



# **Red Hat Infrastructure Migration Solution 1.1**

## **Infrastructure Migration Solution Guide**

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



# Red Hat Infrastructure Migration Solution 1.1 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, 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 Software Collections 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 plan, perform, and troubleshoot a migration of virtual machines from VMware to Red Hat Virtualization or Red Hat OpenStack Platform. The migration is performed with Red Hat CloudForms.

## Table of Contents

<b>CHAPTER 1. INFRASTRUCTURE MIGRATION SOLUTION OVERVIEW</b> .....	<b>3</b>
1.1. VMWARE TO RED HAT VIRTUALIZATION WORKFLOW	3
1.2. VMWARE TO RED HAT OPENSTACK PLATFORM WORKFLOW	4
<b>CHAPTER 2. PLANNING THE MIGRATION</b> .....	<b>7</b>
2.1. PLANNING QUESTIONS	7
2.2. BEST PRACTICES FOR MIGRATION	8
<b>CHAPTER 3. PREPARING THE MIGRATION ENVIRONMENT</b> .....	<b>10</b>
3.1. SOFTWARE COMPATIBILITY MATRIX	10
3.2. PREPARING THE VMWARE ENVIRONMENT	10
3.3. PREPARING THE RED HAT VIRTUALIZATION ENVIRONMENT	11
3.4. PREPARING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT	11
3.5. CONVERSION HOSTS	13
3.5.1. Conversion host considerations	13
3.5.1.1. Number of conversion hosts	13
3.5.1.2. Transformation options: VDDK or SSH	13
3.5.2. Creating conversion hosts	14
3.5.3. Configuring conversion hosts for transformation	15
3.5.3.1. VDDK only: Downloading and copying the VDDK package	15
3.5.3.2. VDDK and SSH: Configuring the conversion hosts	15
3.5.3.3. VDDK and SSH: Validating the transformation configuration	18
3.5.3.4. SSH only: Configuring and validating the SSH connection	18
3.5.4. Enabling conversion hosts in CloudForms	19
3.5.5. Verifying conversion hosts	20
3.5.6. Changing the maximum number of concurrent migrations	20
3.5.7. Upgrading conversion hosts	21
<b>CHAPTER 4. MIGRATING THE INFRASTRUCTURE</b> .....	<b>22</b>
4.1. CREATING AN INFRASTRUCTURE MAPPING	22
4.2. CREATING AND RUNNING A MIGRATION PLAN	23
4.2.1. Pre-migration tasks	24
4.2.1.1. Creating a CSV file to add virtual machines to the migration plan	24
4.2.1.2. Setting up an Ansible environment	24
4.2.1.3. Creating a RHEL pre-migration playbook for migrating RHEL virtual machines	25
4.2.2. Creating a migration plan	25
4.2.3. Running a migration plan	27
4.2.3.1. Running a migration plan immediately	27
4.2.3.2. Running a migration plan in the future	27
4.2.4. Validating a migration plan	27
4.2.5. Canceling a migration plan in progress	28
<b>CHAPTER 5. TROUBLESHOOTING</b> .....	<b>29</b>
5.1. RETRYING AND RE-MIGRATING	29
5.2. LOGS	29
5.3. COMMON ISSUES AND MISTAKES	30
5.3.1. Infrastructure mapping errors	30
5.3.2. Migration plan errors	30
5.3.3. Virtual machine migration failure	31
5.4. KNOWN ISSUES	32



# CHAPTER 1. INFRASTRUCTURE MIGRATION SOLUTION OVERVIEW

Red Hat's Infrastructure Migration Solution (IMS) enables you to migrate your virtual machines from VMware to Red Hat Virtualization or to Red Hat OpenStack Platform.

[Red Hat Virtualization](#) (RHV) is an enterprise-grade virtualization platform built on Red Hat Enterprise Linux. With Red Hat Virtualization, you can manage your entire virtual infrastructure, including hosts, virtual machines, networks, storage, and users, from a centralized graphical user interface or with a REST API.

[Red Hat OpenStack Platform](#) provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

[Red Hat CloudForms](#) provides the user interface required to manage challenging enterprise environments.



## NOTE

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

## Migration phases

The Infrastructure Migration Solution has 4 key phases:

### 1. [Planning](#):

- Assessing the source and target environments to ensure sufficient resources for migration
- Formulating strategies to minimize the impact on users

### 2. [Preparation](#):

- Configuring the source and target environments
- Creating and configuring the conversion hosts

### 3. [Migration](#):

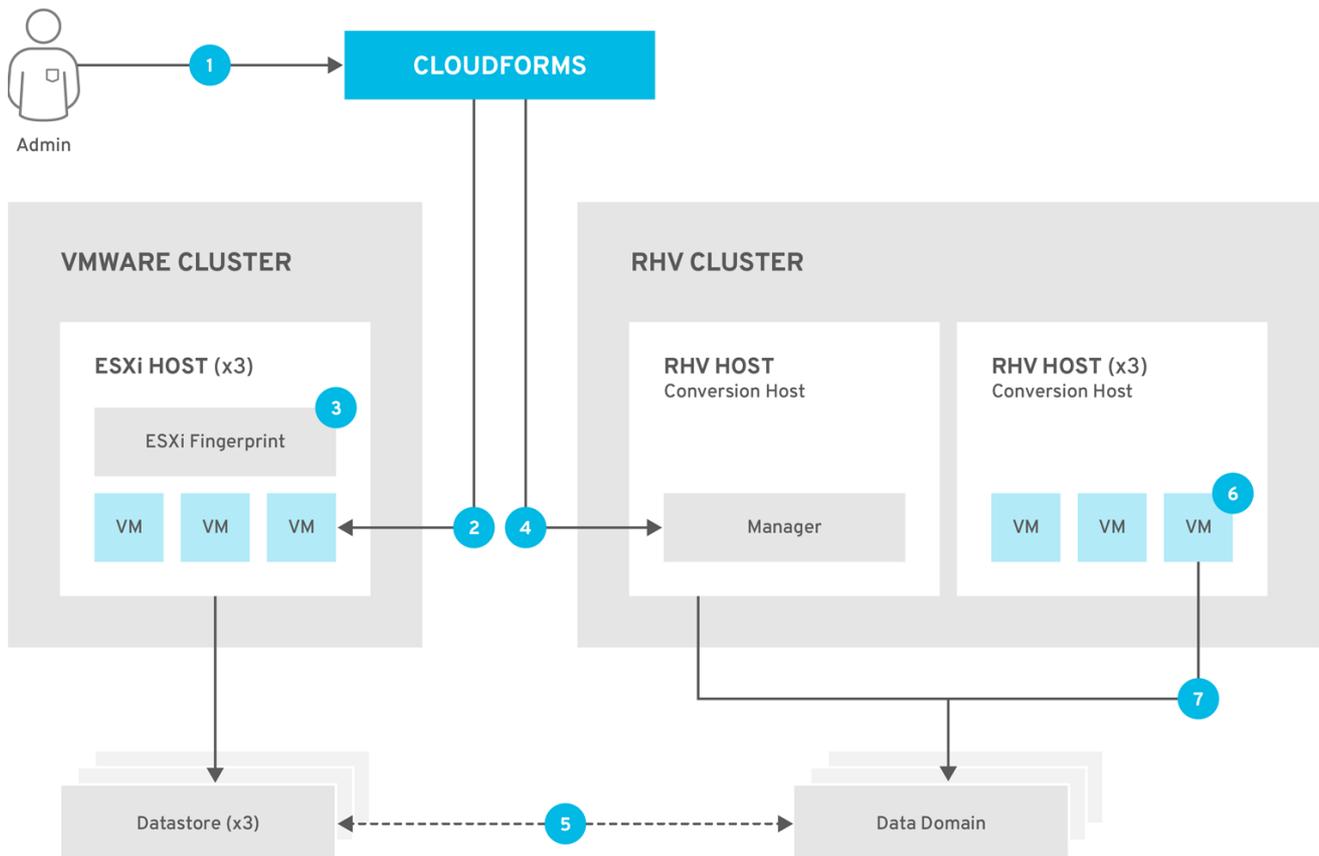
- Mapping the resources of the source environment to the target environment
- Creating and running a migration plan

### 4. [Validation](#):

- Checking that the migration plan completed successfully
- [Troubleshooting](#), if necessary

## 1.1. VMWARE TO RED HAT VIRTUALIZATION WORKFLOW

This workflow describes the process of migrating from VMware to Red Hat Virtualization.

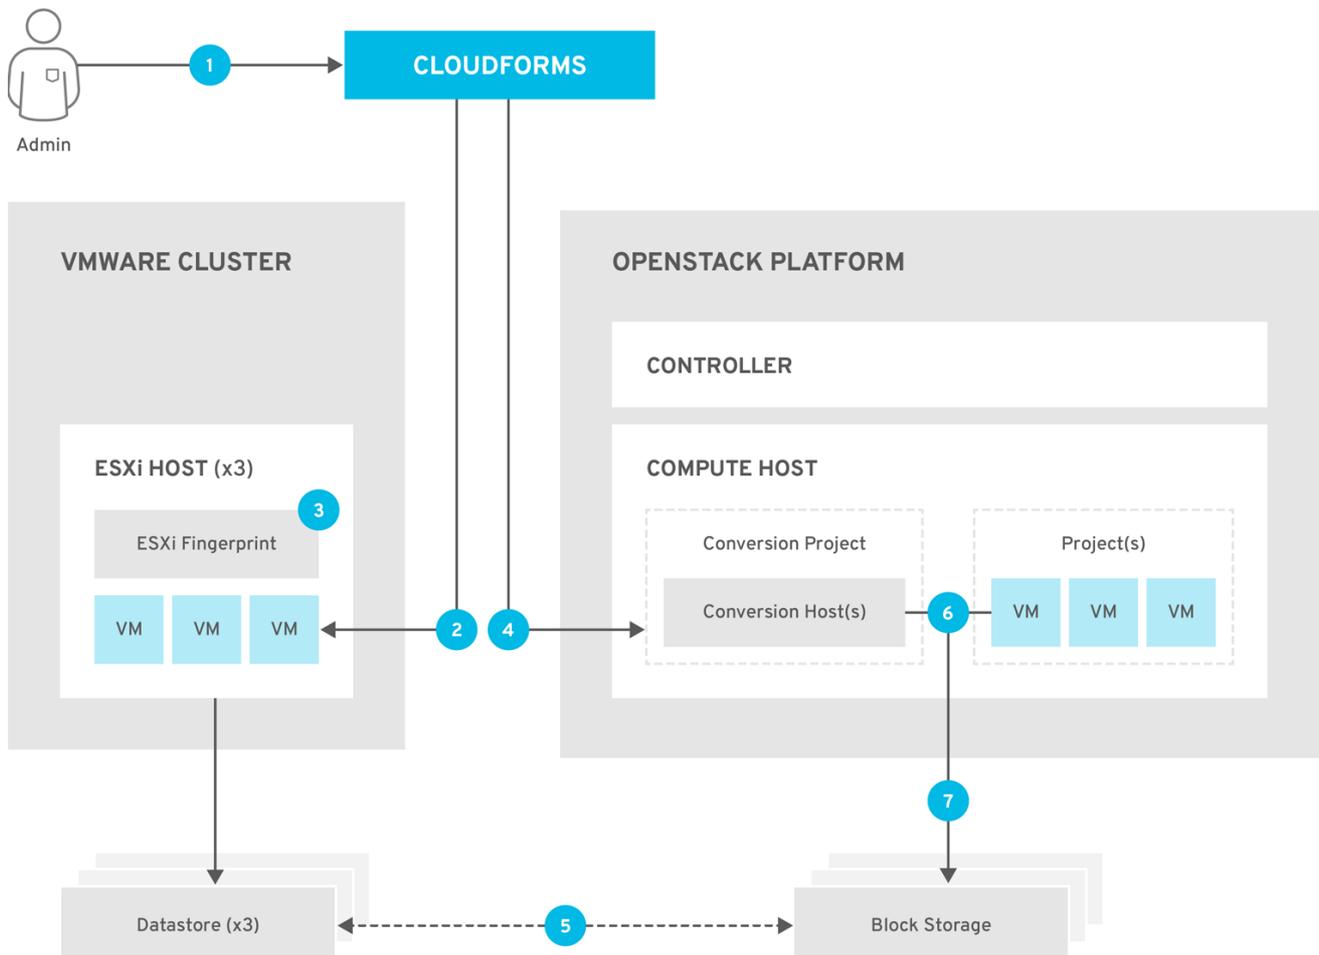


CLOUDFORMS\_474716\_0718

1. After creating an infrastructure mapping and a migration plan, you run the migration plan.
2. CloudForms uses the migration plan to locate the virtual machines to be migrated.
3. If VDDK transformation is used, the ESXi host fingerprint is captured for authentication during the virtual machine conversion process.
4. Using the Red Hat Virtualization attributes defined for the target environment, CloudForms initiates communication with the conversion hosts.
5. The conversion host connects to the source datastore through the ESXi host with **virt-v2v-wrapper**, streams the disks to be converted to the target data storage domain with **virt-v2v**, and converts the disks.
6. The target virtual machine is created in the Red Hat Virtualization environment. The target virtual machine uses the source virtual machine's metadata to maintain the virtual machine's attributes (tags, power state, MAC address, CPU count, memory, disks, and virtual machine name) after migration.
7. The converted disks are attached to the target virtual machine. The target virtual machine's power state is identical to its pre-migration state.  
The migration process is complete and the migration plan's status is displayed in CloudForms.

## 1.2. VMWARE TO RED HAT OPENSTACK PLATFORM WORKFLOW

This workflow describes the process of migrating from VMware to Red Hat OpenStack Platform.



CLOUDFORMS\_474716\_0718

1. After creating an infrastructure mapping and a migration plan, you run the migration plan.
2. CloudForms uses the migration plan to locate the virtual machines to be migrated.

**NOTE**

Source virtual machines must be powered on for migration.

3. If VDDK transformation is used, the ESXi host fingerprint is captured for authentication during the virtual machine conversion process.
4. Using the OpenStack Platform attributes defined for the target environment, CloudForms initiates communication with the conversion hosts.
5. The conversion host connects to the source datastore through the ESXi host with **virt-v2v-wrapper** and streams the disks to be converted to the target block storage with **virt-v2v**. The conversion host creates volumes in the block storage, attaches them to itself, and converts the source disks.
6. Once the disks are converted, **virt-v2v** detaches the volumes from the conversion host. **virt-v2v-wrapper** creates the target instance in the OpenStack Platform environment with the converted disks, using the flavor and security group defined in the migration plan and the network(s) defined in the infrastructure mapping.

7. The disks mapped in the block storage are attached to the instance and the instance is powered on.

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

## CHAPTER 2. PLANNING THE MIGRATION

During the planning phase, formulate a specific migration goal:

Example: "I want to migrate 2000 virtual machines, with 200 TB of data, in less than 6 months."

See the [planning questions](#) for guidance in allocating resources. A [migration example](#) is provided so that you can estimate the time duration of the migration.

[Best practices for migration](#) help you tailor the migration so that it fits your needs and schedule.

### 2.1. PLANNING QUESTIONS

Plan sufficient resources for your migration based on your responses to the following questions:

#### What am I migrating?

Identify the hosts and virtual machines to be migrated. If your deployment is large and complex, you may want to divide your resources into migration groups.

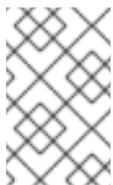
#### What is the maximum number of disks, virtual machines, or NICs that I can migrate?

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. The maximum number of NICs per virtual machine that you can migrate is 4.

If you exceed the capabilities of the environment, the migrations will fail and may negatively affect existing applications running on virtual machines attached to the network and storage. Plan carefully and perform test migrations to determine the capabilities of your environment.

#### What operating systems can I migrate?

For a list of supported guest operating systems, see [Converting Virtual Machines from Other Hypervisors to KVM with virt-v2v in RHEL 7](#). For a list of certified guest operating systems, see [Certified Guest Operating Systems in Red Hat OpenStack Platform and Red Hat Enterprise Virtualization](#).



#### NOTE

Although it is possible to migrate guest operating systems that are not certified for Red Hat Virtualization or OpenStack Platform, running these guests is not officially supported.

#### What resources will my target environment require?

Your target environment must be large enough to run the migrated virtual machines in addition to the original VMware virtual machines. The VMware virtual machines are not affected by the migration process.

#### What am I missing?

Identify resource gaps, such as bandwidth, storage, licenses, or a suitable maintenance window.

#### How long will the migration take?

There are no specific rules to determine how long the actual migration will take. This is determined on a case-by-case basis. A [migration example](#) is provided as a guide. 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.

#### What impact will the migration have on my users?

Assess the effect the migration may have on a production environment. Check whether users will lose access to critical applications.

### Migration example

This migration example, which took 1 hour and 30 minutes to complete, can help you estimate the duration of your migration process:

- Infrastructure:
  - 10 virtual machines (100 GB RAM)
  - Virtual machine disks 2/3 full (total of 1 TB of data)
- 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

## 2.2. BEST PRACTICES FOR MIGRATION

- Create migration groups, so that you are not migrating all of your virtual machines at the same time. The following questions provide guidelines:
  - How are the virtual machines grouped now?
  - Which virtual machines should be migrated together?
  - What workloads should be migrated together?
  - What applications must remain available?
- Consider which parts of the workload to migrate first:
  - Databases
  - Applications
  - Web servers
  - Load balancers
- Keep linked applications together, so that they migrate at the same time.
- Schedule your migration carefully, to minimize the impact on your users.
- Stagger the migration schedules.
- Move critical applications during maintenance windows.
- Plan for redundancy.

- Ensure that you have sufficient space for the migration. The migration process does not delete the original virtual machines.
- Minimize the impact on your environment by creating multiple migration plans to control and throttle the migration process.
- Perform test migrations with different [maximum numbers of concurrent migrations per conversion host](#) to assess the capabilities of your environment's infrastructure.
- Prepare your users for downtime.

## CHAPTER 3. PREPARING THE MIGRATION ENVIRONMENT

Preparing the environment for migration involves the following key steps:

- Preparing the [source environment](#)
- Preparing the target environment: [Red Hat Virtualization](#) or [Red Hat OpenStack Platform](#)
- Configuring the [conversion hosts](#)

### 3.1. SOFTWARE COMPATIBILITY MATRIX

You must use the supported product versions for each IMS release:

IMS Release	1.0	1.1
VMware	5.5 or later	5.5 or later
Red Hat Virtualization	4.2.7	4.2.8
CloudForms	4.6	4.7
OpenStack Platform	N/A	13 or later

### 3.2. PREPARING THE VMWARE ENVIRONMENT

Prepare the VMware source environment for migration:

1. Extend the VMware network to the target environment.



#### IMPORTANT

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

2. Unmount ISO/CDROM disks from the source virtual machines.
3. Check that the source virtual machines do not have encrypted disks.
4. Check that each NIC on each virtual machine has no more than 1 IPv4 and/or 1 IPv6 address.
5. Install [VMware Tools](#) on all source virtual machines to capture the IP addresses.
6. Check that source virtual machine names contain only upper- or lower-case letters, numbers, underscores (`_`), hyphens (`-`), or periods (`.`). Special characters and spaces are not permitted.
7. Red Hat Virtualization only: Check that the names of source virtual machines do not duplicate virtual machine names in the target environment.

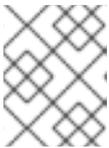
You are ready to prepare the target environment, [Red Hat Virtualization](#) or [OpenStack Platform](#).

### 3.3. PREPARING THE RED HAT VIRTUALIZATION ENVIRONMENT

Check the [software compatibility matrix](#) before you install the software.

To prepare the Red Hat Virtualization environment for migration:

1. Ensure that the datastore has sufficient space for the virtual machines you are migrating.
2. Set the BIOS settings of the physical hosts for optimal performance (rather than power-saving), according to the vendor's recommendations, and disable the C1E halt state, if applicable.
3. Install [Red Hat Virtualization Manager](#) on the Manager machine.
4. Install [Red Hat Virtualization Hosts](#) or [Red Hat Enterprise Linux hosts](#) 7.6 on the host machines.
5. Enable the following [ports](#) in the conversion hosts' network:
  - 22 - SSH
  - 443 - CloudForms, Red Hat Virtualization Manager, and VDDK
  - 902 - CloudForms to VMware
  - 5480 - `virt-v2v` to vCenter
6. Attach [data and ISO storage domains](#) to the data center.



#### NOTE

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

7. Upload the [VirtIO and Guest Tool images](#) to the ISO storage domain. The VirtIO image file name must include the version number: `virtio-win-version.iso`. The Guest Tool image is used to migrate Windows virtual machines.
8. Install Red Hat CloudForms [4.6](#) (with RHV 4.2.7) or [4.7](#) (with RHV 4.2.8) on the Manager machine.
9. Add the [VMware](#) and [Red Hat Virtualization](#) environments as providers in CloudForms.



#### NOTE

Do not remove or change a provider before, during, or after migration because it will have negative effects on CloudForms.

10. Configure the [conversion hosts](#) on physical machines, not nested virtual machines.

### 3.4. PREPARING THE RED HAT OPENSTACK PLATFORM ENVIRONMENT

Check the [software compatibility matrix](#) before you install the software.

To prepare the Red Hat OpenStack Platform environment for migration:

1. Ensure that the storage backends have sufficient space for the virtual machines you are migrating.
2. Set the BIOS settings of the physical hosts for optimal performance (rather than power-saving), according to the vendor's recommendations, and disable the C1E halt state, if applicable.
3. Install [Red Hat OpenStack Platform 13](#) or later.
4. Create [provider networks](#) for the target instances to ensure that the IP addresses of the source virtual machines are preserved.
5. [Create a project](#) for the conversion hosts and whatever destination projects you require for the target instances.
6. Ensure that the [admin user has member and admin roles](#) in the conversion and destination projects.
7. [Set at least one volume type](#) for the target block storage. Otherwise, CloudForms will not detect the storage during infrastructure mapping.
8. Ensure that the user configured in CloudForms for the OpenStack Platform provider has **admin** privileges on the target projects.
9. Configure security groups:
  - A security group for the conversion hosts and CloudForms must have port 22 (SSH) enabled.
  - A security group for CloudForms must have port 443 (HTTPS) enabled.

**NOTE**

Outbound traffic is enabled by default. If you have changed this setting, enable ports 902 (CloudForms to VMware) and 5480 (**virt-v2v** to vCenter).

10. Create flavors suitable for the virtual machines that you are migrating. If you do not create custom flavors, CloudForms will try to map the source virtual machine to an existing flavor.
11. Install [Red Hat CloudForms 4.7](#) or later.
12. Add the [VMware](#) and [OpenStack Platform](#) environments as providers in CloudForms.

**NOTE**

Do not remove or change a provider before, during, or after migration because it will have negative effects on CloudForms.

13. Configure the [conversion hosts](#).

**NOTE**

For best performance, deploy the conversion hosts on a compute node with [nested virtualization](#) enabled. This is a technology preview.

## 3.5. CONVERSION HOSTS

### 3.5.1. Conversion host considerations

Before creating the conversion hosts, you must consider the following questions, taking into account your environment and the size of the migration:

- How many conversion hosts do I need? (see [Section 3.5.1.1, “Number of conversion hosts”](#))
- Should I use VDDK or SSH transformation? (see [Section 3.5.1.2, “Transformation options: VDDK or SSH”](#))

#### 3.5.1.1. Number of conversion hosts

Because all the virtual machines in a migration plan are migrated at the same time, the number of conversion hosts you require depends on the number of virtual machines you are migrating. Multiple conversion hosts are recommended for load-balancing purposes, even for smaller migrations.

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. You should thoroughly test your environment before the migration to determine how many migrations it can support without negative effects.

Each conversion host is limited to **10** concurrent migrations, unless you [change the default value](#).



#### IMPORTANT

The maximum number of concurrent virtual machine migrations per conversion host is **20** for both VDDK and SSH transformation. The SSH limitation, however, applies only to RHV 4.2.5, not to later versions of RHV.

#### 3.5.1.2. Transformation options: VDDK or SSH

You must select the transformation option, VDDK or SSH, that is appropriate for your migration size and environment:

- **VDDK**
  - Recommended because it is approximately twice as fast as SSH
  - Compiled with `nbdkit` to provide direct connectivity to the virtual machine disks
  - Requires VDDK package download
  - Requires the location of the VDDK library for conversion host configuration
- **SSH**
  - Slower than VDDK
  - Better suited for large-scale migrations than VDDK
  - Does not require additional downloads, packages, tools, or compilation
  - SSH access must be enabled on the source hypervisors

- o Source and target environments must share SSH keys

### 3.5.2. Creating conversion hosts

If you are migrating to Red Hat Virtualization, select one or more existing hosts (on physical machines, not virtual machines) to use as conversion hosts.

If you are migrating to OpenStack Platform:

1. Download the OpenStack Platform conversion appliance:
  - a. Go to the [Red Hat Customer Portal Product Downloads](#) site.
  - b. In the **A-Z** tab, click **Red Hat OpenStack Platform**.
  - c. Click the green **Download Latest** button to go to the OpenStack Platform download page.
  - d. Click the **Product Software** tab and download the **RHOSP V2V Image for Red Hat OpenStack Director (x86\_64)**.



#### NOTE

Always download the latest version of the OpenStack Platform conversion appliance because it contains patches and fixes.

The conversion appliance for OpenStack Platform 14 may be used with OpenStack Platform 13.

2. Upload the appliance to Glance storage.
3. Deploy the appliance as a conversion host instance, with the following resources:
  - 4 vCPUs
  - 8 GB RAM
  - **/tmp** (1 GB for each conversion)
  - **/var/tmp** (1 GB for each conversion)
  - **/var/logs** (5 GB)
4. Add the SSH key for **ccloud-user** in CloudForms:
  - a. Click **Compute** → **Clouds** → **Providers** and select the OpenStack Platform provider.
  - b. From the **Configuration** dropdown menu, select **Edit Selected Cloud Provider**.
  - c. Click the **RSA key pair** tab in the **Endpoints** section.
  - d. Enter the **Username**, **ccloud-user**.
  - e. Click **Choose File**, browse to the SSH key file that you use to create OpenStack Platform instances, and click **Save**.

You are ready to configure the conversion hosts for [transformation](#).

### 3.5.3. Configuring conversion hosts for transformation

You can configure your conversion hosts for VDDK or SSH transformation. Configuring the conversion hosts for transformation involves the following key steps:

1. VDDK only: [Downloading and copying the VDDK package](#)
2. VDDK and SSH: [Configuring the conversion hosts](#)
3. VDDK and SSH: [Validating the transformation configuration](#)
4. SSH only: [Configuring and validating the SSH connection](#)

#### 3.5.3.1. VDDK only: Downloading and copying the VDDK package

1. Go to the [VMware Documentation site](#), click **VMware SDK & API Product Documentation** to expand, and click **VMware Virtual Disk Development Kit (VDDK)**.
2. Click **Latest Releases** and select the latest VDDK release.
3. Download the VDDK package, **VMware-vix-disklib-\_version\_.x86\_64.tar.gz**.
4. Copy the VDDK package to a location that is accessible over **HTTP**, for example, **/var/www/html** on the Red Hat Virtualization Manager machine.

You are ready to configure the conversion hosts.

#### 3.5.3.2. VDDK and SSH: Configuring the conversion hosts

1. Go to **/usr/share/ovirt-ansible-v2v-conversion-host/playbooks** on the Manager machine (for Red Hat Virtualization) or on each conversion host (for OpenStack Platform).
2. Generate an SSH key pair to connect to the VMware hypervisors over SSH. You must create the keys regardless of the chosen transformation method.
3. Create **extra\_vars.yml** according to the following examples:
  - Red Hat Virtualization conversion host using VDDK:

```

---
v2v_host_type: rhv

# Transport methods to configure on the conversion host. Valid
values: vddk, ssh (on separate lines)
v2v_transport_methods:
  - vddk

# Maximum number of concurrent conversions per host. Default is
10.
v2v_max_concurrent_conversions: 10

# File name of VDDK package
v2v_vddk_package_name: "VMware-vix-disklib-version.x86_64.tar.gz"

# URL of VDDK package
v2v_vddk_package_url: "http://path/to/downloaded_vddk_package/{{

```

```

v2v_vddk_package_name }}"

manageiq_provider_name: RHV

# Base URL of CloudForms machine
manageiq_url: "https://CloudForms_FQDN"

# Whether to validate certificate of CloudForms server. Default
is true
manageiq_validate_certs: false
manageiq_zone_id: "42000000000001"

# List of cloud or infrastructure providers
# Each provider is a dictionary with 3 attributes: `name`,
`hostname`, and `connection_configurations`
manageiq_providers:
  - name: "RHV"
    hostname: FQDN_or_IP 1
    connection_configurations: 2
      - endpoint:
          role: "default"
          verify_ssl: true
          certificate_authority: |
            -----BEGIN CERTIFICATE-----
MIIDoDCCAoigAwIBAgIBATANBgkqhkiG9w0BAQsFADA9MRswGQYDVQ...
            -----END CERTIFICATE-----

```

- 1 **hostname** is the *FQDN\_or\_IP* of the RHV Manager.
- 2 **RHV connection\_configurations** has a single endpoint, whose role is **default**. You can choose to verify the SSL certificate (**verify\_ssl: true**) and provide the CA chain (**certificate\_authority**), found in `/etc/pki/ovirt-engine/apache-ca.pem` on the Manager machine.

- OpenStack Platform conversion host using SSH:

```

---
v2v_host_type: openstack
v2v_transport_methods:
  - ssh
v2v_max_concurrent_conversions: 10
manageiq_provider_name: OpenStack
manageiq_url: "https://CloudForms_FQDN"
manageiq_validate_certs: false
manageiq_zone_id: "42000000000001"
manageiq_providers:
  - name: "OpenStack"
    hostname: FQDN_or_IP 1
    connection_configurations: 2
      - endpoint:
          role: "default"
# Valid values: `non-ssl`, `ssl-without-validation`, `ssl`. `ssl`
requires `certificate_authority`.

```

```

security_protocol: "ssl"
certificate_authority: |
    -----BEGIN TRUSTED CERTIFICATE-----
MIIDNzCCA8CAQEWdQYJKoZIhvcNAQELBQAwYjELMAkGA1UEBhMCVV...
    -----END TRUSTED CERTIFICATE-----
    -----BEGIN TRUSTED CERTIFICATE-----
MIIDLzCCAn+gAwIBAgIJAOP7AaT7dsLYMA0GCSqGSIb3DQEBCwUAMG...
    -----END TRUSTED CERTIFICATE-----

```

- 1 **hostname** is the *FQDN\_or\_IP* of the OpenStack Platform controller.
- 2 OpenStack Platform **connection\_configurations** has a single endpoint, whose role is **default**. You can specify the connection security: **non-ssl**, **ssl-without-validation**, or **ssl**.

For **ssl**, you must provide a CA chain (`certificate_authority`), which is a concatenation of the following CA files:

- o `/etc/pki/ca-trust/source/anchors/undercloud-cacert.pem` on the undercloud server
- o `/etc/pki/ca-trust/anchors/overcloud-cacert.pem` on one of the overcloud controllers  
If you deploy your own [certificate authority](#), use the chain that signs the OpenStack Platform API certificates.

4. Create an encrypted `secure_vars.yml` file for your passwords and keys:

```
# ansible-vault create secure_vars.yml
```

5. Add the following lines to `secure_vars.yml`:

```

---
# Use this file for sensitive information
# CloudForms `admin` user, for connecting to CloudForms
manageiq_username: "username"

# CloudForms `admin` password:
manageiq_password: "password"

# SSH private key to connect to VMware hypervisors.
# This key is required for both SSH and VDDK transformation

v2v_ssh_private_key: |
    -----BEGIN RSA PRIVATE KEY-----
b3B1bnNzaC1rZXktdjEAAAABAG5vbmUAAAABbm9uZQAAAAAAAAABAAAAIwAAAAAdzc2gt
cn....
    -----END RSA PRIVATE KEY-----

```

6. Run the **conversion\_host\_enable** playbook to configure the conversion hosts for transformation:

```
# ansible-playbook -i hostname, -c local -b \ 1
-e @extra_vars.yml -e @secure_vars.yml --ask-vault-pass \
/usr/share/ovirt-ansible-v2v-conversion-
host/playbooks/conversion_host_enable.yml
```

- 1** **hostname** is the *FQDN* or IP address of the CloudForms provider endpoint. For Red Hat Virtualization, the CloudForms provider endpoint is the RHV Manager. For OpenStack Platform, it is the controller.



### IMPORTANT

Running the **conversion\_host\_enable** playbook more than once on the same host creates multiple entries in the CloudForms database for the same host. If you need to run the **conversion\_host\_enable** playbook again, first remove the existing database entry by running the **conversion\_host\_disable** playbook on the same host:

```
# ansible-playbook /usr/share/ovirt-ansible-v2v-
conversion-host/playbooks/conversion_host_disable.yml
```

#### 3.5.3.3. VDDK and SSH: Validating the transformation configuration

Run the **conversion\_host\_check** playbook to verify that the conversion hosts are configured correctly:

```
# ansible-playbook --ask-vault-pass -i `hostname`,
conversion_host_check.yml
```

#### VDDK only:

- If you are configuring Red Hat Virtualization conversion hosts for VDDK transformation, you must [enable them](#) in CloudForms.
- If you are configuring OpenStack Platform conversion hosts for VDDK transformation, you can create an [infrastructure mapping](#).

**SSH only:** If you are configuring conversion hosts for SSH transformation, you must configure and validate the SSH connection.

#### 3.5.3.4. SSH only: Configuring and validating the SSH connection

1. Enable SSH on all VMware hypervisors.  
For instructions, go to the [VMware vSphere Documentation site](#). In the navigation pane, click **vSphere version** → **ESXi and vCenter Server** → **VMware ESXi Installation and Setup** → **Installing and Setting Up ESXi** → **Setting Up ESXi** → **Enable ESXi Shell and SSH Access with the Direct Console User Interface**.
2. Copy the public SSH key corresponding to **v2v\_ssh\_private\_key** in **secure\_vars.yml** (see [Section 3.5.3.2, “VDDK and SSH: Configuring the conversion hosts”](#)) to each VMware hypervisor:
  - RHV Manager:
    -

```
# ssh root@esx1.example.com sh -c \
    cat >> /etc/ssh/keys-root/authorized_keys <
/var/lib/vdsm/.ssh/id_rsa.pub
```

- OpenStack Platform conversion host:

```
# ssh root@_esx1.example.com_ sh -c \
    'cat >> /etc/ssh/keys-root/authorized_keys' <
/OpenStack_Platform/conversion_host_key/id_rsa.pub
```

### 3. Validate the SSH connection:

- RHV Manager:

```
# sudo -u vdsms ssh-agent
SSH_AUTH_SOCK=/tmp/ssh-11111AAAAA/agent.12345; export
SSH_AUTH_SOCK;
SSH_AGENT_PID=12345; export SSH_AGENT_PID;
echo Agent pid 12345;

# sudo -u vdsms SSH_AUTH_SOCK=/tmp/ssh-11111AAAAA/agent.12345 ssh-
add
# sudo -u vdsms \
    SSH_AUTH_SOCK=/tmp/ssh-123456ABCDE/agent.12345 ssh
root@esx1.example.com
```

- OpenStack Platform conversion host: Connect over SSH to the VM hypervisor as **root**. If the connection is successful, the conversion host is correctly configured for SSH transformation.



## IMPORTANT

**Red Hat Virtualization only:** If you are using SSSD with single sign-on, you must reinstall **ipa-client** without configuring the OpenSSH client. See [SSH transformation failure](#) for details.

You are ready to [enable the conversion hosts](#) in CloudForms.

### 3.5.4. Enabling conversion hosts in CloudForms

OpenStack Platform conversion hosts do not require configuration in CloudForms.

Authenticate the Red Hat Virtualization conversion hosts:

1. Click **Compute** → **Infrastructure** → **Hosts** and select a Red Hat Virtualization conversion host.
2. Click the **Configuration** drop-down button, and select **Edit Selected items**.
3. In the **Default** tab of the Endpoints section, enter the **Username root** and the password.
4. Click **Validate**.
5. Click **Save**.



```
# vmdb
# rails console
irb(main):001:0> $evm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $evm.vmdb(:host).find_by(:name =>
"host1.example.com").custom_set("Max Transformation Runners", 20)
```

### Changing the maximum number of concurrent migrations for a provider

To set a provider's limit to **30**, enter this command on the CloudForms machine:

```
# vmdb
# rails console
irb(main):001:0> $evm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $evm.vmdb(:ext_management_system).find_by(:name =>
"provider").custom_set("Max Transformation Runners", 30)
```

To get a provider's limit, enter this command on the CloudForms machine:

```
# vmdb
# rails console
irb(main):001:0> $evm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $evm.vmdb(:ext_management_system).find_by(:name =>
"provider").custom_get("Max Transformation Runners")
```

### 3.5.7. Upgrading conversion hosts

- To upgrade Red Hat Virtualization conversion hosts, enter the following commands on the Manager machine:

```
# yum update
# cp /usr/share/ansible/roles/oVirt.v2v-conversion-host/files/virt-
v2v-wrapper.py /usr/bin/
```

- To upgrade OpenStack Platform conversion hosts, download the latest conversion appliance and create new conversion hosts.

## CHAPTER 4. MIGRATING THE INFRASTRUCTURE

Migrating the infrastructure involves the following key tasks:

- Creating an [infrastructure mapping](#) to map the clusters, datastores, and networks of your source environment to your target environment
- Creating and running a [migration plan](#)



### NOTE

You can [cancel](#) a migration plan that is in progress and retry the migration in the future.

### 4.1. CREATING AN INFRASTRUCTURE MAPPING

The infrastructure mapping maps your source resources to target resources.

#### Infrastructure Mappings screen

#### Prerequisite

**Red Hat Virtualization only:** Before creating an infrastructure mapping, you must [authenticate](#) the RHV conversion hosts. If the hosts are not authenticated, the mapping displays a [missing networks error](#).

#### Procedure

1. Click **Compute** → **Migration** → **Infrastructure Mappings**.
2. Click **Create Infrastructure Mapping**.
3. Perform the following procedures in each infrastructure mapping screen:

Screen	Procedure
General	<ol style="list-style-type: none"> <li>1. Enter the infrastructure mapping name and (optional) description.</li> <li>2. Select the target provider from the drop-down list.</li> <li>3. Click <b>Next</b>.</li> </ol>

Screen	Procedure
<b>Map Compute</b>	<ol style="list-style-type: none"> <li>1. Select a source cluster and a target Red Hat Virtualization cluster or OpenStack Platform project. <b>Red Hat Virtualization only:</b> A target cluster that does not contain a conversion host displays a warning icon (  ). The icon appears after a few seconds' delay for polling the clusters and refreshing the UI.</li> <li>2. Click <b>Add Mapping</b>. You can map additional clusters.</li> <li>3. Click <b>Next</b>.</li> </ol>
<b>Map Storage</b>	<ol style="list-style-type: none"> <li>1. Select a source cluster from the drop-down list.</li> <li>2. Select a source datastore and a target Red Hat Virtualization datastore or OpenStack Platform volume type. <b>OpenStack Platform only:</b> If the target datastore is missing, see <a href="#">Storage not detected.</a>)</li> <li>3. Click <b>Add Mapping</b>. You can map additional datastores.</li> <li>4. Click <b>Next</b>.</li> </ol>
<b>Map Networks</b>	<ol style="list-style-type: none"> <li>1. Select a source cluster from the drop-down list.</li> <li>2. Select one or more source networks and a target network.</li> <li>3. Click <b>Add Mapping</b>. You can map the networks of additional source clusters.</li> <li>4. Click <b>Create</b>.</li> </ol>
<b>Results</b>	<p>Click <b>Close</b>. Your infrastructure mapping is saved in <b>Compute</b> → <b>Migration</b> → <b>Infrastructure Mappings</b>.</p> <p>If error messages appear, see <a href="#">Infrastructure mapping errors</a>.</p>

**IMPORTANT**

If you add or remove providers or provider objects from an infrastructure mapping, the mapping will have [missing resource errors](#). You must re-create the infrastructure mapping.

After you have created a successful infrastructure mapping, you can create a [migration plan](#).

## 4.2. CREATING AND RUNNING A MIGRATION PLAN

The infrastructure migration involves the following key tasks:

1. [Pre-migration tasks](#):

- **CSV file:** If you are migrating a large number of virtual machines, you can create a [CSV file](#) to add the virtual machines to the migration plan. If you are [re-migrating virtual machines](#), you must use a CSV file to add the virtual machines because the migration plan cannot re-discover them.
- **Ansible environment:** If you are using Ansible to automate pre- and post-migration tasks, you must set up an [Ansible environment](#).
- **RHEL pre-migration playbook:** If you are migrating RHEL or other Linux virtual machines, you must set up an [Ansible environment](#) and create a [RHEL pre-migration playbook](#).

## 2. [Creating the migration plan](#)

### 3. [Running the migration plan](#) immediately or scheduling it to run in the future

You can [cancel](#) a migration plan in progress.

## 4.2.1. Pre-migration tasks

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

If you are migrating a large number of virtual machines, you can create a CSV file to add the source virtual machines to the migration plan. If you are [re-migrating virtual machines](#), you must use a CSV file.

The virtual machine **Name** is mandatory. **Provider**, **Host**, and **Cluster** are optional, to differentiate identical virtual machine names in different environments.

#### Sample CSV file with **Name (mandatory)** and **Provider (optional)**

```
Name,Provider
vm01,vSphere3
vm02,vSphere3
vm03,vSphere3
...
```

Duplicate virtual machines with no optional fields to differentiate them, or with duplicate optional fields, are marked with a red **x** and cannot be added to the migration plan.

### 4.2.1.2. Setting up an Ansible environment

You can perform automated tasks before and after a virtual machine is migrated, such as removing web servers from a load-balancing pool before migration and adding them back to the pool after migration.

Each task is an Ansible service catalog item, which you select in **Advanced Options** when you create the migration plan. You must create the Ansible service catalog item before creating the migration plan. The CloudForms inventory requires a minimum of one repository, one playbook, and one set of credentials.

To set up an Ansible environment and create a service catalog item:

1. Configure an [Ansible repository](#) that is accessible to CloudForms.
2. Create [Ansible credentials](#) for the virtual machines.
3. Create an [Ansible service catalog item](#).

### 4.2.1.3. Creating a RHEL pre-migration playbook for migrating RHEL virtual machines

If you are migrating virtual machines with RHEL or other Linux operating systems, you must create a **RHEL pre-migration** playbook that calls the Ansible `ims.rhel_pre_migration` role. You will select this playbook to run as a **pre-migration task** when you create the migration plan.



#### IMPORTANT

This role is required for both Red Hat Virtualization and OpenStack Platform migrations. Otherwise, the migration will appear to be successful but you will not be able to access the IP addresses of the target virtual machines.

The **RHEL pre-migration** role performs the following tasks on a source virtual machine:

- Preserves the static IP address configuration by creating udev rules to associate the source virtual machine's MAC address with its existing interface name (Red Hat Virtualization and OpenStack Platform)
- Installs the Red Hat Virtualization guest agent, which enables Red Hat Virtualization to report the IP address of the new virtual machine and to integrate it into the new environment (Red Hat Virtualization only)

To create a **RHEL pre-migration** playbook:

1. Set up an [Ansible environment](#).
2. Create a **RHEL pre-migration** playbook that calls the `ims.rhel_pre_migration` role:

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



#### NOTE

The **RHEL pre-migration** role assumes that the `rhel-6-server-rpms` or `rhel-7-server-rpms` repository is enabled in the source virtual machine when it installs `qemu-guest-agent`. If you have disabled the repository, you must enable it in the playbook.

### 4.2.2. Creating a migration plan

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click **Create Migration Plan**.
3. Perform the following procedures in each migration plan screen:

Screen	Procedure
--------	-----------

Screen	Procedure
<b>General</b>	<ol style="list-style-type: none"> <li>1. Select the infrastructure mapping from the drop-down list. If you select an OpenStack Platform mapping, an <b>Instance Properties</b> screen appears after the <b>VMs</b> screen.</li> <li>2. Enter the migration plan name and (optional) description.</li> <li>3. Select a virtual machine discovery method: <ul style="list-style-type: none"> <li>• <b>Choose from a list of VMs discovered in the selected infrastructure mapping</b></li> <li>• <b>Import a CSV file with a list of VMs to be migrated</b></li> </ul> </li> <li>4. Click <b>Next</b>.</li> </ol>
<b>VMs</b>	<ol style="list-style-type: none"> <li>1. If you selected <b>Choose from a list of VMs discovered in the selected infrastructure mapping</b>, select the virtual machines to be migrated. If you selected <b>Import a CSV file with a list of VMs to be migrated</b> : <ol style="list-style-type: none"> <li>a. Click <b>Import</b>.</li> <li>b. Browse to the CSV file and click <b>Open</b>.</li> </ol> </li> <li>2. Click <b>Next</b>.</li> </ol> <p>If you cannot discover or import the source virtual machines, see <a href="#">Migration plan errors</a>.</p>
<b>Instance Properties</b>  (OpenStack Platform only)	<ol style="list-style-type: none"> <li>1. Click the pencil icon to edit the network or flavor of each selected virtual machine. If a flavor is too small for the source virtual machine, it is marked with an asterisk (*).</li> <li>2. Click <b>Next</b>.</li> </ol>
<b>Advanced Options</b>  (Optional)	<ol style="list-style-type: none"> <li>1. Select a pre-migration playbook service.</li> <li>2. Select a post-migration playbook service.</li> <li>3. Select the virtual machines on which to run the playbook services.</li> <li>4. Click <b>Next</b>.</li> </ol>
<b>Schedule</b>	<ol style="list-style-type: none"> <li>1. Select a schedule option: <ul style="list-style-type: none"> <li>• <b>Save migration plan to run later</b></li> <li>• <b>Start migration immediately</b></li> </ul> </li> <li>2. Click <b>Create</b>.</li> </ol>

Screen	Procedure
Results	Click <b>Close</b> .

### 4.2.3. Running a migration plan

You can run the migration plan immediately or in the future, depending on the schedule option that you selected.

#### 4.2.3.1. Running a migration plan 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.

To view the migration progress of individual virtual machines:

1. Click **Migration Plans in Progress**.
2. Click the migration plan name to display the migration plan details.



#### NOTE

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

#### 4.2.3.2. Running a migration plan in the future

The migration plan is saved in **Migration Plans Not Started** and will not run until you schedule or run it.

- To schedule a migration plan:
  1. Click the **Schedule** button.
  2. 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.
- To run a scheduled plan immediately, click **Migrate**.
- To edit or delete a migration plan, click the **More Actions** icon.

### 4.2.4. Validating a migration plan

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

A **Denied State** error indicates a problem with the conversion hosts (see [Denied state error](#) for more information).

To troubleshoot a failed migration, see [Chapter 5, Troubleshooting](#).

## 4.2.5. Canceling a migration plan in progress

You can cancel a migration plan that has started:

1. In **Compute** → **Migration** → **Migration Plans**, 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.

You can click **Retry** to run the migration plan immediately or **Schedule** to run the migration plan in the future.

Click the **More Actions** button to archive, edit, or delete the migration plan.

## CHAPTER 5. TROUBLESHOOTING

If errors occur for reasons unrelated to the infrastructure mapping or migration plan (for example, network outage), you can retry the migration plan or re-migrate previously migrated virtual machines. See [Section 5.1, “Retrying and re-migrating”](#).

If the cause of the failed migration is not clear, see [Section 5.2, “Logs”](#) and [Section 5.3, “Common issues and mistakes”](#).

### 5.1. RETRYING AND RE-MIGRATING

When virtual machines fail to migrate, the migration plan’s status is **Failed**. You must clean up the target environment if the migration failed during the conversion phase retrying a failed migration plan or re-migrating virtual machines.

#### Cleaning up a failed migration

1. Delete newly created target virtual machines or instances to avoid name conflicts with source virtual machines.
2. Delete newly created disks in the target datastore to free up space.
3. OpenStack Platform only: Delete newly created network ports of failed instances.

#### Retrying a failed migration plan

If the source virtual machines have never been migrated, you can retry the failed migration plan:

1. Click **Compute** → **Migration** → **Migration Plans**.
2. Click **Migration Plans Complete**.
3. Click the **Retry** button beside the failed migration plan.

#### Re-migrating virtual machines

If the virtual machines have been migrated previously, you must re-migrate the virtual machines using a CSV file:

1. Create a [CSV file](#) with the names of the source virtual machines.
2. Create a new migration plan using the CSV file to add the virtual machines.



#### NOTE

The source virtual machines cannot be re-discovered by the migration plan because they are marked in the CloudForms database as **migrated**.

### 5.2. LOGS

You can check the following logs to troubleshoot a failed migration.

#### CloudForms

`/var/www/miq/vmdb/log/automation.log`. This log contains the steps for orchestrating the virtual machine migrations in CloudForms, in order to troubleshoot migration errors.

## Conversion hosts

- Red Hat Virtualization: `/var/log/vdsm/import/v2v-import-YYYYMMDDTHHMMSS-1243.log` and `v2v-import-YYYYMMDDTHHMMSS-1243-wrapper.log`
- OpenStack Platform: `/var/log/virt-v2v/v2v-import-YYYYMMDDTHHMMSS-1243.log` and `v2v-import-YYYYMMDDTHHMMSS-1243-wrapper.log`

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

- **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.
- **virt-v2v**: Debug output from **virt-v2v** itself. This log tracks the core of the virtual machine process, including **libguestfs** traces and disk migration details.



### IMPORTANT

Both logs are required for bug analysis and must be provided if you open a [Red Hat Support](#) call.

## Virtual machine

1. In **Compute** → **Migration Plans** → **Migration Plans Complete**, click a migration plan name to go to the details view.
2. Click **Download Log** → **Migration log** of a virtual machine.

## 5.3. COMMON ISSUES AND MISTAKES

### 5.3.1. Infrastructure mapping errors

- Missing datastores or clusters. This error occurs when you create a mapping and subsequently add or remove providers or provider objects, which changes the object IDs. Delete the mapping and create a new one.

#### Red Hat Virtualization only

- Missing networks. This error occurs if you do not add the conversion host credentials in CloudForms. Delete the mapping, [enable the conversion hosts in CloudForms](#), and create a new mapping.
- Missing networks and datastore. This error occurs when you create a mapping and subsequently refresh the hosts, which changes the object IDs. Delete the mapping and create a new one.

#### OpenStack Platform only

- Storage not detected. If the target block storage does not appear while you are creating an infrastructure mapping, check that you have [set at least one volume type](#) for the storage.

### 5.3.2. Migration plan errors

- Virtual machines cannot be discovered. If a migration plan cannot discover source virtual machines that have never been migrated in the past, check that the datastores and networks of the source virtual machines have been added to the Infrastructure Mapping. If a migration plan cannot discover source virtual machines that have been migrated in the past, [re-migrate](#) the virtual machines. Previously migrated virtual machines cannot be discovered because they are marked as **migrated** in the CloudForms VMDB.
- Virtual machines cannot be added to the migration plan with a CSV file. If you cannot add the source virtual machines into a migration plan with a CSV file, check the [CSV file's format and fields](#). Create a new migration plan using the updated CSV file.
- **Denied State.** If a migration plan fails immediately and the migration plan displays a **Denied State** error message, check that you have [created the conversion hosts](#). (Red Hat Virtualization only: ensure that you have [authenticated the conversion hosts](#)). Cancel the migration plan and [retry it](#).
- IP address of RHEL or other Linux virtual machine not accessible after migration. If you cannot access the IP address of a RHEL or other Linux virtual machine after an apparently successful migration, [re-migrate](#) the virtual machine while running the [RHEL pre-migration playbook](#).
- No IP address. VMware Tools not installed on source virtual machine. VMware Tools preserves the IP address for migration.

### 5.3.3. Virtual machine migration failure

If a virtual machine fails to migrate, check its [conversion log](#) and review the possible causes below.

The following scenarios can cause a failed migration:

- Mounted ISO/CDROM disk(s)
- Encrypted disk(s)
- Invalid name. The source virtual machine name must not contain spaces or special characters.

#### Red Hat Virtualization only

- Virtual machine name duplicates an existing target virtual machine name.
- SSH transformation failure. If you are using SSSD with single sign-on, SSH fails for the **vds** account. Reinstall **ipa-client** without configuring the OpenSSH client:

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



#### NOTE

You cannot resolve this issue by changing the configuration file because the file is restored during upgrades.

#### OpenStack Platform only

- Virtual machine powered off. The virtual machine must be powered on at the time of migration.

- No IP address. Source virtual machine has an interface configuration file that is mapped to a non-existent interface (for example, `/etc/sysconfig/network-scripts/ifcfg-eth1` exists, but `eth1` interface does not).

```
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 error
```

- Not Authorized. The 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.

```
ERROR: Command exited with non-zero return code 1, output:
HttpException: 403: Client Error for url:
https://123.123.123.123:13696/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": ""}}
```

## 5.4. KNOWN ISSUES

After installing IMS, you must obtain critical updates for Red Hat Virtualization and the OpenStack Platform conversion appliance by doing the following:

- Red Hat Virtualization: Run `yum update`.
- OpenStack Platform: Contact [Red Hat Support](#) for a newer version of the appliance build.

These items will be addressed in a future release.

### All platforms

- Canceling a migration does not stop the creation of volumes, instances, and network ports (OpenStack Platform), or virtual machines (Red Hat Virtualization). [BZ#1666799](#)
- Virtual machines whose names contain international characters or spaces fail to migrate using SSH transformation. [BZ#1669240](#)
- Virtual machines whose names contain international characters or spaces fail to migrate using VDDK transformation. [BZ#1669286](#)

### Red Hat Virtualization

- Refreshing the hosts causes the network(s) and datastore to disappear from infrastructure mappings. [BZ#1669176](#)

### OpenStack Platform

- Instance is not created after disk conversion. [BZ#1668049](#)
- Names of virtual machines that were migrated using SSH transformation change during migration. [BZ#1669133](#)