a Revert your target vRealize Automation environment to the state you captured when you took a snapshot before migration.

b Restore your target IaaS Microsoft SQL database using the backup of the source IaaS database.

What to do next

Chapter 5, "Post-Migration Tasks," on page 29.

Migrate vRealize Automation Source Data to a vRealize Automation 7.3 High-Availability Environment

You can migrate your current vRealize Automation environment to a new installation of vRealize Automation 7.3 configured as a high-availability environment.

Prerequisites

- "Gather Information Required for Migration," on page 13.
- "Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.
- "Add Each Tenant from the Source vRealize Automation Environment to the Target Environment," on page 16.
- "Create an Administrator for Each Added Tenant," on page 17.
- "Synchronize Users and Groups for an Active Directory Link Before Migration to a High-Availability Environment," on page 20.
- "Manually Clone the Source vRealize Automation IaaS Microsoft SQL Database," on page 22.
- "Snapshot the Target vRealize Automation Environment," on page 22.

Procedure

1 In your target vRealize Automation 7.3 environment, open a browser and go to the master vRealize Automation appliance management console at https://*vra-va-hostname.domain.name*:5480.

2 Log in with the user name **root** and the password you entered when you deployed the appliance.

3 Select **vRA Settings > Migration**.

4 Enter the information for the source vRealize Automation appliance.

Option	Description
Host name	The host name for the source vRealize Automation appliance.
Root username	`root`
Root password	The root password that you entered when you deployed the source vRealize Automation appliance.

5 Enter the information for the target vRealize Automation appliance.

Option	Description
Root username	`root`
Root password	The root password that you entered when you deployed the target vRealize Automation appliance.
Default tenant	The default tenant you created when you configured single sign-on in the Installation wizard, usually vsphere.local.
Administrator username	The tenant administrator user name that you entered when you deployed the target vRealize Automation appliance. Change existing value if necessary.
Administrator password	The password that you entered for the default tenant administrator when you deployed the target vRealize Automation appliance.

6 Enter the information for the target IaaS database server.

Option	Description
Database server	The location of the Microsoft SQL Server instance where the restored vRealize Automation IaaS Microsoft SQL database resides. If a named instance and a non-default port are used, enter in *SERVER,PORT\INSTANCE-NAME* format.
Cloned database name	Name of the source vRealize Automation 6.2.x or 7.x IaaS Microsoft SQL database that you backed up on the source and restored on the target environment.
Authentication mode	■ **Windows** If you use the Windows authentication mode, the IaaS service user must have the SQL Server db_owner role. The same permissions apply when using SQL Server authentication mode. ■ **SQL Server** **SQL Server** opens the **Login name** and **Password** text boxes.
Login name	Login name of the SQL Server user with the db_owner role for the cloned IaaS Microsoft SQL database.
Password	Password for the SQL Server user with the db_owner role for the cloned IaaS Microsoft SQL database.
Original encryption key	Original encryption key that you retrieve from the source environment. See "Obtain the Encryption Key from the Source vRealize Automation Environment," on page 15.
New passphrase	A series of words used to generate a new encryption key. You use this passphrase each time you install a new IaaS component in the target vRealize Automation environment.

7 Click **Validate**.

The page displays the validation progress.

- If all the items validate successfully, go to step 8.

- If an item fails to validate, inspect the error message and the validation log file on the IaaS nodes. For log file locations, see "Migration Log Locations," on page 52. Click **Edit Settings** and edit the problem item. Go to step 7.

8 Click **Migrate**.

The page displays the migration progress.

- If migration is successful, the page displays information about the Software Agent post-migration update.
- If migration is unsuccessful, inspect the migration log files on the virtual appliance and the IaaS nodes. For log file locations, see "Migration Log Locations," on page 52.

Finish these steps before you restart migration.

a Revert your target vRealize Automation environment to the state you captured when you took a snapshot before migration.

b Restore your target IaaS Microsoft SQL database using the backup of the source IaaS database.

What to do next

Chapter 5, "Post-Migration Tasks," on page 29.

Post-Migration Tasks 5

After you migrate vRealize Automation, perform the post-migration tasks that pertain to your situation.

Note After you migrate the identity stores, users of vRealize Code Stream must manually reassign vRealize Code Stream roles.

This chapter includes the following topics:

Add Tenant and IaaS Administrators from the Source vRealize Automation 6.2.x Environment

You must delete and restore the vRealize Automation 6.2.x tenant administrators in each tenant after migration.

Perform the following procedure for each tenant in the target vRealize Automation console.

Note If you migrate from a vRealize Automation 7.x environment, you do not need to perform this procedure.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Log in to the target vRealize Automation console.

  a

Procedure

1 Select **Administration > Tenants**.

2 Click a tenant name.

3 Click **Administrators**.

4 Make a list of each tenant administrator name and user name.

5 Point to each administrator and click the delete icon (Delete) until you delete all administrators.

6 Click **Finish**.

7 On the Tenants page, click the tenant name again.

8 Click **Administrators**.

9 Enter the name of each user that you deleted in the appropriate search box and press Enter.

10 Click the name of the appropriate user from the search returns to add the user back as an administrator.

   When you finish, the list of tenant administrators administrators looks the same as the list of administrators you deleted.

11 Click **Finish**.

Run Test Connection and Verify Migrated Endpoints

Migrating to vRealize Automation 7.3 makes changes to endpoints in the target environment.

After you migrate to vRealize Automation 7.3, you must use the **Test Connection** action for all applicable endpoints. You might also need to make adjustments to some migrated endpoints. For more information, see *Considerations When Working With Upgraded or Migrated Endpoints* in *Configuring vRealize Automation*.

The default security setting for upgraded or migrated endpoints is to not accept untrusted certificates.

After upgrading or migrating from pre-vRealize Automation 7.3, if you were using untrusted certificates you must perform the following steps for all vSphere and NSX endpoints to enable certificate validation. Otherwise, the endpoint operations fail with certificate errors. For more information see VMware Knowledge Base articles *Endpoint communication is broken after upgrade to vRA 7.3 (2150230)* at http://kb.vmware.com/kb/2150230 and *How to download and install vCenter Server root certificates to avoid Web Browser certificate warnings (2108294)* at http://kb.vmware.com/kb/2108294.

1 After upgrade or migration, log in to the vRealize Automation vSphere agent machine and restart your vSphere agents by using the **Services** tab.

   Migration might not restart all agents, so manually restart them if needed.

2 Wait for at least one ping report to finish. It takes a minute or two for a ping report to finish.

3 When the vSphere agents have started data collection, log in to vRealize Automation as an IaaS administrator.

4 Click **Infrastructure > Endpoints > Endpoints**.

5 Edit a vSphere endpoint and click **Test Connection**.

6 If a certificate prompt appears, click **OK** to accept the certificate.

   If a certificate prompt does not appear, the certificate might currently be correctly stored in a trusted root authority of the Windows machine hosting service for the endpoint, for example as a proxy agent machine or DEM machine.

7 Click **OK** to apply the certificate acceptance and save the endpoint.

8 Repeat this procedure for each vSphere endpoint.

9 Repeat this procedure for each NSX endpoint.

If the **Test Connection** action is successful but some data collection or provisioning operations fail, you can install the same certificate on all the agent machines that serve the endpoint and on all DEM machines. Alternatively, you can uninstall the certificate from existing machines and repeat the above procedure for the failing endpoint.

Run NSX Network and Security Inventory Data Collection in Your Target vRealize Automation 7.3 Environment

After you migrate, you must run NSX Network and Security Inventory data collection in the target VMware vRealize ™ Automation 7.3 environment.

This data collection is necessary for the load balancer reconfigure action to work in vRealize Automation 7.3 for 7.1 and 7.2 deployments.

Note Do not perform this data collection if you migrated from vRealize Automation 6.2.x to 7.3.

Prerequisites

- "Run NSX Network and Security Inventory Data Collection in the Source vRealize Automation Environment," on page 21 .
- Successfully migrate to vRealize Automation 7.3.

Procedure

- Run NSX Network and Security Inventory data collection in your target vRealize Automation environment before you migrate to vRealize Automation 7.3. See *Start Endpoint Data Collection Manually* in *Managing vRealize Automation*.

Reconfigure Load Balancers After Migration to a High-Availability Environment

When you migrate to a high-availability environment, you must reconfigure each load balancer after you finish migration.

Prerequisites

"Migrate vRealize Automation Source Data to a vRealize Automation 7.3 High-Availability Environment," on page 25.

Procedure

- To restore the original health check settings so replica nodes can accept incoming traffic, configure the load balancers for these items.
  - vRealize Automation appliance.
  - IaaS Web Server that hosts the Model Manager.
  - Manager Service.

Migrating an External vRealize Orchestrator Server to vRealize Automation 7.3

You can migrate your existing external vRealize Orchestrator server to a vRealize Orchestrator instance embedded in vRealize Automation.

You can deploy vRealize Orchestrator as an external server instance and configure vRealize Automation to work with that external instance, or you can configure and use the vRealize Orchestrator server that is included in the vRealize Automation appliance.

VMware recommends that you migrate your external vRealize Orchestrator to the Orchestrator server that is built into vRealize Automation. The migration from an external to embedded Orchestrator provides the following benefits:

- Reduces the total cost of ownership.
- Simplifies the deployment model.
- Improves the operational efficiency.

Note Consider using the external vRealize Orchestrator in the following cases:

- Multiple tenants in the vRealize Automation environment
- Geographically dispersed environment
- Workload handling
- Use of specific plug-ins, such as the Site Recovery Manager plug-in

Control Center Differences Between External and Embedded Orchestrator

Some of the menu items that are available in Control Center of an external vRealize Orchestrator are not included in the default Control Center view of an embedded Orchestrator instance.

In Control Center of the embedded Orchestrator server, a few options are hidden by default.

Menu Item	Details
Licensing	The embedded Orchestrator is preconfigured to use vRealize Automation as a license provider.
Export/Import Configuration	The embedded Orchestrator configuration is included in the exported vRealize Automation components.
Configure Database	The embedded Orchestrator uses the database that is used by vRealize Automation.
Customer Experience Improvement Program	You can join the Customer Experience Improvement Program (CEIP) from the vRealize Automation appliance management interface. See *The Customer Experience Improvement Program* in *Managing vRealize Automation*.

Another options that are hidden from the default Control Center view are the **Host address** text box and the **UNREGISTER** button on the **Configure Authentication Provider** page.

Note To see the full set of Control Center options in vRealize Orchestrator that is built into vRealize Automation, you must access the advanced Orchestrator Management page at https://*vra-va-hostname.domain.name_or_load_balancer_address*:8283/vco-controlcenter/#/?advanced and click the F5 button on the keyboard to refresh the page.

Migrate an External vRealize Orchestrator 6.x on Windows to vRealize Automation 7.3

After you upgrade your vRealize Automation from version 6.x to version 7.3, you can migrate your existing external Orchestrator 6.x installed on Windows to the Orchestrator server that is built into vRealize Automation 7.3.

Note If you have a distributed vRealize Automation environment with multiple vRealize Automation appliance nodes, perform the migration procedure only on the primary vRealize Automation node.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Stop the Orchestrator server service on the external Orchestrator.
- Back up the database, including the database schema, of teh external Orchestrator server.

Procedure

1 Download the migration tool from the target Orchestrator server.

   a Log in to the vRealize Automation appliance over SSH as **`root`**.

   b Download the `migration-tool.zip` archive that is located in the `/var/lib/vco/downloads` directory.

2 Export the Orchestrator configuration from the source Orchestrator server.

   a Set the `PATH` environment variable by pointing it to the `bin` folder of the Java JRE installed with Orchestrator.

   b Upload the migration tool to the Windows server, on which the external Orchestrator is installed.

   c Extract the downloaded archive in the Orchestrator install folder.

   The default path to the Orchestrator install folder in a Windows-based installation is `C:\Program Files\VMware\Orchestrator`.

d Run the Windows command prompt as administrator and navigate to the `bin` folder in the Orchestrator install folder.

By default, the path to the `bin` folder is `C:\Program Files\VMware\Orchestrator\migration-cli\bin`.

e Run the `export` command from the command line.

```
C:\Program Files\VMware\Orchestrator\migration-cli\bin\vro-migrate.bat export
```

This command combines the VMware vRealize Orchestrator configuration files and plug-ins into an export archive.

The archive is created in the same folder as the `migration-cli` folder.

3 Migrate the exported configuration to the Orchestrator server that is built into vRealize Automation 7.3.

a Upload the exported configuration file to the `/usr/lib/vco/tools/configuration-cli/bin` directory on the vRealize Automation appliance.

b Under the `/usr/lib/vco/tools/configuration-cli/bin` directory, change the ownership of the exported Orchestrator configuration file.

```
chown vco:vco orchestrator-config-export-orchestrator_ip_address-date_hour.zip
```

c Import the Orchestrator configuration file to the built-in vRealize Orchestrator server, by running the `vro-configure` script with the `import` command.

```
./vro-configure.sh import --skipDatabaseSettings --skipLicense --skipSettings --skipSslCertificate --notForceImportPlugins --notRemoveMissingPlugins --skipTrustStore --path orchestrator-config-export-orchestrator_appliance_ip-date_hour.zip
```

4 Migrate the database to the internal PostgreSQL database, by running the `vro-configure` script with the `db-migrate` command.

```
./vro-configure.sh db-migrate --sourceJdbcUrl JDBC_connection_URL --sourceDbUsername database_user --sourceDbPassword database_user_password
```

Note Enclose passwords that contain special characters in quotation marks.

The *JDBC_connection_URL* depends on the type of database that you use.

```
PostgreSQL: jdbc:postgresql://host:port/database_name
```

```
MSSQL: jdbc:jtds:sqlserver://host:port/database_name\;domain=domain
```

```
Oracle: jdbc:oracle:thin:@host:port:database
```

5 If you migrated vRealize Automation instead of upgrading it, delete the trusted Single Sign-On certificates from the database of the embedded Orchestrator instance.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DELETE FROM vmo_keystore WHERE id='cakeystore-id';"
```

You successfully migrated an external vRealize Orchestrator 6.x installed on Windows to a vRealize Orchestrator instance embedded in vRealize Automation 7.3.

What to do next

Set up the built-in vRealize Orchestrator server. See "Configure the Built-In vRealize Orchestrator Server," on page 38.

Migrate an External vRealize Orchestrator 6.x Virtual Appliance to vRealize Automation 7.3

After you upgrade your vRealize Automation from version 6.x to version 7.3, you can migrate your existing external Orchestrator 6.x Virtual Appliance to the Orchestrator server that is built into vRealize Automation 7.3.

Note If you have a distributed vRealize Automation environment with multiple vRealize Automation appliance nodes, perform the migration procedure only on the primary vRealize Automation node.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Stop the Orchestrator server service on the external Orchestrator.
- Back up the database, including the database schema, of teh external Orchestrator server.

Procedure

1 Download the migration tool from the target Orchestrator server to the source Orchestrator.

   a Log in to the vRealize Orchestrator 6.x Virtual Appliance over SSH as **root**.

   b Under the `/var/lib/vco` directory, run the `scp` command to download the `migration-tool.zip` archive.

   ```
   scp root@vra-va-hostname.domain.name:/var/lib/vco/downloads/migration-tool.zip ./
   ```

   c Run the `unzip` command to extract the migration tool archive.

   ```
   unzip migration-tool.zip
   ```

2 Export the Orchestrator configuration from the source Orchestrator server.

   a In the `/var/lib/vco/migration-cli/bin` directory, run the `export` command.

   ```
   ./vro-migrate.sh export
   ```

   This command combines the VMware vRealize Orchestrator configuration files and plug-ins into an export archive.

   An archive with file name `orchestrator-config-export-`*`orchestrator_ip_address-date_hour`*`.zip` is created in the `/var/lib/vco` folder.

3 Migrate the exported configuration to the Orchestrator server that is built into vRealize Automation 7.3.

   a Log in to the vRealize Automation appliance over SSH as **root**.

   b Under the `/usr/lib/vco/tools/configuration-cli/bin` directory, run the `scp` command to download the exported configuration archive.

   ```
   scp root@orchestrator_ip_or_DNS_name:/var/lib/vco/orchestrator-config-export-
   orchestrator_ip_address-date_hour.zip ./
   ```

   c Change the ownership of the exported Orchestrator configuration file.

   ```
   chown vco:vco orchestrator-config-export-orchestrator_ip_address-date_hour.zip
   ```

d Stop the Orchestrator server service and the Control Center service of the built-in vRealize Orchestrator server.

```
service vco-server stop && service vco-configurator stop
```

e Import the Orchestrator configuration file to the built-in vRealize Orchestrator server, by running the `vro-configure` script with the `import` command.

```
./vro-configure.sh import --skipDatabaseSettings --skipLicense --skipSettings --
skipSslCertificate --notForceImportPlugins --notRemoveMissingPlugins --skipTrustStore --
path orchestrator-config-export-orchestrator_appliance_ip-date_hour.zip
```

4 If the external Orchestrator server from which you want to migrate uses the built-in PostgreSQL database, edit the database configuration files.

a In the `/storage/db/pgsql/data/postgresql.conf` file, uncomment the `listen_addresses` line.

b Set the values of `listen_addresses` to a wildcard (*).

```
listen_addresses ='*'
```

c Append a line to the `/storage/db/pgsql/data/pg_hba.conf` file.

```
host all all vra-va-hostname.domain.name/32 md5
```

Note The `pg_hba.conf` file requires using a CIDR prefix format instead on an IP address and a subnet mask.

d Restart the PostgreSQL server service.

```
service postgresql restart
```

5 Migrate the database to the internal PostgreSQL database, by running the `vro-configure` script with the `db-migrate` command.

```
./vro-configure.sh db-migrate --sourceJdbcUrl JDBC_connection_URL --sourceDbUsername
database_user --sourceDbPassword database_user_password
```

Note Enclose passwords that contain special characters in quotation marks.

The *JDBC_connection_URL* depends on the type of database that you use.

```
PostgreSQL: jdbc:postgresql://host:port/database_name
```

```
MSSQL: jdbc:jtds:sqlserver://host:port/database_name\;domain=domain
```

```
Oracle: jdbc:oracle:thin:@host:port:database
```

6 If you migrated vRealize Automation instead of upgrading it, delete the trusted Single Sign-On certificates from the database of the embedded Orchestrator instance.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DELETE FROM
vmo_keystore WHERE id='cakeystore-id';"
```

7 Revert to the default configuration of the `postgresql.conf` and the `pg_hba.conf` file.

a Restart the PostgreSQL server service.

You successfully migrated an external vRealize Orchestrator 6.x Virtual Appliance to a vRealize Orchestrator instance embedded in vRealize Automation 7.3.

What to do next

Set up the built-in vRealize Orchestrator server. See "Configure the Built-In vRealize Orchestrator Server," on page 38.

Migrate an External vRealize Orchestrator 7.x to vRealize Automation 7.3

You can export the configuration from your existing external Orchestrator instance and import it to the Orchestrator server that is built into vRealize Automation.

Note If you have multiple vRealize Automation appliance nodes, perform the migration procedure only on the primary vRealize Automation node.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Stop the Orchestrator server service on the external Orchestrator.
- Back up the database, including the database schema, of teh external Orchestrator server.

Procedure

1 Export the configuration from the external Orchestrator server.

   a Log in to Control Center of the external Orchestrator server as **root** or as an **administrator**, depending on the source version.

   b Stop the Orchestrator server service from the **Startup Options** page to prevent unwanted changes to the database.

   c Go to the **Export/Import Configuration** page.

   d On the **Export Configuration** page, select **Export server configuration**, **Bundle plug-ins** and **Export plug-in configurations**.

2 Migrate the exported configuration into the embedded Orchestrator instance.

   a Upload the exported Orchestrator configuration file to the `/usr/lib/vco/tools/configuration-cli/bin` directory of the vRealize Automation appliance.

   b Log in to the vRealize Automation appliance over SSH as **root**.

   c Stop the Orchestrator server service and the Control Center service of the built-in vRealize Orchestrator server.

   ```
   service vco-server stop && service vco-configurator stop
   ```

   d Navigate to the `/usr/lib/vco/tools/configuration-cli/bin` directory.

   e Change the ownership of the exported Orchestrator configuration file.

   ```
   chown vco:vco orchestrator-config-export-orchestrator_appliance_ip-date_hour.zip
   ```

   f Import the Orchestrator configuration file to the built-in vRealize Orchestrator server, by running the `vro-configure` script with the `import` command.

   ```
   ./vro-configure.sh import --skipDatabaseSettings --skipLicense --skipSettings --
   skipSslCertificate --notForceImportPlugins --notRemoveMissingPlugins --skipTrustStore --
   path orchestrator-config-export-orchestrator_appliance_ip-date_hour.zip
   ```

3 If the external Orchestrator server from which you want to migrate uses the built-in PostgreSQL database, edit the database configuration files.

   a In the `/storage/db/pgsql/data/postgresql.conf` file, uncomment the `listen_addresses` line.

   b Set the values of `listen_addresses` to a wildcard (*).

   ```
   listen_addresses ='*'
   ```

c Append a line to the `/storage/db/pgsql/data/pg_hba.conf` file.

```
host all all vra-va-hostname.domain.name/32 md5
```

Note The `pg_hba.conf` file requires using a CIDR prefix format instead on an IP address and a subnet mask.

d Restart the PostgreSQL server service.

```
service postgresql restart
```

4 Migrate the database to the internal PostgreSQL database, by running the `vro-configure` script with the `db-migrate` command.

```
./vro-configure.sh db-migrate --sourceJdbcUrl JDBC_connection_URL --sourceDbUsername
database_user --sourceDbPassword database_user_password
```

Note Enclose passwords that contain special characters in quotation marks.

The *JDBC_connection_URL* depends on the type of database that you use.

```
PostgreSQL: jdbc:postgresql://host:port/database_name
```

```
MSSQL: jdbc:jtds:sqlserver://host:port/database_name\;domain=domain
```

```
Oracle: jdbc:oracle:thin:@host:port:database
```

5 If you migrated vRealize Automation instead of upgrading it, delete the trusted Single Sign-On certificates from the database of the embedded Orchestrator instance.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DELETE FROM
vmo_keystore WHERE id='cakeystore-id';"
```

6 Revert to the default configuration of the `postgresql.conf` and the `pg_hba.conf` file.

a Restart the PostgreSQL server service.

You successfully migrated an external Orchestrator server instance to a vRealize Orchestrator instance embedded in vRealize Automation.

What to do next

Set up the built-in vRealize Orchestrator server. See "Configure the Built-In vRealize Orchestrator Server," on page 38.

Configure the Built-In vRealize Orchestrator Server

After you export the configuration of an external Orchestrator server and import it to vRealize Automation 7.3, you must configure the Orchestrator server that is built into vRealize Automation.

Prerequisites

Migrate the configuration from the external to the internal vRealize Orchestrator.

Procedure

1 Log in to the vRealize Automation appliance over SSH as **root**.

2 Start the Control Center service and the Orchestrator server service of the built-in vRealize Orchestrator server.

```
service vco-configurator start && service vco-server start
```

3 Log in to Control Center of the built-in Orchestrator server as an **administrator**.

Note If you migrate from an external vRealize Orchestrator 7.3 instance, skip to step 5.

4 Verify that Orchestrator is configured properly at the **Validate Configuration** page in Control Center.

5 If the external Orchestrator was configured to work in cluster mode, reconfigure the Orchestrator cluster in vRealize Automation.

a Go to the advanced **Orchestrator Cluster Management** page, at https://*vra-va-hostname.domain.name_or_load_balancer_address*:8283/vco-controlcenter/#/control-app/ha?remove-nodes.

Note If the **Remove** check boxes next the existing nodes in the cluster do not appear, you must refresh the browser page by clicking the F5 button on the keyboard.

b Select the check boxes next to the external Orchestrator nodes and click **Remove** to remove them from the cluster.

c To exit the advanced cluster management page, delete the `&remove–nodes` string from the URL and refresh the browser page by clicking the F5 button on the keyboard.

d At the **Validate Configuration** page in Control Center, verify that Orchestrator is configured properly.

6 (Optional) Under the **Package Signing Certificate** tab on the **Certificates** page, generate a new package signing certificate.

7 (Optional) Change the values for **Default tenant** and **Admin group** on the **Configure Authentication Provider** page.

8 Verify that the `vco–server` service appears as REGISTERED under the **Services** tab in the vRealize Automation appliance management console.

9 Select the `vco` services of the external Orchestrator server and click **Unregister**.

What to do next

- Import any certificates that were trusted in the external Orchestrator server to the trust store of the built-in Orchestrator.
- Join the vRealize Automation replica nodes to the vRealize Automation cluster to synchronize the Orchestrator configuration.

  For more information, see *Reconfigure the Target Embedded vRealize Orchestrator to Support High Availability* in *Installing or Upgrading vRealize Automation*.

  Note The vRealize Orchestrator instances are automatically clustered and available for use.

- Restart the `vco–configurator` service on all nodes in the cluster.
- Update the vRealize Orchestrator endpoint to point to the migrated built-in Orchestrator server.
- Add the vRealize Automation host and the IaaS host to the inventory of the vRealize Automation plug-in, by running the Add a vRA host and Add the IaaS host of a vRA host workflows.

Migrate the Embedded vRealize Orchestrator Server from vRealize Automation 7.2 to 7.3

You can migrate the vRealize Orchestrator server from your vRealize Automation 7.2 source environment to vRealize Automation 7.3 by performing these procedures.

Prerequisites

Successful migration to vRealize Automation 7.3.

Procedure

1 Temporarily Change the Configuration of the Source vRealize Automation Appliance on page 40

Before you migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3, you must temporarily change the configuration of the source vRealize Automation appliance.

2 Export the Configuration from the Embedded vRealize Orchestrator on the Source vRealize Automation Appliance on page 41

Before you migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3, you must export the configuration of the embedded source vRealize Orchestrator.

3 Import the Configuration and Database of the Embedded Source vRealize Orchestrator to the Embedded Target vRealize Orchestrator on page 42

Use this procedure to migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3.

4 Reconfigure the Target Embedded vRealize Orchestrator to Support High Availability on page 43

For a high-availability deployment, you must manually rejoin each target replica VMware vRealize ™ Automation appliance to the cluster to enable high-availability support for the embedded VMware vRealize ™ Orchestrator ™.

5 Restore the Configuration of the Source vRealize Automation Appliance on page 44

Use this procedure to restore the configuration of the source VMware vRealize ™ Automation appliance.

Temporarily Change the Configuration of the Source vRealize Automation Appliance

Before you migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3, you must temporarily change the configuration of the source vRealize Automation appliance.

Prerequisites

- For a minimal deployment, open an SSH console session to the source vRealize Automation appliance as **`root`** user.
- For a high-availability deployment, open an SSH console session to the master source vRealize Automation appliance as **`root`** user.

Procedure

1 Create a vro_migration user.

a Run this command to create a vro_migration user in the source PostgreSQL server. Before you run the command, replace {*VRO-MIGRATION-USER-PASSWORD*} with the vro_migration user password.

```
sudo –u postgres –i –– /opt/vmware/vpostgres/current/bin/psql vcac
–c "CREATE USER vro_migration WITH PASSWORD
'{VRO–MIGRATION–USER–PASSWORD}';"
```

b Run this command to grant the vro_migration user access to the tables in the vcac database.

```
sudo –u postgres –i –– /opt/vmware/vpostgres/current/bin/psql vcac
–c "GRANT SELECT ON ALL TABLES IN SCHEMA public TO vro_migration;"
```

2 Run this command to create a backup of the source PostgreSQL client authentication configuration file at `/storage/db/pgdata/pg_hba.conf`.

```
cp /storage/db/pgdata/pg_hba.conf /storage/db/pgdata/pg_hba.conf.bak
```

3 Run this command to modify the source PostgreSQL client authentication configuration file to grant vro_migration user remote access to vcac database from the target vRealize Automation appliance. Before you run the command, replace {*TARGET-VRA-APPLIANCE-IPV4-ADDRESS*} with the IP v4 address of the target vRealize Automation appliance.

```
echo "host vcac vro_migration {TARGET–VRA–APPLIANCE–IPV4–ADDRESS}/32 md5"
>> /storage/db/pgdata/pg_hba.conf
```

4 Run this command to restart the source PostgreSQL server.

```
service vpostgres restart
```

What to do next

"Export the Configuration from the Embedded vRealize Orchestrator on the Source vRealize Automation Appliance," on page 41

Export the Configuration from the Embedded vRealize Orchestrator on the Source vRealize Automation Appliance

Before you migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3, you must export the configuration of the embedded source vRealize Orchestrator.

Prerequisites

- For a minimal deployment, open an SSH console session to the source vRealize Automation appliance as **`root`** user.
- For a high-availability deployment, open an SSH console session to the master source vRealize Automation appliance as **`root`** user.

Procedure

◆ Run this command to export vRealize Orchestrator configuration as a ZIP file to `/tmp/vro–config.zip`.

```
/usr/lib/vco/tools/configuration–cli/bin/vro–configure.sh export ––skipLicense ––
path /tmp/vro–config.zip
```

What to do next

"Import the Configuration and Database of the Embedded Source vRealize Orchestrator to the Embedded Target vRealize Orchestrator," on page 42

Import the Configuration and Database of the Embedded Source vRealize Orchestrator to the Embedded Target vRealize Orchestrator

Use this procedure to migrate the VMware vRealize ™ Orchestrator ™ server from your VMware vRealize ™ Automation 7.2 source environment to vRealize Automation 7.3.

Prerequisites

- For a minimal deployment, open an SSH console session to the target vRealize Automation appliance as **root** user.
- For a high-availability deployment, open an SSH console session to the master target vRealize Automation appliance as **root** user.

Procedure

1 Run this command to stop the vRealize Orchestrator server service.

```
service vco-server stop
```

2 Run this command to stop the vRealize Orchestrator Control Center service.

```
service vco-configurator stop
```

For a high-availability deployment, stop the vRealize Orchestrator server service and vRealize Orchestrator Control Center service on the master vRealize Automation appliance and on each replica appliance.

3 Run this command to copy `vro-config.zip` from the source vRealize Automation appliance to the `/tmp` directory on the target vRealize Automation appliance. When prompted, enter the password for the source vRealize Automation appliance root user. Before you run the command, replace {*SOURCE-VRA-APPLIANCE-HOSTNAME*} with the fully qualified domain name of the source vRealize Automation appliance.

```
scp root@{SOURCE-VRA-APPLIANCE-HOSTNAME}:/tmp/vro-config.zip /tmp/vro-config.zip
```

4 Run this command to change the ownership of `/tmp/vro-config.zip`.

```
chown vco:vco /tmp/vro-config.zip
```

5 Run this command to import the configuration file to the embedded target vRealize Orchestrator server.

```
/usr/lib/vco/tools/configuration-cli/bin/vro-configure.sh import --skipDatabaseSettings --
skipLicense --skipSettings --skipSslCertificate --skipTrustStore --notForceImportPlugins --
notRemoveMissingPlugins --path /tmp/vro-config.zip
```

6 Run this command to migrate the source vRealize Orchestrator database to the PostgreSQL Server running on the target vRealize Automation appliance. Before you run the command, replace {*SOURCE-VRA-APPLIANCE-HOSTNAME*} with the fully qualified domain name of the source vRealize Automation appliance and {*VRO-MIGRATION-USER-PASSWORD*} with the vro_migration user password.

```
/usr/lib/vco/tools/configuration-cli/bin/vro-configure.sh db-migrate --sourceJdbcUrl
jdbc:postgresql://{SOURCE-VRA-APPLIANCE-HOSTNAME}:5432/vcac --sourceDbUsername vro_migration
--sourceDbPassword {VRO-MIGRATION-USER-PASSWORD}
```

7 Run this command to delete the old trusted certificates from the migrated database

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DELETE FROM
vmo_keystore WHERE id='cakeystore-id';"
```

8 Run this command to delete old vRealize Orchestrator nodes from the migrated database.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DELETE FROM
vmo_clustermember;"
```

9 Run this command to delete `vro-config.zip` from the `/tmp` directory.

```
rm -rf /tmp/vro-config.zip
```

10 Run this command to start the vRealize Orchestrator server service.

```
service vco-server start
```

For a high-availability deployment, start the vRealize Orchestrator server service only on the master vRealize Automation appliance.

What to do next

"Reconfigure the Target Embedded vRealize Orchestrator to Support High Availability," on page 43

Reconfigure the Target Embedded vRealize Orchestrator to Support High Availability

For a high-availability deployment, you must manually rejoin each target replica VMware vRealize ™ Automation appliance to the cluster to enable high-availability support for the embedded VMware vRealize ™ Orchestrator ™.

Prerequisites

Log in to the target replica vRealize Automation appliance management console.

1 Start a browser and open the target replica vRealize Automation management console using the fully qualified domain name (FQDN) of the target replica virtual appliance: https://*vra-va-hostname.domain.name*:5480.

2 Log in with the user name **root** and the password that you entered when you deployed the target replica vRealize Automation appliance.

Procedure

1 Select **vRA Settings > Cluster**.

2 In the **Leading Cluster Node** text box, enter the FQDN of the target master vRealize Automation appliance.

3 Enter the root password in the **Password** text box.

4 Click **Join Cluster**.

Continue past any certificate warnings. The system restarts services for the cluster.

5 Verify that the services are running.

a On the top tab bar, click **Services**.

b Click **Refresh** to monitor the progress of services startup.

What to do next

"Restore the Configuration of the Source vRealize Automation Appliance," on page 44

Restore the Configuration of the Source vRealize Automation Appliance

Use this procedure to restore the configuration of the source VMware vRealize ™ Automation appliance.

Prerequisites

- For a minimal deployment, open an SSH console session to the source vRealize Automation appliance as **root** user.
- For a high-availability deployment, open an SSH console session to the master source vRealize Automation appliance as **root** user.

Procedure

1 Run this command to delete `vro-config.zip` from the `/tmp` directory.

```
rm -rf /tmp/vro-config.zip
```

2 Run this command to revoke vco_migration user remote access to the vcac database by removing the previously added line from the source PostgreSQL client authentication configuration file.

```
sed -i '/vro_migration/d' /storage/db/pgdata/pg_hba.conf
```

3 Run this command to restart the PostgreSQL server.

```
service vpostgres restart
```

4 Delete vro_migration user from the source PostgreSQL database.

a Run this command to revoke vro_migration user access to the tables in the vcac database.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "REVOKE ALL
PRIVILEGES ON ALL TABLES IN SCHEMA public FROM vro_migration;"
```

b Run this command to remove vro_migration user from the source PostgreSQL server.

```
sudo -u postgres -i -- /opt/vmware/vpostgres/current/bin/psql vcac -c "DROP USER
vro_migration;"
```

Reconfigure the vRealize Automation Endpoint in the Target vRealize Orchestrator

Use the following procedure to reconfigure the vRealize Automation endpoint in the embedded target vRealize Orchestrator.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Connect to the target vRealize Orchestrator using the vRealize Orchestrator client. For information, see *Using the VMware vRealize Orchestrator Client* in the vRealize Orchestrator documentation.

Procedure

1 Select **Design** from the top drop-down menu.

2 Click **Inventory**.

3 Expand **vRealize Automation**.

4 Identify endpoints containing the fully qualified domain name (FQDN) of the source vRealize Automation appliance host or if you migrated from a high-availability deployment, the load-balanced host.

If you find endpoints containing the FQDN of the source vRealize Automation appliance host or if you migrated from a high-availability deployment, the load-balanced host	If you do not find endpoints containing the FQDN of the source vRealize Automation appliance host or if you migrated from a high-availability deployment, the load-balanced host
1 Click **Workflows**. 2 Click the expand button to select **Library > vRealize Automation > Configuration**. 3 Run the **Remove a vRA host** workflow for every endpoint containing the FQDN of the source vRealize Automation appliance host .	1 Click **Resources**. 2 Click the update icon on the top toolbar. 3 Click the expand button to select **Library > vCACCAFE > Configuration**. 4 Delete each resource that has a `URL` property containing the FQDN of the source vRealize Automation appliance host or if you migrated from a high-availability deployment., the load-balanced host.

5 Click **Workflows**.

6 Click the expand button to select **Library > vRealize Automation > Configuration**.

7 To add the target vRealize Automation appliance host or if you migrated to a high-availability deployment, the load-balanced host, run the **Add a vRA host using component registry** workflow.

Reconfigure the vRealize Automation Infrastructure Endpoint in the Target vRealize Orchestrator

Use the following procedure to reconfigure the vRealize Automation infrastructure endpoint in the embedded target vRealize Orchestrator.

Prerequisites

- Successful migration to vRealize Automation 7.3.
- Connect to the target vRealize Orchestrator using the vRealize Orchestrator client. For information, see *Using the VMware vRealize Orchestrator Client* in the vRealize Orchestrator documentation.

Procedure

1 Select **Design** from the top drop-down menu.

2 Click **Inventory**.

3 Expand **vRealize Automation Infrastructure**.

4 Identify endpoints containing the fully qualified domain name (FQDN) of the source vRealize Automation infrastructure host or if you migrated from a high-availability deployment, the load-balanced host.

If you find endpoints containing the FQDN of the source vRealize Automation infrastructure host or if you migrated from a high-availability deployment, the load-balanced host	If you do not find endpoints containing the FQDN of the source vRealize Automation infrastructure host or if you migrated from a high-availability deployment, the load-balanced host
1 Click **Workflows**. 2 Click the expand button to select **Library > vRealize Automation > Infrastructure Administration > Configuration**. 3 Run the **Remove an IaaS host** workflow for every endpoint containing the FQDN of the source vRealize Automation infrastructure host .	1 Click **Resources**. 2 Click the update icon on the top toolbar. 3 Click the expand button to select **Library > vCAC > Configuration**. 4 Delete each resource that has a `host` property containing the FQDN of the source vRealize Automation infrastructure host or if you migrated from a high-availability deployment., the load-balanced host.

5 Click **Workflows**.

6 Click the expand button to select **Library > vRealize Automation > Configuration**.

7 To add the target vRealize Automation infrastructure host, or if you migrated to a high-availability deployment load-balanced host, run the **Add the IaaS host of a vRA host** workflow.

Install vRealize Orchestrator Customization

You can run a workflow to install the customized state change workflow stubs and vRealize Orchestrator menu operation workflows.

For information, see *Install vRealize Orchestrator Customization* in *Life Cycle Extensibility*.

Prerequisites

Successful migration to vRealize Automation 7.3.

Reconfigure Embedded vRealize Orchestrator Infrastructure Endpoint in the Target vRealize Automation

When you migrate from a vRealize Automation 6.2.x environment, you must update the URL of the infrastructure endpoint that points to the target embedded vRealize Orchestrator server.

Prerequisites

- Successfully migrate to vRealize Automation 7.3.
- Log in to the target vRealize Automation console.
  a Open the vRealize Automation console using the fully qualified domain name of the target virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

    For a high-availability environment, open the console using the fully qualified domain name of the target virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

  b Log in as a IaaS administrator user.

Procedure

1 Select **Infrastructure > Endpoints > Endpoints**.

2 On the Endpoints page, select the vRealize Orchestrator endpoint, and click **Edit**.

3 In the Address text box, edit the vRealize Orchestrator endpoint URL.

- If you migrated to a minimal environment, replace the vRealize Orchestrator endpoint URL with `https://vra-va-hostname.domain.name:443/vco`.
- If you migrated to a high-availability environment, replace the vRealize Orchestrator endpoint URL with `https://vra-va-lb-hostname.domain.name:443/vco`.

4 Click **OK**.

5 Manually run a data collection on the vRealize Orchestrator endpoint.

a On the Endpoints page, select the vRealize Orchestrator endpoint.

b Select**Actions > Data Collection**.

Verify that the data collection is successful.

Reconfigure the Azure Endpoint in the Target vRealize Automation Environment

After migration, you must reconfigure your Microsoft Azure endpoint.

Perform this procedure for each Azure endpoint.

Prerequisites

- Successfully migrate to vRealize Automation 7.3.
- Log in to the target vRealize Automation console.

a Open the vRealize Automation console using the fully qualified domain name of the target virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

For a high-availability environment, open the console using the fully qualified domain name of the target virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

b Log in as a IaaS administrator user.

Procedure

1 Select **Administration > vRO Configuration > Endpoints**.

2 Select an Azure endpoint.

3 Click **Edit**.

4 Click **Details**.

5 In the **Client secret** text box, enter the original client secret.

6 Click **Finish**.

7 Repeat for each Azure endpoint.

Migrate vRealize Automation 6.2.x Automation Application Services to 7.3

You can use the VMware vRealize Application Services Migration Tool to migrate your existing application services blueprints and deployment profiles from VMware vRealize Application Services 6.2.x to vRealize Automation 7.3.

Prerequisites

Successful migration to vRealize Automation 7.3.

Procedure

- To download the VMware vRealize Application Services Migration Tool, complete these steps.
  a Click Download VMware vRealize Automation.
  b Select **Drivers & Tools > VMware vRealize Application Services Migration Tool**.

Update Software Agent on Existing Virtual Machines

After migration from vRealize Automation 7.2 to 7.3, the target vRealize Automation console cannot manage software components on existing virtual machines. Before the target console can manage software components on existing virtual machines, you must update the software agent on each virtual machine.

You use the vRealize Orchestrator client to perform these tasks:

- Import the downloaded Software Agent Post-Migration Update package to the source vRealize Orchestrator.
- Update the software agent on an existing virtual machine.
- Re-establish communication with the target vRealize Automation appliance

Note Updating software agents is an irreversible operation. After you do this update, you can no longer manage software components on existing virtual machines with the source vRealize Automation console.

Prerequisites

- Successful migration from source vRealize Automation 7.2 environment to target vRealize Automation 7.3 environment.
- Download the Software Agent Post-Migration Update package.
  a Open the target vRealize Automation appliance Guest and Software Agent Installers page using the target appliance fully qualified domain name: https://*vra-va-hostname.domain.name*/software/index.html.
  b Click **Software Agent Update workflow**.
- Connect to the target vRealize Orchestrator using the vRealize Orchestrator client. For information, see *Using the VMware vRealize Orchestrator Client* in the vRealize Orchestrator documentation.

Procedure

1. On the vRealize Orchestrator client, select **Run** from the top drop-down menu.
2. On the My Orchestrator page, click **Import package**.
3. Navigate to the directory where you downloaded the Software Agent Post-Migration Update package, `com.vmware.vra.sct.update.package`.
4. Select the package name and click **Open**.
5. Click **Import and trust provider**.
6. Click **Import selected elements**.

   The **Packages** tab opens showing the imported package.
7. Click the **Workflows** tab.
8. Click the expand button to select **Library > vRealize Automation > Migration > Software Agents**.
9. Double-click **Re-Parent Software Agents with Target vRealize Automation**.

   Run this workflow for each tenant in the source vRealize Automation environment.

10 To run the wizard, click the green **Start workflow** button at the top of the right pane.

11 Provide the requested information for the source vRealize Automation environment.

12 Provide the requested information for the target vRealize Automation environment.

This target environment information is provided on the target vRealize Automation management console Migration Status page.

- Virtual appliance IP address.
- Virtual appliance certificate.
- Software agent JAR SHA256 checksum.

13 Click **Submit**.

The workflow performs these tasks on the source vRealize Automation environment.

- Authenticates the user on the tenant to get an API token.
- Installs the software agent update scripts as new software components in the source vRealize Automation environment. System installs one software component for each supported operating system, Windows or Linux respectively.
- Obtains a list of running virtual machines with software agent installed.
- Updates the software agent by running the appropriate software agent update script on each virtual machine in the list.
- Uninstalls previously added software components from the source vRealize Automation environment.

Delete Original Target vRealize Automation IaaS Microsoft SQL Database

You can delete the original IaaS database after migration is complete.

Prerequisites

Successful migration to vRealize Automation 7.3.

Your migrated environment does not use the original vRealize Automation IaaS Microsoft SQL database that you created when you installed the target vRealize Automation 7.3 environment. You can safely delete this original IaaS database from the Microsoft SQL Server after you complete migration.

Update Data Center Location Menu Contents After Migration

After migration, you must add any missing custom data center locations to the **Location** drop-down menu.

After migration to vRealize Automation 7.3, the data center locations in the **Location** drop-down menu on the Compute Resources page revert to the default list. Although custom data center locations are missing, all compute resource configurations migrate successfully and the `Vrm.DataCenter.Location` property is not affected. You can still add custom data center locations to the **Location** menu.

Prerequisites

Migrate to vRealize Automation 7.3.

Procedure

- Add missing data center locations to the **Location** drop-down menu. See *Scenario: Add Datacenter Locations for Cross Region Deployments* in *Configuring vRealize Automation*.

Validate the Target vRealize Automation 7.3 Environment

You can verify that all data is migrated successfully to the target vRealize Automation 7.3 environment.

Prerequisites

- Migrate to vRealize Automation 7.3.
- Log in to the target vRealize Automation console.
  a Open the vRealize Automation console using the fully qualified domain name of the target virtual appliance: https://*vra-va-hostname.domain.name*/vcac.

    For a high-availability environment, open the console using the fully qualified domain name of the target virtual appliance load balancer: https://*vra-va-lb-hostname.domain.name*/vcac.

  b Log in with the tenant administrator user name and password.

Procedure

1 Select **Infrastructure > Managed Machines** and verify that all the managed virtual machines are present.

2 Click **Compute Resources**, select each endpoint, and click **Data Collection**, **Request now**, and **Refresh** to verify that the endpoints are working.

3 Click **Design**, and on the Blueprints page, verify the elements of each blueprint.

4 Click **XaaS** and verify the contents of **Custom Resources**, **Resource Mappings**, **XaaS Blueprints**, and **Resource Actions**.

5 Select **Administration > Catalog Management** and verify the contents of **Services**, **Catalog Items**, **Actions**, and **Entitlements**.

6 Select **Items > Deployments** and verify the details for the provisioned virtual machines.

7 On the Deployments page, select a provisioned, powered off, virtual machine and select **Actions > Power On**, click **Submit**, and click **OK**. Verify that the virtual machine powers on correctly.

8 Click **Catalog** and request a new catalog item.

9 On the **General** tab, enter the request information.

10 Click the Machine icon, accept all the default settings, click **Submit**, and click **OK**.

11 Verify that the request finishes successfully.

Troubleshooting Migration 6

Migration troubleshooting topics provide solutions to problems you might experience when you migrate vRealize Automation.

This chapter includes the following topics:

PostgreSQL Version Causes Error

A source vRealize Automation 6.2.x environment containing an updated PostgreSQL database blocks administrator access.

Problem

If an upgraded PostgreSQL database is used by vRealize Automation 6.2.x, an administrator must add an entry to the `pg_hba.conf` file that provides access to this database from vRealize Automation.

Solution

1 Open the `pg_hba.conf` file.

2 To grant access to this database, add the following entry.

host all *vcac-database-user vra-va-ip trust-method*

Some Virtual Machines Do Not Have a Deployment Created during Migration

Virtual machines in the missing state at the time of migration do not have a corresponding deployment created in the target environment.

Problem

If a virtual machine is in the missing state in the source environment during migration, a corresponding deployment is not created in the target environment.

Solution

- If a virtual machine goes out of the missing state after migration, you can import the virtual machine to the target deployment using bulk import.

Load Balancer Configuration Causes Timeout for Long-Running Operations

A load balancer can cause an unexpected connection termination.

Problem

Some load balancers have very short timeouts for keeping a connection alive during execution of an HTTP/HTTPS request. This short timeout can result in unexpected connection termination when migration performs long-running operations.

Solution

- Increase the timeout on the load balancer or update the load balancer DNS record to point to the appropriate active node for the duration of the migration. Once migration is complete, revert the load balancer DNS record.

Migration Log Locations

You can troubleshoot validation or migration problems by viewing the logs that record the migration process.

Table 6-1. Source vRealize Automation Appliance

Log	Location
Package creation log	`/var/log/vmware/vcac/migration-package.log`

Table 6-2. Target vRealize Automation Appliance

Log	Location
Migration log	`/var/log/vmware/vcac/migrate.log`
Migration execution log	`/var/log/vmware/vcac/mseq.migration.log`
Migration execution output log	`/var/log/vmware/vcac/mseq.migration.out.log`
Validation execution log	`/var/log/vmware/vcac/mseq.validation.log`
Validation execution output log	`/var/log/vmware/vcac/mseq.validation.out.log`

Table 6-3. Target vRealize Automation Infrastructure Nodes

Log	Location
Migration log	`C:\Program Files (x86)\VMware\vCAC\InstallLogs-YYYYMMDDHHMMXX\Migrate.log`
Validation log	`C:\Program Files (x86)\VMware\vCAC\InstallLogs-YYYYMMDDHHMMXX\Validate.log`

Index