



Red Hat Infrastructure Migration Solution 1.2

Infrastructure Migration Solution Guide

Migrating from VMware to Red Hat Virtualization or Red Hat OpenStack Platform

Red Hat Infrastructure Migration Solution 1.2 Infrastructure Migration Solution Guide

Migrating from VMware to Red Hat Virtualization or Red Hat OpenStack Platform

Red Hat IMS Documentation Team
ims-docs@redhat.com

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide describes how to migrate VMware virtual machines to Red Hat Virtualization 4.3 or Red Hat OpenStack Platform 13 or 14, using Red Hat CloudForms 4.7.6 or later.

Table of Contents

PREFACE	4
PART I. MIGRATING FROM VMWARE TO RED HAT VIRTUALIZATION	5
CHAPTER 1. PLANNING THE MIGRATION	6
1.1. QUESTIONS TO ASK BEFORE MIGRATION	7
1.2. RECOMMENDATIONS AND BEST PRACTICES	9
CHAPTER 2. PREPARING THE ENVIRONMENT FOR MIGRATION	11
2.1. PREPARING THE VMWARE ENVIRONMENT	11
2.1.1. Preparing the VMware network	11
2.1.2. Preparing the VMware virtual machines	11
2.1.3. Configuring a VMware hypervisor for more than ten concurrent migrations	12
2.2. PREPARING THE RED HAT VIRTUALIZATION ENVIRONMENT	12
2.2.1. Prerequisites	13
2.2.2. Installing and configuring Red Hat Virtualization 4.3	13
2.2.3. Installing and configuring CloudForms 4.7.6	14
2.2.4. Configuring the conversion hosts	15
Downloading the VMware Virtual Disk Development Kit	15
2.2.5. Verifying the conversion hosts in a browser	17
CHAPTER 3. MIGRATING THE VIRTUAL MACHINES	18
3.1. CREATING AN INFRASTRUCTURE MAPPING	18
3.2. CHECKING FOR MIGRATION PREREQUISITES	20
3.2.1. Creating a CSV file to add virtual machines to the migration plan	21
3.2.2. Adding Ansible playbooks to CloudForms for premigration and postmigration tasks	21
3.2.3. Creating a RHEL premigration playbook for RHEL/Linux source virtual machines	22
3.3. CREATING AND RUNNING A MIGRATION PLAN	22
3.3.1. Scheduling a saved migration plan	25
3.3.2. Viewing a migration plan in progress	25
3.3.3. Canceling a migration plan in progress	25
3.3.4. Retrying a failed migration plan	25
3.4. CHANGING THE MAXIMUM NUMBER OF CONCURRENT MIGRATIONS	26
CHAPTER 4. TROUBLESHOOTING	28
4.1. MIGRATION LOGS	28
4.1.1. Downloading the conversion host and playbook logs	28
4.1.2. Accessing the CloudForms migration log	29
4.2. INFRASTRUCTURE MAPPING ERRORS	29
4.3. MIGRATION PLAN ERRORS	29
4.4. IP ADDRESS ERRORS	29
4.5. ENVIRONMENT CONFIGURATION ERRORS	30
4.6. KNOWN ISSUES	30
APPENDIX A. CONFIGURING THE RED HAT VIRTUALIZATION ENVIRONMENT FOR SSH TRANSFORMATION	32
A.1. CONFIGURING THE VMWARE HYPERVISORS FOR SSH TRANSFORMATION	32
A.2. REINSTALLING IPA-CLIENT	32
A.3. CONFIGURING THE CONVERSION HOSTS	33
A.4. COPYING THE VMWARE SSH KEYS TO THE CONVERSION HOSTS	34
A.4.1. Copying keys collected during VMware hypervisor configuration	34
A.4.2. Copying keys collected with ssh-keyscan	35

PART II. MIGRATING FROM VMWARE TO RED HAT OPENSTACK PLATFORM	36
CHAPTER 5. PLANNING THE MIGRATION	37
5.1. QUESTIONS TO ASK BEFORE MIGRATION	39
5.2. RECOMMENDATIONS AND BEST PRACTICES	40
CHAPTER 6. PREPARING THE ENVIRONMENT FOR MIGRATION	42
6.1. PREPARING THE VMWARE ENVIRONMENT	42
6.1.1. Preparing the VMware network	42
6.1.2. Preparing the VMware virtual machines	42
6.1.3. Configuring a VMware hypervisor for more than ten concurrent migrations	43
6.2. PREPARING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT	43
6.2.1. Prerequisites	44
6.2.2. Installing and configuring Red Hat OpenStack Platform 13 or 14	44
6.2.3. Installing and configuring CloudForms 4.7.6	45
Viewing the cloud provider timeline in CloudForms	46
6.2.4. Creating the conversion hosts	46
6.2.5. Configuring the conversion hosts	47
Downloading the VMware Virtual Disk Development Kit	48
6.2.6. Verifying the conversion hosts in a browser	49
CHAPTER 7. MIGRATING THE VIRTUAL MACHINES	51
7.1. CREATING AN INFRASTRUCTURE MAPPING	51
7.2. CHECKING FOR MIGRATION PREREQUISITES	53
7.2.1. Creating a CSV file to add virtual machines to the migration plan	54
7.2.2. Adding Ansible playbooks to CloudForms for premigration and postmigration tasks	54
7.2.3. Creating a RHEL premigration playbook for RHEL/Linux source virtual machines	55
7.3. CREATING AND RUNNING A MIGRATION PLAN	55
7.3.1. Scheduling a saved migration plan	58
7.3.2. Viewing a migration plan in progress	58
7.3.3. Canceling a migration plan in progress	58
7.3.4. Retrying a failed migration plan	59
7.4. CHANGING THE MAXIMUM NUMBER OF CONCURRENT MIGRATIONS	59
CHAPTER 8. TROUBLESHOOTING	61
8.1. MIGRATION LOGS	61
8.1.1. Downloading the conversion host and playbook logs	61
8.1.2. Accessing the CloudForms migration log	62
8.2. INFRASTRUCTURE MAPPING ERRORS	62
8.3. MIGRATION PLAN ERRORS	62
8.4. IP ADDRESS ERRORS	62
8.5. ENVIRONMENT CONFIGURATION ERRORS	63
8.6. KNOWN ISSUES	63
APPENDIX B. CONFIGURING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT FOR SSH TRANSFORMATION	65
B.1. CONFIGURING THE VMWARE HYPERVISORS FOR SSH TRANSFORMATION	65
B.2. CONFIGURING THE CONVERSION HOSTS	65
B.3. COPYING THE VMWARE SSH KEYS TO THE CONVERSION HOSTS	67
B.3.1. Copying keys collected during VMware hypervisor configuration	67
B.3.2. Copying keys collected with ssh-keyscan	67

PREFACE

Red Hat Infrastructure Migration Solution (IMS) 1.2 enables you to migrate virtual machines from VMware 5.5 (and later) to Red Hat Virtualization or Red Hat OpenStack Platform, using Red Hat CloudForms:

Red Hat Virtualization 4.3

Red Hat Virtualization provides a virtualization platform built on Red Hat Enterprise Linux. You can manage your virtual infrastructure, including hosts, virtual machines, networks, storage, and users, from a centralized graphical user interface or with a REST API. See [Red Hat Virtualization 4.3](#) documentation for more information.

See [Part I, "Migrating from VMware to Red Hat Virtualization"](#) .

Red Hat OpenStack Platform 13 and 14

Red Hat OpenStack Platform provides a scaleable, fault-tolerant, private or public cloud based on Red Hat Enterprise Linux. See [Red Hat OpenStack Platform 13](#) or [Red Hat OpenStack Platform 14](#) documentation for more information.

See [Part II, "Migrating from VMware to Red Hat OpenStack Platform"](#) .

Red Hat CloudForms 4.7.6 or later

Red Hat CloudForms is the environment in which you perform the migration. Red Hat CloudForms is a unified set of management tools for use across virtualized, private cloud, and public cloud platforms. See [Red Hat CloudForms 4.7](#) documentation for more information.

PART I. MIGRATING FROM VMWARE TO RED HAT VIRTUALIZATION

The migration process involves the following tasks:

1. Planning the migration. See [Chapter 1, *Planning the migration*](#).
2. Preparing the VMware, Red Hat Virtualization, and CloudForms environments. See [Chapter 2, *Preparing the environment for migration*](#).
3. Migrating the virtual machines. See [Chapter 3, *Migrating the virtual machines*](#).
4. Troubleshooting, if necessary. See [Chapter 4, *Troubleshooting*](#).

CHAPTER 1. PLANNING THE MIGRATION

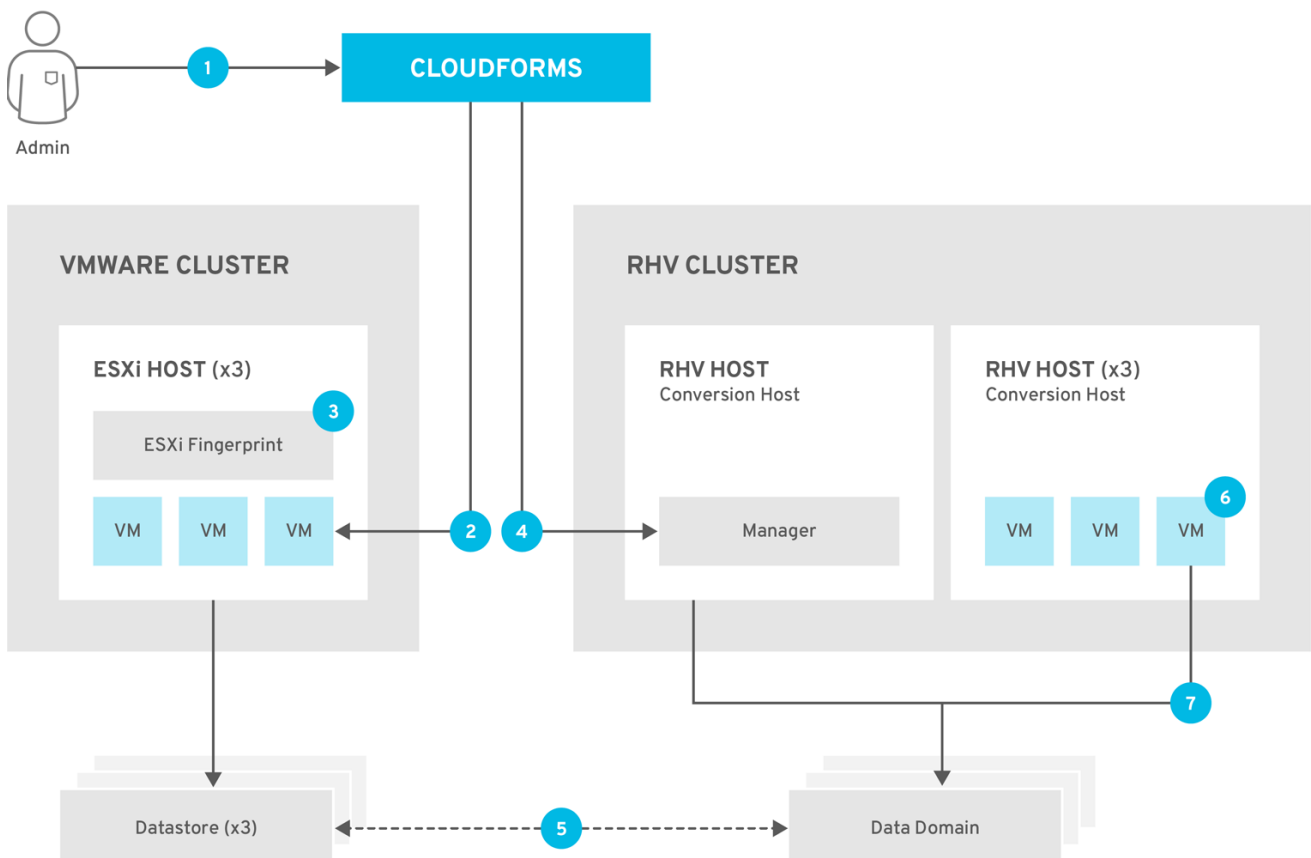
During the planning phase, you will formulate a specific migration goal, for example, "I want to migrate 2000 virtual machines, with 200 TB of data, in less than 6 months".

Review the following information to plan your migration:

- [Section 1.1, "Questions to ask before migration"](#) includes the following:
 - What operating systems are supported?
 - How long will the migration take?
 - How many conversion hosts do I need?
- [Section 1.2, "Recommendations and best practices"](#) provides guidelines on the following topics:
 - Scheduling the migration
 - Distributing the workload
 - Deploying the conversion hosts
 - Controlling the migration process

The migration workflow describes the migration process in greater detail.

Figure 1.1. VMware to Red Hat Virtualization migration workflow



CLOUDFORMS_474716_0718

- 1 You create and run a migration plan in CloudForms.
- 2 CloudForms uses the migration plan to locate the source virtual machines.
- 3 CloudForms captures the ESXi host fingerprint for authentication during the virtual machine conversion process.
- 4 Using the attributes defined for the Red Hat Virtualization environment, CloudForms initiates communication with the conversion hosts (Red Hat Virtualization hosts with **virt-v2v** and **virt-v2v-wrapper** installed).
- 5 **virt-v2v-wrapper** connects to the source datastore through the ESXi host. **virt-v2v** streams the source disks to the target data domain and converts the source disks.
- 6 **virt-v2v-wrapper** creates a target Red Hat Virtualization virtual machine, using the source virtual machine's metadata in order to maintain its attributes (tags, power state, MAC address, CPU count, memory, disks, and virtual machine name) after migration.
- 7 **virt-v2v** attaches the converted disks to the Red Hat Virtualization virtual machine. (The virtual machine's power state is the same as the source virtual machine's power state.)

The migration process is complete and the migration plan's status is displayed in CloudForms.

1.1. QUESTIONS TO ASK BEFORE MIGRATION

The following questions can help you to estimate the resources and time required for migration.

What am I migrating?

- Identify the VMware virtual machines that you will be migrating.

What is the maximum number of disks or virtual machines that I can migrate?

- There is no maximum number of disks or virtual machines that you can migrate. However, you may not want to migrate all your virtual machines at the same time, in order to minimize the impact on your users.



IMPORTANT

If you exceed the capabilities of your environment, the migrations will fail. This situation could affect existing applications running on virtual machines attached to the network and storage.

What operating systems can I migrate?

- You can migrate any guest operating system that is certified and supported for use with the Red Hat Virtualization Hypervisor. See [Certified Guest Operating Systems in Red Hat OpenStack Platform and Red Hat Virtualization](#).

What am I missing?

- Identify resource gaps, such as bandwidth, storage, licenses, or a suitable maintenance window, before you begin the migration.

What impact will the migration have on my users?

- Assess the effects the migration may have on a production environment. It may be possible to migrate your applications in phases, without downtime at the application layer, if the applications are distributed in a high-availability architecture.
- Check whether users will lose access to critical applications.

How long will the migration take?

There is no formula to estimate how long the actual migration will take. This is determined on a case-by-case basis.

The following example is provided as a guide:

- Duration of migration: 1:15:53 (hh:mm:ss)
- 10 virtual machines
- Total data migrated: 1000 GB
- Hardware:
 - Strong host (40 cores, 500 GB RAM)
 - Fast SSD XtremIO storage
 - Fibre Channel 8 interface for host-to-storage connection
 - 10 GbE network interface cards for all other connections

How many conversion hosts do I need?

The number of conversion hosts you create depends on the size of your migration. All the virtual machines in a migration plan are migrated at the same time, in parallel. The number of virtual machines that you can migrate simultaneously depends on your infrastructure capabilities. Each migration requires a certain amount of network bandwidth, I/O throughput, and processing power for the conversion process.

Multiple conversion hosts provide load-balancing and better performance, even for small migrations.

Conversion hosts are limited to a maximum of ten concurrent migrations, unless you change the default values.

You should test your environment thoroughly before the migration to determine how many migrations it can support without negative effects, for example, five conversion hosts, each running ten concurrent migrations.

Should I migrate my virtual machines with VDDK or SSH?

You can migrate your virtual machines with either the VMware Virtual Disk Development Kit (VDDK) or SSH. VDDK is the default because it is much faster than SSH and easier to configure.

VDDK is limited to 20 concurrent migrations per conversion host, because of network limitations, and 10 concurrent migrations per VMware hypervisor, unless you increase the hypervisor's NFC service memory.


If you cannot use VDDK, SSH transformation is a fallback option.

1.2. RECOMMENDATIONS AND BEST PRACTICES

The following recommendations will help to minimize the impact of the migration on your environment.

Scheduling the migration

- Schedule your migration carefully, to minimize the impact on your users.
- Prepare your users for downtime.



IMPORTANT

Currently, IMS supports only cold migration. Virtual machines are powered off gracefully as part of the migration process.

- Stagger the migration schedules.
- Move critical applications during maintenance windows.

Distributing the migration workload

- Create migration groups, so that you are not migrating all of your virtual machines at the same time, keeping in mind the following considerations:
 - How are the virtual machines grouped now?
 - Which virtual machines should be migrated together?
 - Which workloads or linked applications should be migrated together?
 - What applications must remain available?
- Consider which parts of the workload to migrate first:
 - Databases
 - Applications
 - Web servers
 - Load balancers

Deploying the conversion hosts

- Create a sufficient number of conversion hosts for your migration, with sufficient resources.
- Create multiple conversion hosts for load-balancing. The virtual machines in a migration plan are automatically distributed among the conversion hosts. This decreases the load on the conversion hosts and allows you to increase the concurrent migrations beyond the limits of a single conversion host.

Controlling the migration process

- Create multiple migration plans for finer control.
- Perform test migrations with different maximum numbers of concurrent migrations to assess the capabilities of your environment's infrastructure.

CHAPTER 2. PREPARING THE ENVIRONMENT FOR MIGRATION

Preparing your environment for migration involves the following tasks:

1. Preparing the VMware environment. See [Section 2.1, “Preparing the VMware environment”](#).
2. Preparing the target environment. See [Section 2.2, “Preparing the Red Hat Virtualization environment”](#).
3. Configuring the conversion hosts. See [Section 2.2.4, “Configuring the conversion hosts”](#).



NOTE

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix A, *Configuring the Red Hat Virtualization environment for SSH transformation*](#).

2.1. PREPARING THE VMWARE ENVIRONMENT

Preparing the VMware environment for migration involves the following tasks:

1. Extending the VMware network. See [Section 2.1.1, “Preparing the VMware network”](#).
2. Preparing the VMware virtual machines. See [Section 2.1.2, “Preparing the VMware virtual machines”](#).
3. (Optional) Configuring a VMware hypervisor for more than ten concurrent migrations from that hypervisor. See [Section 2.1.3, “Configuring a VMware hypervisor for more than ten concurrent migrations”](#).

2.1.1. Preparing the VMware network

Extend the VMware network to the Red Hat Virtualization environment.



IMPORTANT


- The network configuration must not be changed in any way during the migration.
- IP addresses, VLANs, and other network configuration must not be changed before or after migration because the conversion process preserves the source virtual machine MAC addresses.

2.1.2. Preparing the VMware virtual machines

Perform the following steps on each VMware virtual machine that you are migrating:

1. Install VMware Tools to capture IP addresses.
To download and install VMware Tools, see [VMware Workstation 5.0: Installing VMware Tools](#).
2. Unmount mounted ISO/CDROM disks.
3. Ensure that attached disks are not encrypted.

4. Ensure that each NIC has no more than one IPv4 and/or one IPv6 address.
5. Ensure that the virtual machine names contain only upper- or lower-case letters, numbers, underscores (`_`), hyphens (`-`), or periods (`.`).

**NOTE**

International characters and spaces are not permitted.

6. Ensure that the virtual machine names do not duplicate names of virtual machines in the Red Hat Virtualization environment.

If you are performing more than ten concurrent migrations from a VMware hypervisor, you must configure the hypervisor to support the additional connections. See [Section 2.1.3, "Configuring a VMware hypervisor for more than ten concurrent migrations"](#).

2.1.3. Configuring a VMware hypervisor for more than ten concurrent migrations

If you are performing more than ten concurrent migrations from a VMware hypervisor using VDDK transformation, the migration will fail because the hypervisor's NFC service memory buffer is limited to ten parallel connections. See [VMware vSphere 6.5 NFC session connection limits](#) and [Virt-v2v. VDDK: ESXi NFC service memory limits](#) for details.

You can increase the hypervisor's NFC service memory to enable additional connections for migrations.

Procedure

1. Log in to a VMware hypervisor.
2. Change the value of **maxMemory** to **1000000000** in **/etc/vmware/hostd/config.xml**:

```
<nfcsvc>
  <path>libnfcsvc.so</path>
  <enabled>true</enabled>
  <maxMemory>1000000000</maxMemory>
  <maxStreamMemory>10485760</maxStreamMemory>
</nfcsvc>
```

3. Restart **hostd**:

```
# /etc/init.d/hostd restart
```

You do not need to reboot the VMware hypervisor.

2.2. PREPARING THE RED HAT VIRTUALIZATION ENVIRONMENT

Preparing the Red Hat Virtualization environment involves the following key steps:

1. Installing and configuring Red Hat Virtualization 4.3. See [Section 2.2.2, "Installing and configuring Red Hat Virtualization 4.3"](#).
2. Installing and configuring CloudForms 4.7.6 or later. See [Section 2.2.3, "Installing and configuring CloudForms 4.7.6"](#).

**NOTE**

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix A, Configuring the Red Hat Virtualization environment for SSH transformation](#).

2.2.1. Prerequisites

The following prerequisites apply:

- BIOS settings of physical hosts adjusted for optimal performance (rather than power-saving), according to the vendor's recommendations
- C1E halt state disabled, if applicable
- Compatible software versions:

Table 2.1. Software compatibility

Software	IMS 1.1	IMS 1.2
VMware	5.5 or later	
Red Hat Virtualization	4.2.8	4.3.4 or later
CloudForms	4.7.0, with CFME 5.10.3	4.7.6 or later, with CFME 5.10.5
	CFME 5.10.4 does not support migration. You can use CFME 5.10.4 to manage the Red Hat Virtualization environment. Only the migration functionality is affected.	

2.2.2. Installing and configuring Red Hat Virtualization 4.3

Procedure

1. Install Red Hat Virtualization Manager 4.3.4 on the Manager machine. See [Installing the Red Hat Virtualization Manager](#) in the *Red Hat Virtualization Installation Guide*.
2. Install Red Hat Virtualization Host 4.3.4 or Red Hat Enterprise Linux 7.6 on the host machines. See [Installing Red Hat Virtualization Host](#) or [Installing Red Hat Enterprise Linux Hosts](#) in the *Red Hat Virtualization Installation Guide*.

**NOTE**

Some of these hosts will be deployed as conversion hosts. The number of conversion hosts depends on your migration size and infrastructure capabilities.

3. Enable the following ports in the conversion host network:
 - 22 - SSH
 - 443 - CloudForms, Red Hat Virtualization Manager, and VDDK


- 902 - CloudForms to VMware
 - 5480 - Conversion hosts to vCenter
For details, see [Ports used by Red Hat CloudForms Management Engine 5.1 and above](#) .
4. Create and attach data and ISO storage domains to the data center, ensuring that the data domains have sufficient space for the migrated virtual machines. See [Storage](#) in the *Red Hat Virtualization Administration Guide*.

**NOTE**

IMS only supports shared storage, such as NFS, iSCSI, or FCP. Local storage is not supported.

Although the ISO storage domain has been deprecated in RHV 4.3, it is required for migration.

5. Upload the VirtIO and RHV Guest Tools image files to the ISO storage domain. See [Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain](#) in the *Red Hat Virtualization Administration Guide*.
The VirtIO file name must include the version number (**virtio-win-version.iso**). The guest tools are required for migrating Windows virtual machines.
6. Optionally, you can create a MAC address pool that includes the MAC addresses of the VMware virtual machines to be migrated. See [Creating MAC Address Pools](#) in the *Red Hat Virtualization Administration Guide*.

**IMPORTANT**

If the Red Hat Virtualization MAC address pool range overlaps the VMware MAC address range, you must ensure that the MAC addresses of the migrating virtual machines do not duplicate those of existing virtual machines. Otherwise, the migration will fail.

If you do not create a MAC address pool, the migrated virtual machines will not have MAC addresses in the same range as virtual machines created in Red Hat Virtualization.

2.2.3. Installing and configuring CloudForms 4.7.6

CAUTION

Removing or changing a provider will cause errors in the infrastructure mappings and migration plans.

Procedure

1. Install Red Hat CloudForms 4.7.6 or later on the Manager machine. See [Installing Red Hat CloudForms on Red Hat Virtualization](#).

**NOTE**

CFME 5.10.4 does not support migration.

2. Add VMware to CloudForms as a provider. See [Adding a VMware vCenter Provider](#) .
3. Add Red Hat Virtualization to CloudForms as a provider. [Adding a Red Hat Virtualization Provider](#) in *Red Hat CloudForms: Managing Providers* .

Configure your conversion hosts with CloudForms to use the VMware Virtual Disk Development Kit (VDDK).



NOTE

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix A, Configuring the Red Hat Virtualization environment for SSH transformation](#) .

2.2.4. Configuring the conversion hosts

Prerequisite

- Download and save the VMware Virtual Disk Development Kit.

Downloading the VMware Virtual Disk Development Kit

Procedure


1. Navigate to [VMware Documentation](#).
2. Click **VMware SDK & API Product Documentation** to expand.
3. Click **VMware Virtual Disk Development Kit (VDDK)**
4. Click **Latest Releases** and select the latest VDDK release.
5. Click **Download SDKs** to download the VDDK archive file.
6. Save the VDDK archive file in an HTTP-accessible location and record its path.

Procedure

1. In CloudForms, click **Compute** → **Migration** → **Migration Settings**.
2. In the **Conversion Hosts** tab, click **Configure Conversion Host**. The **Configure Conversion Host** wizard is displayed.

The screenshot shows a wizard window titled "Configure Conversion Host". At the top, there is a progress bar with four steps: "Location" (1), "Host(s)" (2), "Authentication" (3), and "Results" (4). The "Location" step is highlighted with a blue circle. Below the progress bar, there is a dropdown menu labeled "* Provider Type" with the text "<Choose>". At the bottom right of the window, there are three buttons: "Cancel", "< Back", and "Next >".

3. In the **Location** screen, add the provider details:
 - a. Select a **Provider Type**.
 - b. Select a **Provider**.
 - c. Select a **Cluster** and click **Next**.
4. In the **Host(s)** screen, select one or more hosts from the dropdown list and click **Next**.
5. In the **Authentication** screen, add the conversion hosts' SSH key and the transformation method:
 - a. Click **Browse** to browse to the SSH private key or paste it in the **Conversion host SSH private key** field.
The Manager deploys a private SSH key on the conversion hosts in order to send commands and run playbooks. The default key file is **/etc/pki/ovirt-engine/keys/engine_id_rsa** on the Manager machine.
 - b. Select **VDDK** as the **Transformation method**.
 - c. Enter the path of the VDDK package in the **VDDK library path** field and click **Validate**. Wait for validation to complete.
 - d. Click **Configure**.
6. In the **Results** screen, wait for the conversion host configuration to finish and click **Close**.
The configured conversion hosts and status information, including error messages, appear in the **Configured Conversion Hosts** list.

If an error occurs, you can download a host's log by clicking the **More Actions** icon () and selecting **Download Log**.

You can click the **Retry** button if the conversion host configuration failed for reasons unconnected with your environment.

(Optional) Click **Remove** to remove the configuration from a configured conversion host.

You can migrate your virtual machines.

(Optional) You can verify the conversion hosts in a browser. See [Section 2.2.5, "Verifying the conversion hosts in a browser"](#).

2.2.5. Verifying the conversion hosts in a browser

You can verify your conversion hosts in a browser by using the CloudForms API:

1. In the address bar of a browser, enter the following:

```
https://CloudForms_FQDN/api/conversion_hosts 1
```

- 1** **CloudForms_FQDN** is the FQDN of the CloudForms machine.

A log-in screen is displayed.

2. Enter your CloudForms **Username** and **Password** and click **Sign in**.

The conversion hosts and their IDs are displayed in JSON format:

```
{ "name": "conversion_hosts", "count": 3, "subcount": 3, "pages": 1, "resources":
  [ { "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000001" },
    { "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000002" },
    { "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000003" } ], "actions":
  [ { "name": "create", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts" },
    { "name": "edit", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts" },
    { "name": "delete", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts" } ], "links": { "self": "https://cloudforms.example.com/api/conversion_hosts?offset=0", "first": "https://cloudforms.example.com/api/conversion_hosts?offset=0", "last": "https://cloudforms.example.com/api/conversion_hosts?offset=0" } }
```

CHAPTER 3. MIGRATING THE VIRTUAL MACHINES

Migrating the virtual machines involves the following key tasks:

1. Mapping the resources of your VMware and Red Hat Virtualization environments. See [Section 3.1, "Creating an infrastructure mapping"](#).
2. Checking for migration conditions with prerequisites. See [Section 3.2, "Checking for migration prerequisites"](#).
3. Creating and running a migration plan. See [Section 3.3, "Creating and running a migration plan"](#).

Optionally, you can change the maximum number of concurrent migrations for conversion hosts or providers to control the migration process.

3.1. CREATING AN INFRASTRUCTURE MAPPING

The infrastructure mapping maps the resources of your VMware and Red Hat Virtualization environments.




IMPORTANT

If you add or remove providers or provider objects from an infrastructure mapping, the mapping will have missing resource errors. You must delete the infrastructure mapping and create a new one.

Procedure

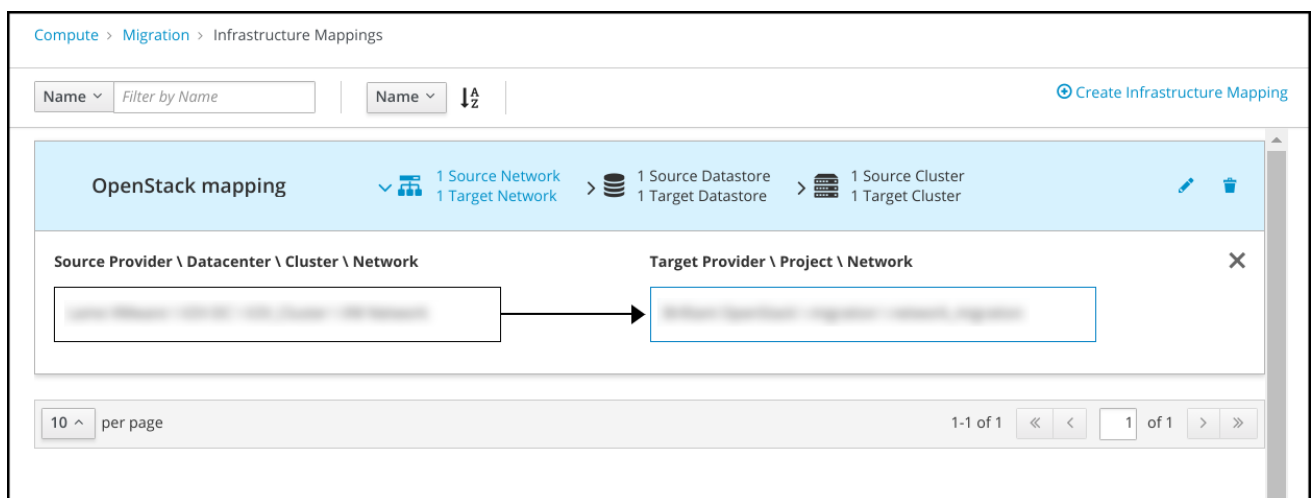
1. Click **Compute** → **Migration** → **Infrastructure Mappings**.
2. Click **Create Infrastructure Mapping** The **Create Infrastructure Mapping** wizard is displayed.

3. In the **General** screen, add the infrastructure mapping details:
 - a. Enter the infrastructure mapping **Name** and (optional) **Description**.
 - b. Select the **Target Provider**.
 - c. Click **Next**.
4. In the **Map Compute** screen, map the source and target clusters:
 - a. Select a **Source Provider \ Datacenter \ Cluster** and a **Target Provider \ Datacenter \ Cluster**.
 If the target cluster does not contain a conversion host, a warning icon () appears. You can create and save an infrastructure mapping, but you must configure the conversion hosts before running a migration plan.
 - b. Click **Add Mapping**. You can map additional clusters.
 - c. Click **Next**.
5. In the **Map Storage** screen, map the source and target storage:
 - a. Select a **Source Provider \ Datacenter \ Datastore** and **Target Datastores**.
 - b. Click **Add Mapping**. You can map additional datastores.
 - c. Click **Next**.

6. In the **Map Networks** screen, map the source and target networks:
 - a. Select a source cluster from the drop-down list.
 - b. Select one or more networks from **Source Provider \ Datacenter \ Network** and **Target Project \ Network**.
 - c. Click **Add Mapping**. You can map the networks of additional source clusters.
 - d. Click **Create**.
7. In the **Results** screen, click **Close**. The infrastructure mapping is saved in **Compute → Migration → Infrastructure Mappings**.

You can click an infrastructure mapping element to view its details:

Infrastructure Mappings list



After you have created an infrastructure mapping, check for migration prerequisites.

If these conditions do not apply, you can create a migration plan. See [Section 3.3, “Creating and running a migration plan”](#).

3.2. CHECKING FOR MIGRATION PREREQUISITES

Check your migration for the following conditions, which have prerequisites:

You are migrating previously migrated virtual machines

Add previously migrated machines to the migration plan with a CSV file. A CSV file is also recommended for large migrations.

See [Section 3.2.1, “Creating a CSV file to add virtual machines to the migration plan”](#).

You are using Ansible playbooks for premigration/postmigration tasks

Create an Ansible repository and add credentials and playbooks to CloudForms.

See [Section 3.2.2, “Adding Ansible playbooks to CloudForms for premigration and postmigration tasks”](#).

You are migrating virtual machines running RHEL or other Linux operating system

Create a RHEL premigration playbook to preserve IP addresses and select this playbook when you create a migration plan.

See [Section 3.2.2, “Adding Ansible playbooks to CloudForms for premigration and postmigration tasks”](#) and [Section 3.2.3, “Creating a **RHEL premigration** playbook for RHEL/Linux source virtual machines”](#).

3.2.1. Creating a CSV file to add virtual machines to the migration plan

If you are migrating virtual machines that were migrated in the past, you should create a CSV file to add the virtual machines to the migration plan, because the migration plan cannot discover them automatically.



NOTE

A CSV file is recommended for large migrations because it is faster than manually selecting individual virtual machines.

Table 3.1. CSV file fields

Field	Comments
Name	Virtual machine name. Required
Host	Optional. Only required if virtual machines have identical Name fields.
Provider	Optional. Only required if virtual machines have identical Name and Host fields.

CSV file example

```
Name,Host,Provider
vm01,host1,vSphere3
vm02,host1,vSphere3
vm03,host1,vSphere3
```

3.2.2. Adding Ansible playbooks to CloudForms for premigration and postmigration tasks

You can add Ansible playbooks to CloudForms to perform automated premigration and postmigration tasks on specific virtual machines, for example:

- Removing webserver from a load-balancing pool before migration and returning them to the pool after migration
- Running **fstrim** after migration to reduce the space required by virtual machines migrating to Red Hat OpenStack Platform with Ceph storage

Procedure

1. Enable the **Embedded Ansible** server role in CloudForms. See [Enabling the Embedded Ansible Server Role](#) in *Red Hat CloudForms: Managing Providers*.
2. Add an Ansible playbook repository. See [Adding a Playbook Repository](#) in *Red Hat CloudForms: Managing Providers*.

3. Add the credentials of each virtual machine that you are migrating. See [Credentials](#) in *Red Hat CloudForms: Managing Providers*.
4. Add your playbook as an Ansible service catalog item. See [Creating an Ansible Playbook Service Catalog Item](#) in *Red Hat CloudForms: Provisioning Virtual Machines and Instances*.

You will select the playbooks and the virtual machines on which they run in the **Advanced Options** screen when you create the migration plan.

3.2.3. Creating a RHEL premigration playbook for RHEL/Linux source virtual machines

If you are migrating virtual machines running RHEL or other Linux operating system, you can create a **RHEL premigration** playbook to ensure that the IP addresses are accessible after migration. The **RHEL premigration** playbook calls the Ansible **ims.rhel_premigration** role.

To install the role with Ansible Galaxy, see [ims_rhel_pre_migration](#). This role is not included in the IMS installation.

The **ims.rhel_premigration** role performs the following tasks on the VMware virtual machines:

- Preserves the static IP address configuration by creating udev rules to associate the virtual machine's MAC address with its interface name
- Installs the Red Hat Virtualization guest agent. The guest agent reports the new virtual machine's IP address and installed applications to the Manager.



NOTE

The **ims.rhel_premigration** role assumes that either the **rhel-6-server-rpms** or the **rhel-7-server-rpms** repository is enabled in the source virtual machine when it installs **qemu-guest-agent**. If you have disabled the repository, re-enable it in the RHEL premigration playbook.

RHEL premigration playbook example

```
---
- hosts: all
  roles:
    - role: ims.rhel_pre_migration
```

You can create a migration plan. See [Section 3.3, "Creating and running a migration plan"](#).

3.3. CREATING AND RUNNING A MIGRATION PLAN

Before you perform a large migration, you should perform several test migrations with different maximum numbers of concurrent migrations for your conversion hosts or providers. This will enable you to assess the capabilities of your environment's infrastructure.

You can create and run a migration plan in CloudForms with the following options:

- Running a migration plan immediately
- [Scheduling a migration plan to run in the future](#)

- [Viewing a migration plan in progress](#)
- [Canceling a migration plan in progress](#)
- [Retrying a failed migration plan](#)



NOTE

A CSV file is optional, but recommended, for large migrations because it is faster than manually selecting each virtual machine.

Procedure

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click **Create Migration Plan**. The **Create Migration Plan** wizard is displayed.

Create Migration Plan

General VMs Advanced Options Schedule Results

1 2 3 4 5

Infrastructure Mapping * <Choose>

Name * _____
Maximum length is 24 characters.

Description _____
Maximum length is 128 characters.

Select VMs

- Choose from a list of VMs discovered in the selected infrastructure mapping
- Import a CSV file with a list of VMs to be migrated

Cancel < Back Next >

3. In the **General** screen, add the details of the migration plan:
 - a. Select an infrastructure mapping from the drop-down list.
 - b. Enter the migration plan **Name** and (optional) **Description**.
 - c. Select a virtual machine discovery method:
 - **Choose from a list of VMs discovered in the selected infrastructure mapping**

If the virtual machines cannot be discovered, check that the source datastores and networks in the infrastructure mapping are correct.

- **Import a CSV file with a list of VMs to be migrated**

A CSV file is required for previously migrated source virtual machines and recommended for large migrations.

d. Click **Next**.

4. In the **VMs** screen, select the virtual machines for migration:

- If you selected **Choose from a list of VMs discovered in the selected infrastructure mapping**, select the virtual machines for migration.

You can search for virtual machines by **VM Name**, **Data Center**, **Cluster**, and **Folder**.

- If you selected **Import a CSV file with a list of VMs to be migrated**

a. Click **Import**.

b. Browse to the CSV file and click **Open**.

If the virtual machines cannot be added to the migration plan, check the CSV file format and fields for errors.



NOTE

If the **Create Migration Plan** wizard freezes, refresh the web page, check the CSV file for errors (for example, virtual machines with duplicate **Name** fields and no other fields to distinguish them), and create a new migration plan.

c. Click **Next**.

5. In the **Advanced Options** screen, select the playbook service options:

a. Select a premigration and/or postmigration playbook service from the dropdown lists.

b. Select the virtual machines on which to run the playbook services.

c. Click **Next**.

6. In the **Schedule** screen, select a schedule option and click **Create**:

- **Save migration plan to run later**

The migration plan is saved in **Migration Plans Not Started** and will not run unless you schedule it or click **Migrate** to run the scheduled migration plan immediately.

- **Start migration immediately**

The migration plan may take some time to complete. Progress bars indicate the amount of transferred data, the number of migrated virtual machines, and the elapsed time. See [Section 3.3.2, "Viewing a migration plan in progress"](#) for details.

7. In the **Results** screen, click **Close**.

When the migration plan has finished, click **Migration Plans Complete** to view the status of the migration plan. The completed migration plan shows the status of the migrated virtual machines.

In the migration plans list, you can click the **More Actions** icon () to archive, edit, or delete a migration plan.

3.3.1. Scheduling a saved migration plan

To schedule a saved migration plan to run in the future:

1. Click **Migration Plans Not Started**
2. Click the **Schedule** button of a migration plan.
3. In the **Schedule Migration Plan** window, select a date and time and click **Schedule**. The plan's status is **Migration Scheduled** with the date and time.

3.3.2. Viewing a migration plan in progress

To view the progress of a migration plan:

1. Click **Migration Plans in Progress**
2. Click a migration plan name to view its details, including the status of the migrating virtual machines.



NOTE

The counter in **Compute** → **Migration** → **Migration Plans** may be a few seconds ahead of the counter in the migration plan details view. This is because the **Migration Plans** counter displays the total time for running the migration plan, while the details counter displays the time for migrating the virtual machines.

3.3.3. Canceling a migration plan in progress

To cancel a migration plan in progress:

1. Click **Migration Plans in Progress**
2. Select a migration plan and click **Cancel Migration**.
3. Click **Cancel Migrations** to confirm the cancellation. The canceled migration appears in **Migration Plans Complete** with a red **x** indicating that the plan did not complete successfully.

3.3.4. Retrying a failed migration plan

To retry a migration plan that failed because of external circumstances (for example, power outage):

1. Delete all objects created by the failed migration plan:
 - Delete newly created virtual machines to avoid name conflicts with migrating VMware virtual machines.
 - Delete converted disks to free up space.
2. Click **Compute** → **Migration** → **Migration Plans**.
3. Click **Migration Plans Complete**.

4. Click the **Retry** button beside the failed migration plan.

3.4. CHANGING THE MAXIMUM NUMBER OF CONCURRENT MIGRATIONS

You can change the maximum number of concurrent migrations for conversion hosts or providers to control the impact of the migration process on your infrastructure.

The provider setting has priority over the conversion host setting. For example, if the maximum number of concurrent migrations is **20** for a provider and **3** for five conversion hosts, the maximum number of concurrent migrations is **20**, not **15** (**5** conversion hosts **x** **3** concurrent migrations per host).

An increase in the maximum number of concurrent migrations affects all migration plans immediately. Virtual machines that are queued to migrate will migrate in greater numbers.

A decrease maximum number of concurrent migrations affects only future migration plans. Migration plans that are in progress will use the limit that was set when the plan was created.

Changing the maximum number of concurrent migrations for all conversion hosts or providers

CAUTION

If you are using VDDK transformation, the number of concurrent migrations must not exceed **20**. Otherwise, network overload will cause the migration to fail.

1. Click **Compute** → **Migration** → **Migration Settings**.
2. In the **Migration Throttling** tab, select a value for **Maximum concurrent migrations per conversion host** or **Maximum concurrent migrations per provider** and click **Apply**.
The value of **Maximum concurrent migrations per conversion host** is constrained so that it cannot be greater than **Maximum concurrent migrations per provider**.



NOTE

In the current release, the **Maximum concurrent migrations per conversion host** interface control does not work.

Changing the maximum number of concurrent migrations for a specific conversion host

1. Enter the following command:

```
# curl -sk -u username:password_
https://CloudForms_FQDN/api/conversion_hosts/conversion_host_id -X POST -d {"action": "edit",
"resource": {"max_concurrent_tasks": 15}} 1 2 3 4
```

- 1 **username** and **password** are the username and password for CloudForms.
- 2 **CloudForms_FQDN** is the FQDN of the CloudForms machine.
- 3 To obtain a **conversion_host_id**, enter the following command:

```
# curl -sk -u username:password https://CloudForms_FQDN/api/conversion_hosts/
```

- 4 **max_concurrent_tasks** is the maximum number of concurrent migrations. The default is **10**.

CHAPTER 4. TROUBLESHOOTING

To identify errors, you can review the migration logs. See [Section 4.1, “Migration logs”](#).

You can check these common issues and mistakes:

- Infrastructure mapping errors. See [Section 4.2, “Infrastructure mapping errors”](#).
- Migration plan errors. See [Section 4.3, “Migration plan errors”](#).
- IP address errors. See [Section 4.4, “IP address errors”](#).
- Environment configuration errors. See [Section 4.5, “Environment configuration errors”](#).

[Section 4.6, “Known Issues”](#) provides information about issues that will be addressed in a future release.

4.1. MIGRATION LOGS

You can check the following logs to identify the cause of a migration error:


- Conversion host logs. See [Section 4.1.1, “Downloading the conversion host and playbook logs”](#).
- Playbook logs, if you are using premigration or postmigration Ansible playbooks. See [Section 4.1.1, “Downloading the conversion host and playbook logs”](#).
- CloudForms migration log. See [Section 4.1.2, “Accessing the CloudForms migration log”](#).

4.1.1. Downloading the conversion host and playbook logs

You can download the conversion host and playbooks logs in CloudForms.

When disk migration starts, two logs are created in the conversion host:

- **virt-v2v**: Debug output from **virt-v2v** itself. This log tracks the core of the virtual machine migration process, including **libguestfs** traces and disk migration details.
- **virt-v2v-wrapper**: Log of the daemonizing wrapper for **virt-v2v**. This log traces the orchestration of the virtual machine conversion on the conversion host, including disk migration percentages and **virt-v2v** error reporting.



IMPORTANT

If you open a [Red Hat Support call](#), you will need to submit both the **virt-v2v** log and **virt-v2v-wrapper** log for analysis and troubleshooting.

To access the conversion host logs and the playbook logs:

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click a completed migration plan to view its details.
3. Click **Download Log** of a virtual machine and select a log from the dropdown list:
 - **Premigration log** This option only appears if a premigration playbook is used.

- **Migration log** The **Migration log** is the **virt-v2v** log.
- **Virt-v2v-wrapper log**
- **Postmigration log** This option only appears if a postmigration playbook is used.

4.1.2. Accessing the CloudForms migration log

1. Log into the CloudForms machine using SSH.
2. The migration log is **/var/www/miq/vmdb/log/automation.log**.

4.2. INFRASTRUCTURE MAPPING ERRORS

- **Networks missing**, **Datastores missing**, and **Clusters missing** error messages: If you create an infrastructure mapping and then change a provider or refresh the Red Hat Virtualization hosts, the provider's object IDs change. Delete the infrastructure mapping and create a new one.

4.3. MIGRATION PLAN ERRORS

- If the virtual machines are being migrated for the first time and are not discovered by the migration plan, check the source datastores and networks in the infrastructure mapping.
- If the virtual machines have been migrated in the past, they cannot be discovered by the migration plan. Use a CSV file to add the virtual machines to the migration plan.
- If the virtual machines cannot be added to the migration plan with a CSV file, check the CSV file format and fields. Create a new migration plan with the updated CSV file.
- **Create Migration Plan** wizard hangs while importing a CSV file. This error is caused by an invalid CVS file (for example, virtual machines with a duplicate **Name** field and no **Host/Provider** field to distinguish them, or with a duplicate **Name** field and duplicate **Host/Provider** fields). Refresh the web page.
- **Denied State** error (IMS 1.1). If a migration plan fails immediately and the migration plan displays a **Denied State** error message, check that you have created and configured the conversion hosts correctly. Cancel the migration plan and run it again.
- **Unable to migrate VMs because no conversion host was configured at the time of the attempted migration. See the product documentation for information on configuring conversion hosts.** (IMS 1.2) You can create and save a migration plan whose infrastructure mapping does not contain conversion hosts, but you cannot run the migration plan without conversion hosts. Cancel the migration plan, create the conversion hosts, and run the migration plan again.

4.4. IP ADDRESS ERRORS

- If the IP address of a migrated RHEL (or other Linux-based operating system) virtual machine is not accessible, you must create a **RHEL premigration** playbook and add it to the migration plan.
- If a migrated virtual machine does not have an IP address:
 - Check that you installed VMware Tools on the VMware virtual machine before migration.

4.5. ENVIRONMENT CONFIGURATION ERRORS

VMware

- A VMware virtual machine cannot be migrated if it has any of the following conditions:
 - Mounted ISO/CDROM disk
 - Encrypted disk
 - Invalid name, containing spaces or special characters
- If you are performing more than ten concurrent migrations from a single VMware hypervisor, you must increase the hypervisor's NFC service memory.

Red Hat Virtualization

- Name conflict: VMware virtual machine has the same name as a Red Hat Virtualization virtual machine.
- MAC address conflict: VMware virtual machine has the same MAC address as a Red Hat Virtualization virtual machine in a MAC address pool.
- SSH transformation only:
 - If you are using SSSD with single sign-on, you must reinstall **ipa-client** without OpenSSH.
 - Check that the conversion host does not have an existing private SSH key in `/var/lib/vdsm/.ssh/id_rsa`. Conversion host configuration does not overwrite old SSH keys. They must be deleted manually.
 - Check that you enabled SSH access on the VMware hypervisors and correctly configured your conversion hosts for SSH transformation.

4.6. KNOWN ISSUES

The following issues will be addressed in a future release:

- [BZ#1666799: Canceling migration does not stop creating volume, instance, and network port on OpenStack Platform or VMs on Red Hat Virtualization](#)
- [BZ#1678385: Virtual machine with name containing spaces \(*rhel 7*\) fails to migrate using SSH and VDDK transformation](#)
- [BZ#1699343: Migration plan CSV import validation does not work if file contains empty/archived/orphan/invalid VM name](#)
- [BZ#1698761: "Maximum concurrent migrations per conversion host" interface control does not work](#)
- [BZ#716283: Migrating virtual machines are not distributed correctly among the conversion hosts](#)
- [BZ#1726939: Run the preflight check of migration task before waiting for a conversion host](#) . Currently, the preflight check that monitors the migration is performed after a conversion host is assigned to the task. As a result, the total volume of the Datastores reported in **Migration**

Plans In Progress reflects the total volume of the virtual machines that are currently migrating, not the total volume of the migration plan. When all the virtual machines have started to migrate, the correct value of the total volume is displayed.

- BZ#666799: [Canceling migration does not stop creating virtual machines on RHV](#) . If you cancel a migration, you must delete migrated virtual machines and disks in the Administration Portal.
- BZ#1669176: [Refreshing the hosts causes the network\(s\) and datastore to disappear from infrastructure mappings](#)
- CloudForms CFME 5.10.4 does not support migration. Use the following versions for migration:
 - Red Hat Virtualization 4.2: CFME 5.10.3
 - Red Hat Virtualization 4.3: CFME 5.10.5

**NOTE**

You can use CFME 5.10.4 to manage Red Hat Virtualization 4.3. Only the migration functionality is affected.

APPENDIX A. CONFIGURING THE RED HAT VIRTUALIZATION ENVIRONMENT FOR SSH TRANSFORMATION

You can configure your environment for SSH transformation if you cannot use VDDK.

Configuring your environment involves the following additional steps:

1. Configuring the VMware hypervisors. See [Section A.1, "Configuring the VMware hypervisors for SSH transformation"](#).
2. (Optional) Reinstalling **ipa-client** if you are configuring your conversion hosts for SSSD with single sign-on. See [Section A.2, "Reinstalling ipa-client"](#).
3. Configuring the conversion hosts. See [Section A.3, "Configuring the conversion hosts"](#).
4. Copying the VMware keys to the conversion hosts. See [Section A.4, "Copying the VMware SSH keys to the conversion hosts"](#).

A.1. CONFIGURING THE VMWARE HYPERVISORS FOR SSH TRANSFORMATION

For SSH transformation, you must configure the VMware hypervisors for passwordless access by sharing a conversion host's public SSH key with the hypervisors.



IMPORTANT

A single SSH key pair is recommended because the key pair is used only for virtual machine conversion and it simplifies conversion host management.

If you wish to use a dedicated SSH key pair for each conversion host, you can copy the public key of each conversion host to all the VMware hypervisors.

Procedure

1. Enable SSH access on each VMware hypervisor. For instructions, navigate to [VMware vSphere Documentation](#) and enter **Enable ESXi Shell and SSH Access with the Direct Console User Interface** in the **Search** field.

You can collect the SSH public keys of the VMware hypervisors at this stage, to copy to the conversion hosts.

2. Generate an SSH key pair without a passphrase:

```
# ssh-keygen -N "
```

3. Copy the public SSH key to **/etc/ssh/keys-root/authorized_keys** on **each** VMware hypervisor. You will use the private SSH key to configure the conversion hosts.

A.2. REINSTALLING IPA-CLIENT

If you are using SSH transformation and configuring your conversion hosts for SSSD with single sign-on, you must reinstall **ipa-client** without the OpenSSH client. Otherwise, SSH will fail for the **vdsms** user. See [BZ#1544379: ipa-client-install changes system-wide SSH configuration](#) for more information. This

issue cannot be resolved by modifying the configuration file because the file is restored during upgrades.

1. Log in to the Manager machine using SSH.
2. Uninstall **ipa-client**:

```
# ipa-client-install --uninstall
```

3. Reinstall **ipa-client** without OpenSSH:

```
# ipa-client-install --no-ssh
```

A.3. CONFIGURING THE CONVERSION HOSTS

Procedure

1. In CloudForms, click **Compute** → **Migration** → **Migration Settings**.
2. In the **Conversion Hosts** tab, click **Configure Conversion Host**. The **Configure Conversion Host** wizard is displayed.

The screenshot shows the 'Configure Conversion Host' wizard interface. At the top, the title bar reads 'Configure Conversion Host' with a close button. Below the title bar is a progress indicator with four steps: 'Location' (1), 'Host(s)' (2), 'Authentication' (3), and 'Results' (4). The 'Location' step is highlighted with a blue circle. Below the progress bar, there is a field labeled '* Provider Type' with a dropdown menu showing '<Choose>'. At the bottom right, there are three buttons: 'Cancel', '< Back', and 'Next >'.

3. In the **Location** screen, add the provider details:
 - a. Select a **Provider Type**.
 - b. Select a **Provider**.


- c. Select a **Cluster**.
 - d. Click **Next**.
4. In the **Host(s)** screen, select one or more hosts from the dropdown list and click **Next**.
 5. In the **Authentication** screen, add the conversion hosts' SSH key and the transformation method:
 - a. Click **Browse** to browse to the SSH private key or paste it in the **Conversion host SSH private key** field.
The Manager deploys a private SSH key on the conversion hosts in order to send commands and run playbooks. The default key file is **/etc/pki/ovirt-engine/keys/engine_id_rsa** on the Manager machine.
 - b. Select **SSH** as the **Transformation method**.
 - c. Click **Browse** to browse to the SSH private key you created for enabling SSH access on the VMware hypervisors or paste it in the **VMware hypervisors SSH private key** field.



IMPORTANT

If a host already has an SSH private key, you must delete the key manually in **/var/lib/vdsm/.ssh/id_rsa** before configuring it as a conversion host. Conversion host configuration does not overwrite existing keys.

- d. Click **Configure**.
6. In the **Results** screen, wait for the conversion host configuration to finish and click **Close**.
The configured conversion hosts and status information, including error messages, appear in the **Configured Conversion Hosts** list.

If an error occurs, you can download a host's log by clicking the **More Actions** icon () and selecting **Download Log**.

You can click the **Retry** button if the conversion host configuration failed for reasons unconnected with your environment.

(Optional) Click **Remove** to remove the configuration from a configured conversion host.

A.4. COPYING THE VMWARE SSH KEYS TO THE CONVERSION HOSTS

Copy the SSH public keys of the VMware hypervisors to the conversion hosts. You can collect the VMware keys either when you configure the VMware hypervisors for SSH transformation or by using **ssh-keyscan**:

- [Copying keys collected during VMware hypervisor configuration](#)
- [Copying keys collected with **ssh-keyscan**](#)

A.4.1. Copying keys collected during VMware hypervisor configuration

1. Copy the VMware keys to **/var/lib/vdsm/.ssh/known_hosts** on each conversion host.
2. Verify the SSH connection by connecting to each VMware hypervisor as **vdsm**:

```
# sudo -u vdsm ssh root@_esx1.example.com_
```

If the connection fails, check that the VMware hypervisor has SSH access enabled and that you copied the correct keys.

A.4.2. Copying keys collected with `ssh-keyscan`

CAUTION

You must run **ssh-keyscan** for each VMware hypervisor. Otherwise your conversion hosts will not have all the VMware hypervisor keys and the migration will fail.

1. Run **ssh-keyscan** for each VMware hypervisor and copy its public key to **known_hosts**, as in the following example:

```
# ssh-keyscan esx1_IP > /var/lib/vdsm/.ssh/known_hosts ❶  
# ssh-keyscan esx2_IP >> /var/lib/vdsm/.ssh/known_hosts  
# ssh-keyscan esx3_IP >> /var/lib/vdsm/.ssh/known_hosts
```

❶ You must use the IP address, not the host name, of the VMware hypervisor.

2. Change the ownership of the **known_hosts** file to **vdsm** user and **kvm** group:

```
# chown 36:36 /var/lib/vdsm/.ssh/known_hosts
```

3. Verify the SSH connection by connecting to each VMware hypervisor as **vdsm**:

```
# sudo -u vdsm ssh root@_esx1.example.com_
```

If the connection fails, check that the VMware hypervisor has SSH access enabled and that you copied the correct keys.

PART II. MIGRATING FROM VMWARE TO RED HAT OPENSTACK PLATFORM

The migration process involves the following tasks:

1. Planning the migration. See [Chapter 5, *Planning the migration*](#).
2. Preparing the VMware, Red Hat OpenStack Platform, and CloudForms environments. See [Chapter 6, *Preparing the environment for migration*](#).
3. Migrating the virtual machines. See [Chapter 7, *Migrating the virtual machines*](#).
4. Troubleshooting, if necessary. See [Chapter 8, *Troubleshooting*](#).



NOTE

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix B, *Configuring the Red Hat OpenStack Platform environment for SSH transformation*](#).

CHAPTER 5. PLANNING THE MIGRATION

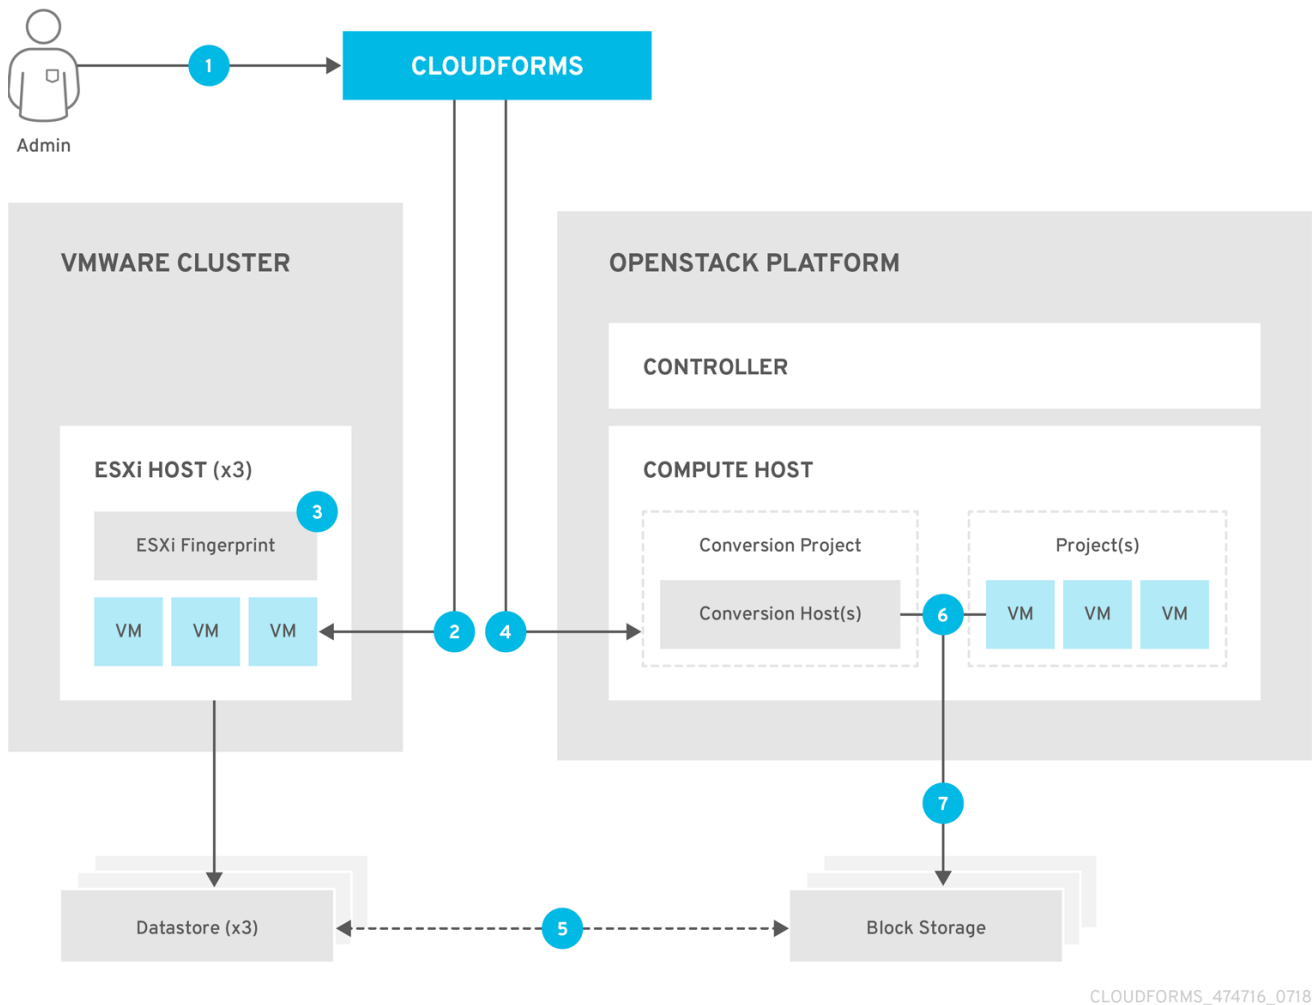
During the planning phase, you will formulate a specific migration goal, for example, "I want to migrate 2000 virtual machines, with 200 TB of data, in less than 6 months".

Review the following information to plan your migration:

- [Section 5.1, "Questions to ask before migration"](#) includes the following:
 - What operating systems are supported?
 - How long will the migration take?
 - How many conversion hosts do I need?
- [Section 5.2, "Recommendations and best practices"](#) provides guidelines on the following topics:
 - Scheduling the migration
 - Distributing the workload
 - Deploying the conversion hosts
 - Controlling the migration process

The migration workflow describes the migration process in greater detail.

Figure 5.1. VMware to Red Hat OpenStack Platform migration workflow



- 1 You create and run a migration plan in CloudForms.
- 2 CloudForms uses the migration plan to locate the source virtual machines.
- 3 CloudForms captures the ESXi host fingerprint for authentication during the virtual machine conversion process.
- 4 Using the attributes defined for the Red Hat OpenStack Platform environment, CloudForms initiates communication with the conversion hosts (Red Hat OpenStack Platform instances created from a conversion host appliance, with **virt-v2v** and **virt-v2v-wrapper** installed).
- 5 **virt-v2v-wrapper** connects to the source datastore through the ESXi host. **virt-v2v** streams the source disks to the target data domain and converts the source disks.
- 6 After the source disks are converted, **virt-v2v** detaches the volumes from the conversion host, migrates the volumes to the destination project, and creates the network ports defined in the infrastructure mapping.

7

virt-v2v-wrapper creates the target Red Hat OpenStack Platform instance with the flavor and security group defined in the migration plan. **virt-v2v** attaches the newly created network ports and the disks mapped in the block storage to the instance and the instance is powered on.

The migration process is complete and the migration plan's status is displayed in CloudForms.

5.1. QUESTIONS TO ASK BEFORE MIGRATION

The following questions can help you to estimate the resources and time required for migration.

What am I migrating?

- Identify the VMware virtual machines that you will be migrating.

What is the maximum number of disks or virtual machines that I can migrate?

- There is no maximum number of disks or virtual machines that you can migrate. However, you may not want to migrate all your virtual machines at the same time, in order to minimize the impact on your users.



IMPORTANT

If you exceed the capabilities of your environment, the migrations will fail. This situation could affect existing applications running on virtual machines attached to the network and storage.

What operating systems can I migrate?

- You can migrate any guest operating system that is certified and supported for use with the Red Hat Virtualization Hypervisor. See [Certified Guest Operating Systems in Red Hat OpenStack Platform and Red Hat Virtualization](#).

What am I missing?

- Identify resource gaps, such as bandwidth, storage, licenses, or a suitable maintenance window, before you begin the migration.

What impact will the migration have on my users?

- Assess the effects the migration may have on a production environment. It may be possible to migrate your applications in phases, without downtime at the application layer, if the applications are distributed in a high-availability architecture.
- Check whether users will lose access to critical applications.

How long will the migration take?

There is no formula to estimate how long the actual migration will take. This is determined on a case-by-case basis.

The following example is provided as a guide:

- Duration of migration: 2:13:00 (hh:mm:ss)
- 20 virtual machines

- 2 conversion hosts, maximum of 10 concurrent conversions
- Total data migrated: 1000 GB
- Hardware:
 - Strong host (40 cores, 500 GB RAM)
 - Fast SSD XtremIO storage
 - Fibre Channel 8 interface for host-to-storage connection
 - 10 GbE network interface cards for all other connections

How many conversion hosts do I need?

The number of conversion hosts you create depends on the size of your migration. All the virtual machines in a migration plan are migrated at the same time, in parallel. The number of virtual machines that you can migrate simultaneously depends on your infrastructure capabilities. Each migration requires a certain amount of network bandwidth, I/O throughput, and processing power for the conversion process.

Multiple conversion hosts provide load-balancing and better performance, even for small migrations.

Conversion hosts are limited to a maximum of ten concurrent migrations, unless you change the default values.

You should test your environment thoroughly before the migration to determine how many migrations it can support without negative effects, for example, five conversion hosts, each running ten concurrent migrations.

Should I migrate my virtual machines with VDDK or SSH?

You can migrate your virtual machines with either the VMware Virtual Disk Development Kit (VDDK) or SSH. VDDK is the default because it is much faster than SSH and easier to configure.

VDDK is limited to 20 concurrent migrations per conversion host, because of network limitations, and 10 concurrent migrations per VMware hypervisor, unless you increase the hypervisor's NFC service memory.


If you cannot use VDDK, SSH transformation is a fallback option.

5.2. RECOMMENDATIONS AND BEST PRACTICES

The following recommendations will help to minimize the impact of the migration on your environment.

Scheduling the migration

- Schedule your migration carefully, to minimize the impact on your users.
- Prepare your users for downtime.



IMPORTANT

Currently, IMS supports only cold migration. Virtual machines are powered off gracefully as part of the migration process.

- Stagger the migration schedules.
- Move critical applications during maintenance windows.

Distributing the migration workload

- Create migration groups, so that you are not migrating all of your virtual machines at the same time, keeping in mind the following considerations:
 - How are the virtual machines grouped now?
 - Which virtual machines should be migrated together?
 - Which workloads or linked applications should be migrated together?
 - What applications must remain available?
- Consider which parts of the workload to migrate first:
 - Databases
 - Applications
 - Web servers
 - Load balancers

Deploying the conversion hosts

- Create a sufficient number of conversion hosts for your migration, with sufficient resources.
- Create multiple conversion hosts for load-balancing. The virtual machines in a migration plan are automatically distributed among the conversion hosts. This decreases the load on the conversion hosts and allows you to increase the concurrent migrations beyond the limits of a single conversion host.

Controlling the migration process

- Create multiple migration plans for finer control.
- Perform test migrations with different maximum numbers of concurrent migrations to assess the capabilities of your environment's infrastructure.

CHAPTER 6. PREPARING THE ENVIRONMENT FOR MIGRATION

Preparing your environment for migration involves the following tasks:

1. Preparing the VMware environment. See [Section 6.1, “Preparing the VMware environment”](#).
2. Preparing the target environment. See [Section 6.2, “Preparing the Red Hat OpenStack Platform environment”](#).
3. Configuring the conversion hosts. See [Section 6.2.5, “Configuring the conversion hosts”](#).



NOTE

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix B, *Configuring the Red Hat OpenStack Platform environment for SSH transformation*](#).

6.1. PREPARING THE VMWARE ENVIRONMENT

Preparing the VMware environment for migration involves the following tasks:

1. Extending the VMware network. See [Section 6.1.1, “Preparing the VMware network”](#).
2. Preparing the VMware virtual machines. See [Section 6.1.2, “Preparing the VMware virtual machines”](#).
3. (Optional) Configuring a VMware hypervisor for more than ten concurrent migrations from that hypervisor. See [Section 6.1.3, “Configuring a VMware hypervisor for more than ten concurrent migrations”](#).

6.1.1. Preparing the VMware network

Extend the VMware network to the Red Hat OpenStack Platform environment.



IMPORTANT

- The network configuration must not be changed in any way during the migration.
- IP addresses, VLANs, and other network configuration must not be changed before or after migration because the conversion process preserves the source virtual machine MAC addresses.

6.1.2. Preparing the VMware virtual machines

Perform the following steps on each VMware virtual machine that you are migrating:

1. Install VMware Tools to capture IP addresses.
To download and install VMware Tools, see [VMware Workstation 5.0: Installing VMware Tools](#).
2. Unmount mounted ISO/CDROM disks.
3. Ensure that attached disks are not encrypted.

4. Ensure that each NIC has no more than one IPv4 and/or one IPv6 address.
5. Ensure that the virtual machine names contain only upper- or lower-case letters, numbers, underscores (`_`), hyphens (`-`), or periods (`.`).

**NOTE**

International characters and spaces are not permitted.

6. Ensure that the virtual machine names do not duplicate names of virtual machines in the Red Hat OpenStack Platform tenant.

If you are performing more than ten concurrent migrations from a VMware hypervisor, you must configure the hypervisor to support the additional connections. See [Section 6.1.3, “Configuring a VMware hypervisor for more than ten concurrent migrations”](#).

6.1.3. Configuring a VMware hypervisor for more than ten concurrent migrations

If you are performing more than ten concurrent migrations from a VMware hypervisor using VDDK transformation, the migration will fail because the hypervisor’s NFC service memory buffer is limited to ten parallel connections. See [VMware vSphere 6.5 NFC session connection limits](#) and [Virt-v2v. VDDK: ESXi NFC service memory limits](#) for details.

You can increase the hypervisor’s NFC service memory to enable additional connections for migrations.

Procedure

1. Log in to a VMware hypervisor.
2. Change the value of **maxMemory** to **1000000000** in **/etc/vmware/hostd/config.xml**:

```
<nfcsvc>
  <path>libnfcsvc.so</path>
  <enabled>>true</enabled>
  <maxMemory>1000000000</maxMemory>
  <maxStreamMemory>10485760</maxStreamMemory>
</nfcsvc>
```

3. Restart **hostd**:

```
# /etc/init.d/hostd restart
```

You do not need to reboot the VMware hypervisor.

6.2. PREPARING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT

Preparing the Red Hat OpenStack Platform environment involves the following key steps:

1. Installing and configuring Red Hat OpenStack Platform 13 or 14. See [Section 6.2.2, “Installing and configuring Red Hat OpenStack Platform 13 or 14”](#).
2. Installing and configuring CloudForms 4.7.6 or later. See [Section 6.2.3, “Installing and configuring CloudForms 4.7.6”](#).

3. Creating the conversion hosts. See [Section 6.2.4, “Creating the conversion hosts”](#).



NOTE

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix B, Configuring the Red Hat OpenStack Platform environment for SSH transformation](#).

6.2.1. Prerequisites

The following prerequisites apply:

- BIOS settings of physical hosts adjusted for optimal performance (rather than power-saving), according to the vendor’s recommendations
- C1E halt state disabled, if applicable
- Compatible software versions:

Table 6.1. Software compatibility

Software	IMS 1.1	IMS 1.2
VMware	5.5 or later	
CloudForms	4.7.0, with CFME 5.10.3	4.7.6 or later, with CFME 5.10.3
	CFME 5.10.4 does not support migration.	
Red Hat OpenStack Platform	13 or 14	
RHOSP V2V Image for Red Hat OpenStack Director	14.0.2	14.0.3

6.2.2. Installing and configuring Red Hat OpenStack Platform 13 or 14

1. Install Red Hat OpenStack Platform 13 or 14. See [Red Hat OpenStack Platform Director Installation and Usage 13](#) or [Red Hat OpenStack Platform Director Installation and Usage 14](#).
2. Create provider networks for the target instances to preserve the IP addresses of the source virtual machines. See [Create a network](#) in the *Red Hat OpenStack Platform Networking Guide*.
3. Create a project for the conversion hosts and whatever destination projects you require for the target instances. See [Create a Project](#) in the *Red Hat OpenStack Platform Users and Identity Management Guide*.
4. Ensure that the **admin** user has **member** and **admin** roles in the conversion and destination projects. See [Edit a Project](#) in the *Red Hat OpenStack Platform Users and Identity Management Guide*.
5. Set at least one volume type for the target block storage. See [Create a Volume](#) and [Changing a Volume’s Type \(Volume Re-typing\)](#) in the *Red Hat OpenStack Platform Storage Guide*. Otherwise, CloudForms cannot detect the storage when you create the infrastructure mapping.

6. Ensure that the storage backends have sufficient space for the migrated virtual machines.



IMPORTANT

If you are using Red Hat Ceph Storage, you will require three times the space of the source virtual machines for the migrated virtual machines. A Ceph storage cluster, by default, creates two copies of an object in a replicated storage pool, for a total of three copies. See [Data Copies](#) in the *Red Hat Ceph Storage Architecture Guide*.

The migrated disks use all of the space because it is preallocated. For example, a source virtual machine with a 100 GB disk requires 300 GB of storage, regardless of how much data the disk actually contains. To save storage space, you can use the **fstrim** command on the migrated virtual machines as a postmigration task or playbook.

7. Configure security groups with the following ports enabled:

- For the conversion hosts and CloudForms: port 22 (SSH)
- For CloudForms: port 443 (HTTPS)



NOTE

Outbound traffic is enabled by default. If you have changed this setting, enable ports 902 (CloudForms to VMware) and 5480 (conversion hosts to vCenter).

8. Create flavors for the source virtual machines. If you do not create custom flavors, CloudForms will try to map each source virtual machine to an existing flavor.
9. Ensure that the event monitor workers are running. See [Server Diagnostics](#) in *Red Hat CloudForms General Configuration*.

6.2.3. Installing and configuring CloudForms 4.7.6

CAUTION

Removing or changing a provider will cause errors in the infrastructure mappings and migration plans.

Procedure

1. Install Red Hat CloudForms 4.7.6 or later. See [Installing Red Hat CloudForms on Red Hat OpenStack Platform](#).




NOTE

CFME 5.10.4 does not support migration.


2. Enable SmartState Analysis:
 - a. In the header bar, click **Administrator** → **Configuration**.

- b. In the left pane, in **Settings**, click **CFME Regions** → **Zones** → **Server**.
- c. In the right pane, in the **Server** tab, set **SmartState Analysis** to **On** and click **Save**.
3. Add VMware to CloudForms as a provider. See [Adding a VMware vCenter Provider](#) in *Red Hat CloudForms: Managing Providers*.
4. Add Red Hat OpenStack Platform to CloudForms as a cloud provider. See [Adding an OpenStack Infrastructure Provider](#) in *Red Hat CloudForms: Managing Providers*.



IMPORTANT

Do not complete the fields in the **RSA key pair** tab. You will add the SSH private key when you configure the conversion hosts.



NOTE

If Red Hat OpenStack Platform has been active for a while, you must wait for CloudForms to update its event history before attempting to use the provider. Check the cloud provider timeline to verify that all events have been processed.

Viewing the cloud provider timeline in CloudForms

To view the cloud provider timeline:

1. Click **Compute** → **Clouds** → **Providers**.
2. Click the cloud provider to view its details.
3. Click **Monitoring** → **Timelines**.
4. Select the following options:
 - Events: **Management Events**
 - Category: **Select All**
 - Severity: **Select All**
5. Click **Apply**.

6.2.4. Creating the conversion hosts

You can create the Red Hat OpenStack Platform conversion hosts with the conversion appliance (**RHOSP V2V Image for Red Hat OpenStack Director**). The number of conversion hosts you deploy depends on your migration size and infrastructure capabilities.



NOTE

For optimal performance, deploy conversion hosts on compute nodes with nested virtualization enabled. See [Configure DevStack with KVM-based Nested Virtualization](#). Nested virtualization is a technology preview.

To create a Red Hat OpenStack Platform conversion host:

1. Navigate to [Red Hat Product Downloads](#).


2. In the **A-Z** tab, click **Red Hat OpenStack Platform**.
3. Click the green **Download Latest** button to go to the **Download Red Hat OpenStack Platform** page.
4. In the **Product Software** tab, locate the **RHOSP V2V Image for Red Hat OpenStack Director 14.0.x (x86_64)**, click **Download Now**, and save the image.

**IMPORTANT**

IMS 1.1 requires Red Hat OpenStack Director 14.0.2.

IMS 1.2 requires Red Hat OpenStack Director 14.0.3 or later.

5. Upload the image to Red Hat OpenStack Platform.
6. Launch the image as a conversion host instance, with the following resources:
 - 4 vCPUs
 - 10 GB RAM, if you use the default maximum number of concurrent migrations per conversion host, which is **10**. If you increase the number of concurrent migrations, you must add 1 GB RAM for each additional concurrent migration. If you reduce the number, you can reduce the RAM but the conversion host cannot have less than 8 GB RAM.
 - **/tmp** (10 GB, or 1 GB for each concurrent migration)
 - **/var/tmp** (10 GB, or 1 GB for each concurrent migration)
 - **/var/logs** (5 GB)
See [Launch an Instance](#) in the *Red Hat OpenStack Platform Instances and Images Guide* .
7. Increase the disk space of the instance to accommodate its file system. See [Resize an Instance](#) in the *Red Hat OpenStack Platform Instances and Images Guide* .
The instance is created from an image, but the disk space defined in the image will not be sufficient. You can either extend the partition (and subsequently, extend the physical volume in the volume group) to the required size, or you can create a new partition and add it as a physical volume to the volume group.

**NOTE**

You must resize **lv_root** to use all available disk space because the image will not use it by default.

Configure your conversion hosts with CloudForms to use the VMware Virtual Disk Development Kit (VDDK).

**NOTE**

The virtual disks are converted with the VMware Virtual Disk Development Kit (VDDK). If you cannot use VDDK, SSH transformation is a fallback option. See [Appendix B, Configuring the Red Hat OpenStack Platform environment for SSH transformation](#) .

6.2.5. Configuring the conversion hosts

Prerequisite

- Download and save the VMware Virtual Disk Development Kit.

Downloading the VMware Virtual Disk Development Kit

Procedure

1. Navigate to [VMware Documentation](#).
2. Click **VMware SDK & API Product Documentation** to expand.
3. Click **VMware Virtual Disk Development Kit (VDDK)**
4. Click **Latest Releases** and select the latest VDDK release.
5. Click **Download SDKs** to download the VDDK archive file.
6. Save the VDDK archive file in an HTTP-accessible location and record its path.


Procedure

1. In CloudForms, click **Compute** → **Migration** → **Migration Settings**.
2. In the **Conversion Hosts** tab, click **Configure Conversion Host**. The **Configure Conversion Host** wizard is displayed.

The screenshot shows the 'Configure Conversion Host' wizard interface. At the top, the title 'Configure Conversion Host' is displayed with a close button (X). Below the title is a progress bar with four steps: 'Location' (step 1, highlighted in blue), 'Host(s)' (step 2), 'Authentication' (step 3), and 'Results' (step 4). Below the progress bar, there is a field for '* Provider Type' with a dropdown menu currently showing '<Choose>'. At the bottom right, there are three buttons: 'Cancel', '< Back', and 'Next >'.

3. In the **Location** screen, add the provider details:
 - a. Select a **Provider Type**.

- b. Select a **Provider**.
 - c. Select a **Project** and click **Next**.
4. In the **Host(s)** screen, select one or more hosts from the dropdown list and click **Next**.
5. In the **Authentication** screen, add the conversion hosts' SSH key and the transformation method:
 - a. Click **Browse** to browse to the SSH private key or paste it in the **Conversion host SSH private key** field.
The Red Hat OpenStack Platform user uses a private SSH key to connect to the conversion hosts.
 - b. Select **VDDK** as the **Transformation method**.
 - c. Enter the path of the VDDK package in the **VDDK library path** field and click **Validate**. Wait for validation to complete.
 - d. Click **Configure**.
6. In the **Results** screen, wait for the conversion host configuration to finish and click **Close**.
The configured conversion hosts and status information, including error messages, appear in the **Configured Conversion Hosts** list.

If an error occurs, you can download a host's log by clicking the **More Actions** icon () and selecting **Download Log**.

You can click the **Retry** button if the conversion host configuration failed for reasons unconnected with your environment.

(Optional) Click **Remove** to remove the configuration from a configured conversion host.

You can migrate your virtual machines.

(Optional) You can verify the conversion hosts in a browser. See [Section 6.2.6, "Verifying the conversion hosts in a browser"](#).

6.2.6. Verifying the conversion hosts in a browser

You can verify your conversion hosts in a browser by using the CloudForms API:

1. In the address bar of a browser, enter the following:

```
https://CloudForms_FQDN/api/conversion_hosts 
```

-  **CloudForms_FQDN** is the FQDN of the CloudForms machine.

A log-in screen is displayed.

2. Enter your CloudForms **Username** and **Password** and click **Sign in**.
The conversion hosts and their IDs are displayed in JSON format:

```
{ "name": "conversion_hosts", "count": 3, "subcount": 3, "pages": 1, "resources":
  [ { "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000001" },
```

```
{ "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000002"},  
{ "href": "https://cloudforms.example.com/api/conversion_hosts/10000000000003"}], "actions":  
[{"name": "create", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts"},  
{ "name": "edit", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts"},  
{ "name": "delete", "method": "post", "href": "https://cloudforms.example.com/api/conversion_hosts"}], "links": { "self": "https://cloudforms.example.com/api/conversion_hosts?offset=0", "first": "https://cloudforms.example.com/api/conversion_hosts?offset=0", "last": "https://cloudforms.example.com/api/conversion_hosts?offset=0"}}
```

CHAPTER 7. MIGRATING THE VIRTUAL MACHINES

Migrating the virtual machines involves the following key tasks:

1. Mapping the resources of your VMware and Red Hat OpenStack Platform environments. See [Section 7.1, “Creating an infrastructure mapping”](#).
2. Checking for migration conditions with prerequisites. See [Section 7.2, “Checking for migration prerequisites”](#).
3. Creating and running a migration plan. See [Section 7.3, “Creating and running a migration plan”](#).

Optionally, you can change the maximum number of concurrent migrations for conversion hosts or providers to control the migration process.

7.1. CREATING AN INFRASTRUCTURE MAPPING

The infrastructure mapping maps the resources of your VMware and Red Hat OpenStack Platform environments.




IMPORTANT

If you add or remove providers or provider objects from an infrastructure mapping, the mapping will have missing resource errors. You must delete the infrastructure mapping and create a new one.

Procedure

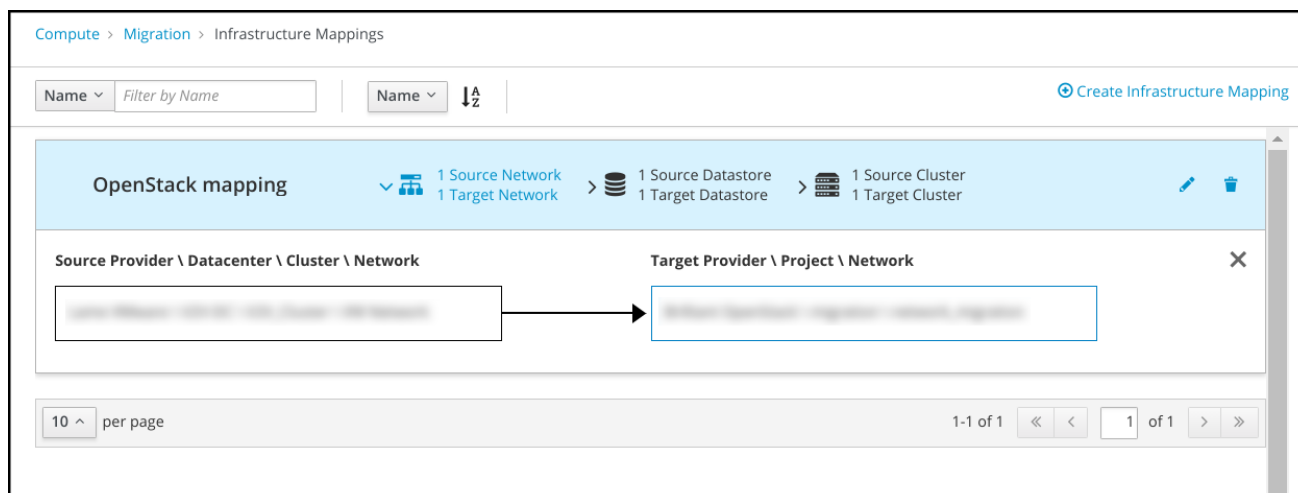
1. Click **Compute** → **Migration** → **Infrastructure Mappings**.
2. Click **Create Infrastructure Mapping** The **Create Infrastructure Mapping** wizard is displayed.

3. In the **General** screen, add the infrastructure mapping details:
 - a. Enter the infrastructure mapping **Name** and (optional) **Description**.
 - b. Select the **Target Provider**.
 - c. Click **Next**.
4. In the **Map Compute** screen, map the source and target clusters:
 - a. Select a **Source Provider \ Datacenter \ Clustersource** cluster and a **Target Provider \ Project**.
 If the target project does not contain a conversion host, a warning icon () appears. You can create and save an infrastructure mapping, but you must configure the conversion hosts before running a migration plan.
 - b. Click **Add Mapping**. You can map additional projects.
 - c. Click **Next**.
5. In the **Map Storage** screen, map the source and target storage:
 - a. Select a **Source Provider \ Datacenter \ Datastore** and **Target Provider \ Volume Type**.
 If the volume type is missing, check that the volume type has been set. Block storage requires at least one volume type. See [Create a Volume](#) and [Changing a Volume's Type \(Volume Re-typing\)](#) in the *Red Hat OpenStack Platform Storage Guide*.

- b. Click **Add Mapping**. You can map additional datastores.
 - c. Click **Next**.
6. In the **Map Networks** screen, map the source and target networks:
 - a. Select a source cluster from the drop-down list.
 - b. Select one or more networks from **Source Provider \ Datacenter \ Network** and **Target Project \ Network**.
IMS supports both provider and tenant networks.
 - c. Click **Add Mapping**. You can map the networks of additional source clusters.
 - d. Click **Create**.
7. In the **Results** screen, click **Close**. The infrastructure mapping is saved in **Compute → Migration → Infrastructure Mappings**.

You can click an infrastructure mapping element to view its details:

Infrastructure Mappings list



After you have created an infrastructure mapping, check for migration prerequisites.

If these conditions do not apply, you can create a migration plan. See [Section 7.3, “Creating and running a migration plan”](#).

7.2. CHECKING FOR MIGRATION PREREQUISITES

Check your migration for the following conditions, which have prerequisites:

You are migrating previously migrated virtual machines

Add previously migrated machines to the migration plan with a CSV file. A CSV file is also recommended for large migrations.

See [Section 7.2.1, “Creating a CSV file to add virtual machines to the migration plan”](#).

You are using Ansible playbooks for premigration/postmigration tasks

Create an Ansible repository and add credentials and playbooks to CloudForms.

See [Section 7.2.2, “Adding Ansible playbooks to CloudForms for premigration and postmigration tasks”](#).

You are migrating virtual machines running RHEL or other Linux operating system

Create a RHEL premigration playbook to preserve IP addresses and select this playbook when you create a migration plan.

See [Section 7.2.2, “Adding Ansible playbooks to CloudForms for premigration and postmigration tasks”](#) and [Section 7.2.3, “Creating a RHEL premigration playbook for RHEL/Linux source virtual machines”](#).

7.2.1. Creating a CSV file to add virtual machines to the migration plan

If you are migrating virtual machines that were migrated in the past, you should create a CSV file to add the virtual machines to the migration plan, because the migration plan cannot discover them automatically.



NOTE

A CSV file is recommended for large migrations because it is faster than manually selecting the security group and flavor of each virtual machine.

Table 7.1. CSV file fields

Field	Comments
Name	Virtual machine name. Required
Host	Optional. Only required if virtual machines have identical Name fields.
Provider	Optional. Only required if virtual machines have identical Name and Host fields.
Security Group	Optional. The default is Default .
Flavor	Optional If you do not create flavors for the migration or if you leave this field blank, CloudForms tries to map the source virtual machines to existing flavors.

CSV file example

```
Name,Host,Provider,Security Group,Flavor
vm01,host1,vSphere3,webservers,x1.medium
vm02,host1,vSphere3,webservers,x1.medium
vm03,host1,vSphere3,webservers,x1.medium
```

7.2.2. Adding Ansible playbooks to CloudForms for premigration and postmigration tasks

You can add Ansible playbooks to CloudForms to perform automated premigration and postmigration tasks on specific virtual machines, for example:

- Removing webservers from a load-balancing pool before migration and returning them to the pool after migration

- Running **fsttrim** after migration to reduce the space required by virtual machines migrating to Red Hat OpenStack Platform with Ceph storage

Procedure

1. Enable the **Embedded Ansible** server role in CloudForms. See [Enabling the Embedded Ansible Server Role](#) in *Red Hat CloudForms: Managing Providers* .
2. Add an Ansible playbook repository. See [Adding a Playbook Repository](#) in *Red Hat CloudForms: Managing Providers*.
3. Add the credentials of each virtual machine that you are migrating. See [Credentials](#) in *Red Hat CloudForms: Managing Providers*.
4. Add your playbook as an Ansible service catalog item. See [Creating an Ansible Playbook Service Catalog Item](#) in *Red Hat CloudForms: Provisioning Virtual Machines and Instances* .

You will select the playbooks and the virtual machines on which they run in the **Advanced Options** screen when you create the migration plan.

7.2.3. Creating a RHEL premigration playbook for RHEL/Linux source virtual machines

If you are migrating virtual machines running RHEL or other Linux operating system, you can create a **RHEL premigration** playbook to ensure that the IP addresses are accessible after migration. The **RHEL premigration** playbook calls the Ansible **ims.rhel_premigration** role.

To install the role with Ansible Galaxy, see [ims_rhel_pre_migration](#). This role is not included in the IMS installation.

The **ims.rhel_premigration** role performs the following tasks on the VMware virtual machines:

- Preserves the static IP address configuration by creating udev rules to associate the virtual machine's MAC address with its interface name



NOTE

The **ims.rhel_premigration** role assumes that either the **rhel-6-server-rpms** or the **rhel-7-server-rpms** repository is enabled in the source virtual machine when it installs **qemu-guest-agent**. If you have disabled the repository, re-enable it in the RHEL premigration playbook.

RHEL premigration playbook example

```
---
- hosts: all
  roles:
    - role: ims.rhel_pre_migration
```


You can create a migration plan. See [Section 7.3, "Creating and running a migration plan"](#) .

7.3. CREATING AND RUNNING A MIGRATION PLAN

Before you perform a large migration, you should perform several test migrations with different maximum numbers of concurrent migrations for your conversion hosts or providers. This will enable you to assess the capabilities of your environment's infrastructure.

You can create and run a migration plan in CloudForms with the following options:

- Running a migration plan immediately
- [Scheduling a migration plan to run in the future](#)
- [Viewing a migration plan in progress](#)
- [Canceling a migration plan in progress](#)
- [Retrying a failed migration plan](#)



NOTE

A CSV file is optional, but recommended, for large migrations because it is faster than manually selecting the security group and flavor of each virtual machine.

Procedure

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click **Create Migration Plan**. The **Create Migration Plan** wizard is displayed.

Create Migration Plan [X]

General (1) VMs (2) Advanced Options (3) Schedule (4) Results (5)

Infrastructure Mapping *

Name *
Maximum length is 24 characters.

Description
Maximum length is 128 characters.

Select VMs

- Choose from a list of VMs discovered in the selected infrastructure mapping
- Import a CSV file with a list of VMs to be migrated

Cancel < Back Next >

3. In the **General** screen, add the details of the migration plan:
 - a. Select an infrastructure mapping from the drop-down list.
 - b. Enter the migration plan **Name** and (optional) **Description**.
 - c. Select a virtual machine discovery method:
 - **Choose from a list of VMs discovered in the selected infrastructure mapping**
If the virtual machines cannot be discovered, check that the source datastores and networks in the infrastructure mapping are correct.
 - **Import a CSV file with a list of VMs to be migrated**
A CSV file is required for previously migrated source virtual machines and recommended for large migrations.
 - d. Click **Next**.

4. In the **VMs** screen, select the virtual machines for migration:
 - If you selected **Choose from a list of VMs discovered in the selected infrastructure mapping**, select the virtual machines for migration.
You can search for virtual machines by **VM Name**, **Data Center**, **Cluster**, and **Folder**.
 - If you selected **Import a CSV file with a list of VMs to be migrated**
 - a. Click **Import**.
 - b. Browse to the CSV file and click **Open**.
If the virtual machines cannot be added to the migration plan, check the CSV file format and fields for errors.



NOTE

If the **Create Migration Plan** wizard freezes, refresh the web page, check the CSV file for errors (for example, virtual machines with duplicate **Name** fields and no other fields to distinguish them), and create a new migration plan.

- c. Click **Next**.
5. In the **Instance Properties** screen, select the networks and/or flavors:
 - a. Click the pencil icon to edit the network or flavor of each selected virtual machine.
Flavors that are too small for the virtual machine are marked with an asterisk (*). If you have not created flavors for the migration, CloudForms tries to map the source virtual machines to existing flavors.
 - b. Click **Next**.
6. In the **Advanced Options** screen, select the playbook service options:
 - a. Select a premigration and/or postmigration playbook service from the dropdown lists.
 - b. Select the virtual machines on which to run the playbook services.
 - c. Click **Next**.

7. In the **Schedule** screen, select a schedule option and click **Create**:

- **Save migration plan to run later**

The migration plan is saved in **Migration Plans Not Started** and will not run unless you schedule it or click **Migrate** to run the scheduled migration plan immediately.

- **Start migration immediately**

The migration plan may take some time to complete. Progress bars indicate the amount of transferred data, the number of migrated virtual machines, and the elapsed time. See [Section 7.3.2, “Viewing a migration plan in progress”](#) for details.

8. In the **Results** screen, click **Close**.

When the migration plan has finished, click **Migration Plans Complete** to view the status of the migration plan. The completed migration plan shows the status of the migrated virtual machines.

In the migration plans list, you can click the **More Actions** icon () to archive, edit, or delete a migration plan.

7.3.1. Scheduling a saved migration plan


To schedule a saved migration plan to run in the future:

1. Click **Migration Plans Not Started**
2. Click the **Schedule** button of a migration plan.
3. In the **Schedule Migration Plan** window, select a date and time and click **Schedule**. The plan's status is **Migration Scheduled** with the date and time.

7.3.2. Viewing a migration plan in progress

To view the progress of a migration plan:

1. Click **Migration Plans in Progress**
2. Click a migration plan name to view its details, including the status of the migrating virtual machines.



NOTE

The counter in **Compute** → **Migration** → **Migration Plans** may be a few seconds ahead of the counter in the migration plan details view. This is because the **Migration Plans** counter displays the total time for running the migration plan, while the details counter displays the time for migrating the virtual machines.

7.3.3. Canceling a migration plan in progress

To cancel a migration plan in progress:

1. Click **Migration Plans in Progress**
2. Select a migration plan and click **Cancel Migration**.

3. Click **Cancel Migrations** to confirm the cancellation. The canceled migration appears in **Migration Plans Complete** with a red **x** indicating that the plan did not complete successfully.

7.3.4. Retrying a failed migration plan

To retry a migration plan that failed because of external circumstances (for example, power outage):

1. Delete all objects created by the failed migration plan:
 - Delete newly created instances to avoid name conflicts with migrating VMware virtual machines.
 - Delete network ports of failed instances.
2. Click **Compute** → **Migration** → **Migration Plans**.
3. Click **Migration Plans Complete**.
4. Click the **Retry** button beside the failed migration plan.

7.4. CHANGING THE MAXIMUM NUMBER OF CONCURRENT MIGRATIONS

You can change the maximum number of concurrent migrations for conversion hosts or providers to control the impact of the migration process on your infrastructure.

The provider setting has priority over the conversion host setting. For example, if the maximum number of concurrent migrations is **20** for a provider and **3** for five conversion hosts, the maximum number of concurrent migrations is **20**, not **15** (**5** conversion hosts **x** **3** concurrent migrations per host).

An increase in the maximum number of concurrent migrations affects all migration plans immediately. Virtual machines that are queued to migrate will migrate in greater numbers.

A decrease maximum number of concurrent migrations affects only future migration plans. Migration plans that are in progress will use the limit that was set when the plan was created.


Changing the maximum number of concurrent migrations for all conversion hosts or providers

CAUTION

Red Hat OpenStack Platform conversion hosts require an additional 1 GB RAM for each additional concurrent migration above **10**.

If you are using VDDK transformation, the number of concurrent migrations must not exceed **20**. Otherwise, network overload will cause the migration to fail.

1. Click **Compute** → **Migration** → **Migration Settings**.
2. In the **Migration Throttling** tab, select a value for **Maximum concurrent migrations per conversion host** or **Maximum concurrent migrations per provider** and click **Apply**.
The value of **Maximum concurrent migrations per conversion host** is constrained so that it cannot be greater than **Maximum concurrent migrations per provider**.

**NOTE**

In the current release, the **Maximum concurrent migrations per conversion host** interface control does not work.

Changing the maximum number of concurrent migrations for a specific conversion host

1. Enter the following command:

```
# curl -sk -u username:password  
https://CloudForms_FQDN/api/conversion_hosts/conversion_host_id -X POST -d {"action": "edit",  
"resource": {"max_concurrent_tasks": 15}} 1 2 3 4
```

- 1 **username** and **password** are the username and password for CloudForms.
- 2 **CloudForms_FQDN** is the FQDN of the CloudForms machine.
- 3 To obtain a **conversion_host_id**, enter the following command:

```
# curl -sk -u username:password https://CloudForms_FQDN/api/conversion_hosts/
```

- 4 **max_concurrent_tasks** is the maximum number of concurrent migrations. The default is **10**.

CHAPTER 8. TROUBLESHOOTING

To identify errors, you can review the migration logs. See [Section 8.1, “Migration logs”](#).

You can check these common issues and mistakes:

- Infrastructure mapping errors. See [Section 8.2, “Infrastructure mapping errors”](#).
- Migration plan errors. See [Section 8.3, “Migration plan errors”](#).
- IP address errors. See [Section 8.4, “IP address errors”](#).
- Environment configuration errors. See [Section 8.5, “Environment configuration errors”](#).

[Section 8.6, “Known Issues”](#) provides information about issues that will be addressed in a future release.

8.1. MIGRATION LOGS

You can check the following logs to identify the cause of a migration error:

- Conversion host logs. See [Section 8.1.1, “Downloading the conversion host and playbook logs”](#).
- Playbook logs, if you are using premigration or postmigration Ansible playbooks. See [Section 8.1.1, “Downloading the conversion host and playbook logs”](#).
- CloudForms migration log. See [Section 8.1.2, “Accessing the CloudForms migration log”](#).

8.1.1. Downloading the conversion host and playbook logs

You can download the conversion host and playbooks logs in CloudForms.

When disk migration starts, two logs are created in the conversion host:

- **virt-v2v**: Debug output from **virt-v2v** itself. This log tracks the core of the virtual machine migration process, including **libguestfs** traces and disk migration details.
- **virt-v2v-wrapper**: Log of the daemonizing wrapper for **virt-v2v**. This log traces the orchestration of the virtual machine conversion on the conversion host, including disk migration percentages and **virt-v2v** error reporting.



IMPORTANT

If you open a [Red Hat Support call](#), you will need to submit both the **virt-v2v** log and **virt-v2v-wrapper** log for analysis and troubleshooting.

To access the conversion host logs and the playbook logs:

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click a completed migration plan to view its details.
3. Click **Download Log** of a virtual machine and select a log from the dropdown list:
 - **Premigration log** This option only appears if a premigration playbook is used.

- **Migration log** The **Migration log** is the **virt-v2v** log.
- **Virt-v2v-wrapper log**
- **Postmigration log** This option only appears if a postmigration playbook is used.

8.1.2. Accessing the CloudForms migration log

1. Log into the CloudForms machine using SSH.
2. The migration log is **/var/www/miq/vmdb/log/automation.log**.

8.2. INFRASTRUCTURE MAPPING ERRORS

- **Networks missing, Datastores missing, and Clusters missing** error messages: If you create an infrastructure mapping and then change a provider, the provider's object IDs change. Delete the infrastructure mapping and create a new one.
- **Storage volume type not detected**: Check that you have set at least one volume type. See [Group Volume Settings with Volume Types](#) in the *Red Hat OpenStack Platform Storage Guide* for the storage.

8.3. MIGRATION PLAN ERRORS

- If the virtual machines are being migrated for the first time and are not discovered by the migration plan, check the source datastores and networks in the infrastructure mapping.
- If the virtual machines have been migrated in the past, they cannot be discovered by the migration plan. Use a CSV file to add the virtual machines to the migration plan.
- If the virtual machines cannot be added to the migration plan with a CSV file, check the CSV file format and fields. Create a new migration plan with the updated CSV file.
- **Create Migration Plan** wizard hangs while importing a CSV file. This error is caused by an invalid CVS file (for example, virtual machines with a duplicate **Name** field and no **Host/Provider** field to distinguish them, or with a duplicate **Name** field and duplicate **Host/Provider** fields). Refresh the web page.
- **Denied State** error (IMS 1.1). If a migration plan fails immediately and the migration plan displays a **Denied State** error message, check that you have created and configured the conversion hosts correctly. Cancel the migration plan and run it again.
- **Unable to migrate VMs because no conversion host was configured at the time of the attempted migration. See the product documentation for information on configuring conversion hosts.** (IMS 1.2) You can create and save a migration plan whose infrastructure mapping does not contain conversion hosts, but you cannot run the migration plan without conversion hosts. Cancel the migration plan, create the conversion hosts, and run the migration plan again.

8.4. IP ADDRESS ERRORS

- If the IP address of a migrated RHEL (or other Linux-based operating system) virtual machine is not accessible, you must create a **RHEL premigration** playbook and add it to the migration plan.
- If a migrated virtual machine does not have an IP address:

- Check that you installed VMware Tools on the VMware virtual machine before migration.
- Check the VMware virtual machine for an interface configuration file mapped to a non-existent interface (for example, `/etc/sysconfig/network-scripts/ifcfg-eth1` exists, but `eth1` interface does not). Log example:

```
CalledProcessError: Command '['openstack', u'--os-username=admin', u'--os-identity-api-version=3', u'--os-user-domain-name=default', u'--os-auth-url=http://osp.example.com:5000/v3', u'--os-project-name=admin', u'--os-password=*****', u'--os-project-id=0123456789abcdef0123456789abcdef', 'port', 'create', '--format', 'json', '--network', u'01234567-89ab-cdef-0123-456789abcdef', '--mac-address', u'00:50:56:01:23:45', '--enable', u'port_0', '--fixed-ip', 'ip-address=None']' returned non-zero exit status 1
date time:ERROR: Command output:
BadRequestException: Unknown errors
```

8.5. ENVIRONMENT CONFIGURATION ERRORS

VMware

- A VMware virtual machine cannot be migrated if it has any of the following conditions:
 - Mounted ISO/CDROM disk
 - Encrypted disk
 - Invalid name, containing spaces or special characters
 - Powered off during migration
- If you are performing more than ten concurrent migrations from a single VMware hypervisor, you must increase the hypervisor's NFC service memory.

Red Hat OpenStack Platform

- **disallowed by policy** error: The Red Hat OpenStack Platform **admin** user in CloudForms does not have **admin** role privileges within the target project. Add the **admin** user as **member** and **admin** to your target project. See [Edit a Project](#) in the *Red Hat OpenStack Platform Users and Identity Management Guide*.

```
ERROR: Command exited with non-zero return code 1, output: HttpException: 403: Client Error for url: https://FQDN:13696/v2.0/ports, {"NeutronError": {"message": "(rule:create_port and rule:create_port:mac_address) and rule:create_port:fixed_ips) is disallowed by policy", "type": "PolicyNotAuthorized", "detail": ""}}
```

8.6. KNOWN ISSUES

The following issues will be addressed in a future release:

- BZ#1666799: [Canceling migration does not stop creating volume, instance, and network port on OpenStack Platform or VMs on Red Hat Virtualization](#)
- BZ#1678385: [Virtual machine with name containing spaces \(rhel 7\) fails to migrate using SSH and VDDK transformation](#)

- [BZ#1699343](#): [Migration plan CSV import validation does not work if file contains empty/archived/orphan/invalid VM name](#)
- [BZ#1698761](#): ["Maximum concurrent migrations per conversion host" interface control does not work](#)
- [BZ#716283](#): [Migrating virtual machines are not distributed correctly among the conversion hosts](#)
- [BZ#1726939](#): [Run the preflight check of migration task before waiting for a conversion host](#) .
Currently, the preflight check that monitors the migration is performed after a conversion host is assigned to the task. As a result, the total volume of the Datastores reported in **Migration Plans In Progress** reflects the total volume of the virtual machines that are currently migrating, not the total volume of the migration plan. When all the virtual machines have started to migrate, the correct value of the total volume is displayed.
- [BZ#1668049](#): [Instance is not created after disk conversion](#)
- [BZ#1669133](#): [Names of virtual machines migrated using SSH transformation are changed](#)
- CloudForms CFME 5.10.4 does not support migration. Use CFME 5.10.3.

APPENDIX B. CONFIGURING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT FOR SSH TRANSFORMATION

You can configure your environment for SSH transformation if you cannot use VDDK.

Configuring your environment involves the following additional steps:

1. Configuring the VMware hypervisors. See [Section B.1, "Configuring the VMware hypervisors for SSH transformation"](#).
2. Configuring the conversion hosts. See [Section B.2, "Configuring the conversion hosts"](#).
3. Copying the VMware keys to the conversion hosts. See [Section B.3, "Copying the VMware SSH keys to the conversion hosts"](#).

B.1. CONFIGURING THE VMWARE HYPERVISORS FOR SSH TRANSFORMATION

For SSH transformation, you must configure the VMware hypervisors for passwordless access by sharing a conversion host's public SSH key with the hypervisors.



IMPORTANT

A single SSH key pair is recommended because the key pair is used only for virtual machine conversion and it simplifies conversion host management.

If you wish to use a dedicated SSH key pair for each conversion host, you can copy the public key of each conversion host to all the VMware hypervisors.

Procedure

1. Enable SSH access on each VMware hypervisor. For instructions, navigate to [VMware vSphere Documentation](#) and enter **Enable ESXi Shell and SSH Access with the Direct Console User Interface** in the **Search** field.

You can collect the SSH public keys of the VMware hypervisors at this stage, to copy to the conversion hosts.

2. Generate an SSH key pair without a passphrase:

```
# ssh-keygen -N "
```

3. Copy the public SSH key to **/etc/ssh/keys-root/authorized_keys** on **each** VMware hypervisor. You will use the private SSH key to configure the conversion hosts.

B.2. CONFIGURING THE CONVERSION HOSTS


Procedure

1. In CloudForms, click **Compute** → **Migration** → **Migration Settings**.
2. In the **Conversion Hosts** tab, click **Configure Conversion Host**. The **Configure Conversion Host** wizard is displayed.

The screenshot shows a window titled "Configure Conversion Host" with a close button (X) in the top right corner. A progress bar at the top indicates four steps: "Location" (1), "Host(s)" (2), "Authentication" (3), and "Results" (4). Step 1 is highlighted with a blue circle. Below the progress bar, there is a field labeled "* Provider Type" with a dropdown menu showing "<Choose>". At the bottom right of the window, there are three buttons: "Cancel", "< Back", and "Next >".

3. In the **Location** screen, add the provider details:
 - a. Select a **Provider Type**.
 - b. Select a **Provider**.
 - c. Select a **Project**.
 - d. Click **Next**.
4. In the **Host(s)** screen, select one or more hosts from the dropdown list and click **Next**.
5. In the **Authentication** screen, add the conversion hosts' SSH key and the transformation method:
 - a. Click **Browse** to browse to the SSH private key or paste it in the **Conversion host SSH private key** field.
The Red Hat OpenStack Platform user uses a private SSH key to connect to the conversion hosts.
 - b. Select **SSH** as the **Transformation method**.
 - c. Click **Browse** to browse to the SSH private key you created for enabling SSH access on the VMware hypervisors or paste it in the **VMware hypervisors SSH private key** field.
 - d. Click **Configure**.
6. In the **Results** screen, wait for the conversion host configuration to finish and click **Close**.

The configured conversion hosts and status information, including error messages, appear in the **Configured Conversion Hosts** list.

If an error occurs, you can download a host's log by clicking the **More Actions** icon () and selecting **Download Log**.

You can click the **Retry** button if the conversion host configuration failed for reasons unconnected with your environment.

(Optional) Click **Remove** to remove the configuration from a configured conversion host.

B.3. COPYING THE VMWARE SSH KEYS TO THE CONVERSION HOSTS

Copy the SSH public keys of the VMware hypervisors to the conversion hosts. You can collect the VMware keys either when you configure the VMware hypervisors for SSH transformation or by using **ssh-keyscan**:

- [Copying keys collected during VMware hypervisor configuration](#)
- [Copying keys collected with **ssh-keyscan**](#)

B.3.1. Copying keys collected during VMware hypervisor configuration

1. Copy the VMware keys to **/root/.ssh/known_hosts** on each conversion host.
2. On each conversion host, verify the SSH connection by connecting to each VMware hypervisor as **cloud-user**:

```
# sudo -u cloud-user ssh root@esx1.example.com
```

If the connection fails, check that the VMware hypervisor has SSH access enabled and that you copied the correct keys.

B.3.2. Copying keys collected with **ssh-keyscan**

CAUTION

You must run **ssh-keyscan** for each VMware hypervisor. Otherwise your conversion hosts will not have all the VMware hypervisor keys and the migration will fail.

1. Run **ssh-keyscan** for each VMware hypervisor and copy its public key to **known_hosts**, as in the following example:

```
# ssh-keyscan esx1_IP > /root/.ssh/known_hosts 1
# ssh-keyscan esx2_IP >> /root/.ssh/known_hosts
# ssh-keyscan esx3_IP >> /root/.ssh/known_hosts
```

1 You must use the IP address, not the host name, of the VMware hypervisor.

2. On each conversion host, verify the SSH connection by connecting to each VMware hypervisor as **cloud-user**:

```
# sudo -u cloud-user ssh root@esx1.example.com
```

If the connection fails, check that the VMware hypervisor has SSH access enabled and that you copied the correct keys.