



Red Hat Infrastructure Migration Solution 1.0

Infrastructure Migration Solution Guide

Migrating from VMware to Red Hat Virtualization

Red Hat Infrastructure Migration Solution 1.0 Infrastructure Migration Solution Guide

Migrating from VMware to Red Hat Virtualization

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

Legal Notice

Copyright © 2018 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 migrate your virtual machine infrastructure from VMware to Red Hat Virtualization using Red Hat CloudForms.

Table of Contents

CHAPTER 1. INTRODUCTION	3
CHAPTER 2. INSTALLATION	4
2.1. PREREQUISITES	4
2.2. RHV CONVERSION HOSTS	5
2.2.1. Creating and Configuring RHV Conversion Hosts for VDDK Transformation	6
2.2.2. Creating and Configuring RHV Conversion Hosts for SSH Transformation	7
Additional resources	9
2.2.3. Configuring RHV Conversion Hosts in CloudForms	9
2.2.4. Configuring the Maximum Number of Concurrent Migrations	9
CHAPTER 3. DISCOVERY AND ASSESSMENT	11
3.1. DISCOVERY QUESTIONS	11
3.2. MIGRATION EXAMPLE	12
3.3. BEST PRACTICES	12
CHAPTER 4. PLANNING	14
4.1. CREATING AN INFRASTRUCTURE MAPPING	14
4.2. CREATING AND RUNNING A MIGRATION PLAN	15
Prerequisites	15
Procedure	15
4.3. ADVANCED MIGRATION OPTION: AUTOMATING PRE- AND POST-MIGRATION TASKS WITH ANSIBLE	17
Creating an Ansible Service Catalog Item	18
IMS RHEL Pre-migration Ansible Playbook Example	18
CHAPTER 5. MIGRATION	19
5.1. WORKFLOW	19
5.2. TROUBLESHOOTING	20
Remigrating Virtual Machines	20
Retrying a Migration Plan	20
Log Files	20
Common Issues and Mistakes	21

CHAPTER 1. INTRODUCTION

The Red Hat Infrastructure Migration Solution enables you to move your virtual machines from VMware to Red Hat Virtualization.

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.

The infrastructure migration is performed with Red Hat CloudForms. CloudForms is a user interface, providing the control and automation required to manage challenging enterprise environments.

The infrastructure migration process has 4 key phases:

- Installation:
 - Installing and configuring Red Hat Virtualization, Red Hat Virtualization Host(s) or Red Hat Enterprise Linux host(s), and Red Hat CloudForms
 - Creating and configuring the RHV conversion host(s) and configuring the maximum number of concurrent migrations.
- Discovery and Assessment:
 - Discovering and assessing your environment
 - Migration example
 - Best practices
- Planning:
 - Creating an infrastructure mapping
 - Creating and running a migration plan
- Migration:
 - Migration workflow
 - Troubleshooting

CHAPTER 2. INSTALLATION

The installation phase requires the following stages:

- Prerequisites: Installing and configuring the Red Hat environment
- Creating the RHV conversion host(s)
- Configuring the RHV conversion host(s)
- Configuring the maximum number of concurrent migrations (optional, only supported for SSH transformation)

2.1. PREREQUISITES

Before migration, your environment must have the following prerequisites:

Target datastore size

The target datastore must have sufficient space for the virtual machines that will be migrated.

BIOS settings of hosts

The BIOS settings of the physical hosts should be set for optimal performance, rather than power-saving, according to the vendor's recommendations. Where applicable, the C1E halt state should be disabled.

Red Hat Virtualization Manager 4.2.4 or later installed

See the [Red Hat Virtualization Planning and Prerequisites Guide](#) and [Installing the Red Hat Virtualization Manager](#).

Red Hat Virtualization Hosts 4.2.4 or RHEL 7.5 hosts or later installed

See [Installing Red Hat Virtualization Host](#) or [Installing Red Hat Enterprise Linux Hosts](#).

Network

The VMware network must be extended to the RHV 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.

Ports

The following ports must be enabled:

- 22 - SSH
- 443 - CloudForms, Red Hat Virtualization Manager, and VDDK
- 902 - CloudForms to VMware
- 5480 - **virt-v2v** to vCenter

For more information, see [Ports used by Red Hat CloudForms Management Engine 5.1 and above](#).

Data and ISO storage domains attached to the RHV data center

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.

VirtIO and Guest Tool images uploaded to the ISO storage domain

See [Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain](#)

**NOTE**

The VirtIO image file name must include the version number: `virtio-win-version.iso`.

Red Hat CloudForms 4.6.4 or later installed

See [Installing Red Hat CloudForms on Red Hat Virtualization](#).

VMware and RHV environments added to CloudForms

See [VMware vCenter Providers](#) and [Red Hat Virtualization Providers](#).

RHV hosts authenticated

See [Authenticating Red Hat Virtualization Hosts](#).

VMware hosts authenticated

See [Authenticating VMware vCenter Hosts](#).


RHV conversion hosts created and configured

See [Section 2.2, “RHV Conversion Hosts”](#).

2.2. RHV CONVERSION HOSTS

All the virtual machines that are specified in a migration plan are migrated at the same time. Therefore, the number of RHV conversion hosts that you create depends on the number of virtual machines you intend to migrate. Red Hat recommends configuring the RHV conversion hosts on physical machines, not nested virtual machines.

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 **10**, for VDDK transformation, and **20**, for SSH transformation. The SSH transformation limitation applies only to Red Hat Virtualization 4.2.5, not to later versions.

Multiple conversion hosts are recommended, even for smaller migrations, for load-balancing purposes. The number of virtual machines that you can migrate at the same time 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 test your environment to understand how many migrations can occur simultaneously without negative effects.

Datapath Transformation Options: VDDK or SSH

Select the best datapath transformation option for your environment, either VDDK or SSH:

VDDK

- Recommended because it is approximately twice as fast as SSH
- Compiled with **nbdkit** to provide direct connectivity to the virtual machine disks
- VDDK package must be downloaded
- Conversion host configuration requires a location
- Each conversion host supports only 10 concurrent migrations

SSH

- Does not require additional downloads, packages, tools, or compilation
- RHV and VMware environments must share SSH keys
- SSH access must be enabled on ESXi hosts
- Slower than VDDK
- Suitable for larger scale migrations than VDDK

2.2.1. Creating and Configuring RHV Conversion Hosts for VDDK Transformation

Create and configure RHV conversion hosts for VDDK transformation:

1. Log in to the [VMware VDDK Documentation](#) site and download the latest VDDK package (for example, **VMware-vix-disklib-6.5.3-8315684.x86_64.tar.gz**) to your local machine.
2. Using SCP, copy the VDDK package to **/var/www/html** on the Manager machine.
3. Go to the Ansible conversion hosts playbooks directory on the Manager machine:

```
# cd /usr/share/ovirt-ansible-v2v-conversion-host/playbooks/
```

4. Create a **conversion_hosts_inventory.yml** inventory file with the following variables:

```
all:
  hosts:
    host1.example.com:
    host2.example.com:
  vars:
    ansible_ssh_private_key_file: /etc/pki/ovirt-
engine/keys/engine_id_rsa
    v2v_vddk_package_name: VMware-vix-disklib-version.x86_64.tar.gz
    v2v_vddk_package_url: http://Manager_FQDN/VMware-vix-
disklib-version.x86_64.tar.gz
```

5. Enable the conversion hosts:


```
# ansible-playbook -i conversion_hosts_inventory.yml
conversion_host_enable.yml
```



NOTE

The Ansible role checks for the presence of the RHV guest tools ISO image in the ISO storage domain. This image is used to migrate Windows virtual machines.

6. Verify that the conversion hosts are enabled:

```
# ansible-playbook -i conversion_hosts_inventory.yml
conversion_host_check.yml
```



NOTE

If you upgrade your RHV environment, you must run the **conversion_host_enable** playbook again, with the option **-e v2v_vddk_override=true**, to update the IMS packages.

After creating the RHV conversion hosts, you must configure them in CloudForms. See [Section 2.2.3, “Configuring RHV Conversion Hosts in CloudForms”](#).

2.2.2. Creating and Configuring RHV Conversion Hosts for SSH Transformation

Create and configure RHV conversion hosts for SSH transformation:

1. Go to the Ansible conversion hosts playbooks directory on the Manager machine:

```
# cd /usr/share/ovirt-ansible-v2v-conversion-host/playbooks/
```

2. Create a **conversion_hosts_inventory.yml** inventory file with the following variables:

```
all:
  hosts:
    host1.example.com:
    host2.example.com:
  vars:
    ansible_ssh_private_key_file: /etc/pki/ovirt-
engine/keys/engine_id_rsa
```

3. Enable the conversion hosts:

```
# ansible-playbook -i conversion_hosts_inventory.yml
conversion_host_enable.yml
```



NOTE

The Ansible role checks for the presence of the RHV guest tools ISO image in the ISO storage domain. This image is used to migrate Windows virtual machines.

4. Verify that the conversion hosts are enabled:


```
# ansible-playbook -i conversion_hosts_inventory.yml
conversion_host_check.yml
```

5. Enable SSH on all the VMware hypervisors. See [Enable SSH from the vSphere Web Client](#) for details.



NOTE

If you created the RHV conversion hosts manually, without using the **ovirt-ansible-v2v-conversion-host** role, you must create the key pairs on each conversion host for the **vdsm** user:

```
# install -o vdsd -g kvm -m 0700 -d /var/lib/vdsd/.ssh
# ssh-keygen -t rsa -N '' -C "vdsd@hostname" -f
/var/lib/vdsd/.ssh/id_rsa
# chown vdsd:kvm /var/lib/vdsd/.ssh/id_rsa
```

6. Copy the SSH keys to the VMware hypervisors:

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



NOTE

This step ensures that each conversion host has the SSH key of the ESXi host in its **known_hosts** file.

7. Validate the configuration by connecting to the VMware hypervisor using **ssh-agent**, the program that **virt-v2v** uses to connect to the hypervisor:

```
# sudo -u vdsd ssh-agent
SSH_AUTH_SOCKET=/tmp/ssh-Gi2oSn44DHNL/agent.65904; export
SSH_AUTH_SOCKET;
SSH_AGENT_PID=65905; export SSH_AGENT_PID;
echo Agent pid 65905;

# sudo -u vdsd SSH_AUTH_SOCKET=/tmp/ssh-Gi2oSn44DHNL/agent.65904 ssh-
add
# sudo -u vdsd \
    SSH_AUTH_SOCKET=/tmp/ssh-Gi2oSn44DHNL/agent.65904 ssh
root@esx1.example.com
```

If the connection is successful, the RHV conversion host is correctly configured for SSH transformation.



NOTE

If you are using SSSD with single sign-on, see [Troubleshooting](#) for additional information.

After creating the RHV conversion hosts, you must configure them in CloudForms. See [Section 2.2.3, “Configuring RHV Conversion Hosts in CloudForms”](#).

Additional resources

- For more information, see [V2V - Conversion Host - Ansible](#).

2.2.3. Configuring RHV Conversion Hosts in CloudForms

Configure the RHV conversion hosts in CloudForms:

1. Go to **Compute** → **Infrastructure** → **Hosts** and select a Red Hat Virtualization Host.
2. Click the **Policy** drop-down button and select **Edit Tags**.
3. Select **V2V - Transformation Host** from the tag drop-down list and **t** as the assigned value.
4. Select **V2V - Transformation Method** from the tag drop-down list and either **VDDK** or **SSH** as the assigned value. See [Datapath Transformation Options: VDDK vs. SSH](#) for more information.
5. Click **Save**.
6. Select the same host, click the **Configuration** drop-down button, and select **Edit this item**.
7. In the **Default** tab of the Endpoints section, enter the **Username root** and the password.
8. Click **Validate**. If a success message is displayed, click **Save**.

Conversion Host Properties

The screenshot displays the Red Hat CloudForms Management Engine interface. The left sidebar shows the navigation menu with options like Cloud Intel, Red Hat Insights, Services, Compute, Configuration, Networks, Storage, Control, Automation, Optimize, and Monitor. The main content area shows the configuration for a specific host, identified as 'rhel 7.5'. The 'Properties' section lists various attributes such as Hostname, IP Address, IPMI IP Address, Hypervisor Hostname, VMM Information, Manufacturer / Model, Operating System, Power State, Lockdown Mode, Maintenance Mode, Devices, Network, Storage Adapters, Number of CPUs, Number of CPU Cores, CPU Cores Per Socket, Memory, and Management Engine GUID. The 'Compliance' section shows Status (Never Verified) and History (Not Available). The 'Security' section shows Users (0), Groups (0), Patches (0), Firewall Rules (0), and SSH Root. The 'Configuration' section shows Packages (0), Services (0), Files (0), and Advanced Settings (0). The 'Smart Management' section shows My Company Tags with V2V - Transformation Host: t and V2V - Transformation Method: VDDK.

2.2.4. Configuring the Maximum Number of Concurrent Migrations

The default limit for concurrent migrations is **10** for a conversion host and **10** for a provider. You can change this value by setting the **Max Transformation Runners** attribute for the conversion host or for the provider, or both. The limit that is set for the provider has priority over the limit set for the conversion host. For example, if you set the number of concurrent migrations to **20** at the provider level, while having **3** conversion hosts, the number of concurrent migrations is limited to **20** because the provider's setting overrides the conversion hosts' setting.



IMPORTANT

The maximum number of concurrent virtual machine migrations per conversion host is **10**, for VDDK transformation, and **20**, for SSH transformation. The SSH transformation limitation applies only to Red Hat Virtualization 4.2.5, not to later versions.

The following examples show how to set **Max Transformation Runners** using the Rails console on the CloudForms appliance.

To set **Max Transformation Runners** for the conversion host *host1.example.com* to **20**:

```
# 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)
```

To set **Max Transformation Runners** for the provider *My RHV* to **30**:

```
# 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 => "My
RHV").custom_set("Max Transformation Runners", 30)
```

The following examples show how to get the current configuration of **Max Transformation Runners** using the **custom_get** method. This method returns either the current configuration value or **nil**, if the value has not been set.

To get the configuration of **Max Transformation Runners** for conversion host *host1.example.com*:

```
# 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_get("Max Transformation Runners")
```

To get the configuration of **Max Transformation Runners** for provider *My RHV*:

```
# 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 => "My
RHV").custom_get("Max Transformation Runners")
```


CHAPTER 3. DISCOVERY AND ASSESSMENT

The discovery and assessment phase is the foundation of the planning phase of your migration. At the end of this 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." To help you formulate a migration goal, you must discover and assess your current environment.

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

The power state of the migrated virtual machines is preserved: a virtual machine that is powered on or off before migration will be in the same power state after migration.

3.1. DISCOVERY QUESTIONS

The following questions will guide you through the discovery process:

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 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. 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 RHV, running these guests is not officially supported.

What resources will my target environment require?

Your target environment must be large enough to hold the migrated virtual machines, in addition to the original VMware virtual machines. The VMware virtual machines are not affected by the migration process. Their power state after the migration will be the same as before.

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. It may be possible to do a phased migration of your applications, if they are distributed in a high-availability architecture, without downtime at the application layer.

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.

3.2. MIGRATION EXAMPLE

This migration example can help you to estimate how long your migration will take.

This environment took 1 hour and 30 minutes to migrate:

- 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

3.3. BEST PRACTICES

- Run the [IMS RHEL pre-migration Ansible playbook](#) to preserve static IP address configuration and to install the RHV guest agent. See [Section 4.3, “Advanced Migration Option: Automating Pre- and Post-Migration Tasks with Ansible”](#) for more information.
- 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.
- Reboot or delete snapshots before the migration, to reduce the risk of failure.
- 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.
- Prepare your users for downtime.

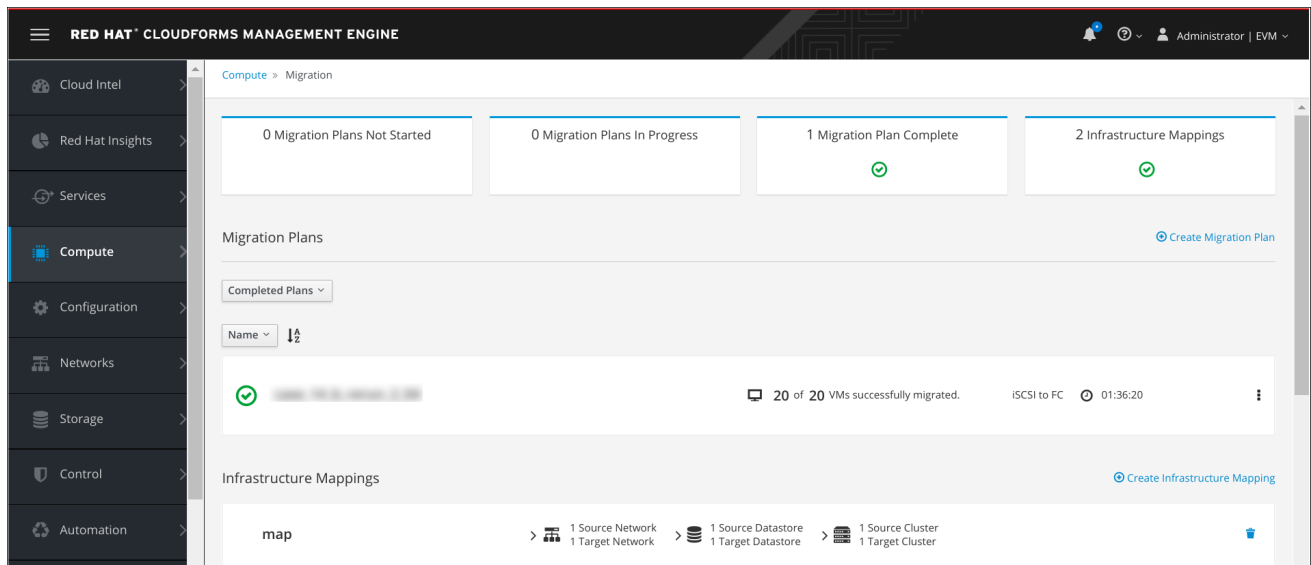
CHAPTER 4. PLANNING

The planning phase of the migration has 2 key stages:

- Creating an infrastructure mapping
- Creating and running a migration plan

4.1. CREATING AN INFRASTRUCTURE MAPPING

Migration Screen



In CloudForms, create an infrastructure mapping to map your VMware resources to RHV resources:

1. Go to **Compute** → **Migration** and click **Create Infrastructure Mapping**.
2. Enter the name of your mapping and click **Next**.
3. Select a source cluster and a target cluster.

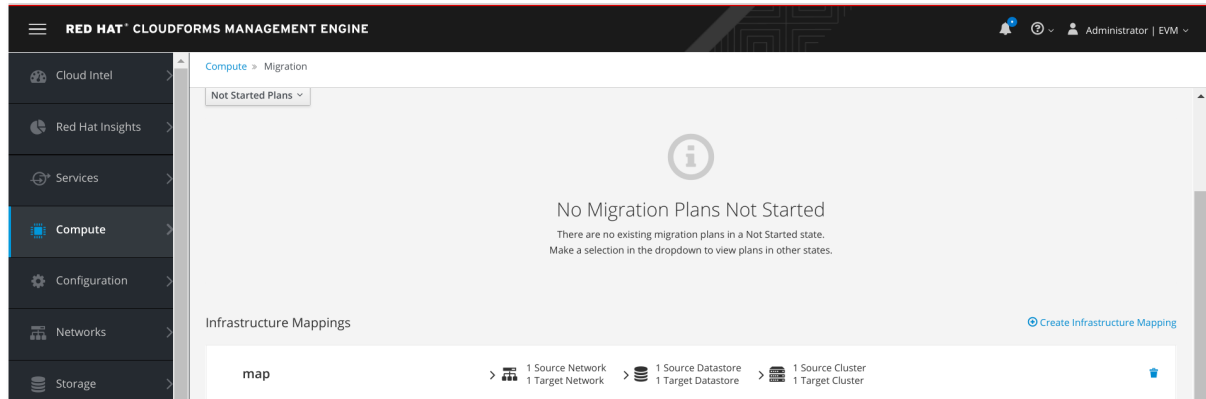


NOTE

A target cluster that does not contain a conversion host is marked with a warning icon (⚠), which appears after a few seconds' delay for polling the clusters and refreshing the UI.

4. Click **Add Mapping** and **Next**.
5. Map the datastore of each cluster:
 - a. Select the target cluster from the drop-down list.
 - b. Select a source datastore and a target datastore.
 - c. Click **Add Mapping** and **Next**.
6. Map the network of each cluster:
 - a. Select the target cluster from the drop-down list.

- b. Select a source network and a target network.
 - c. Click **Add Mapping**.
7. Click **Create**.
 8. If a successful mapping message is displayed, click **Close**.
Your infrastructure mapping is displayed in **Compute** → **Migration**.



After creating an infrastructure mapping, create a migration plan.

4.2. CREATING AND RUNNING A MIGRATION PLAN

Prerequisites

- If you are migrating a large number of virtual machines or you are re-migrating virtual machines after a failed migration, create a CSV file containing a list of source virtual machines to migrate. The CSV file specifies the virtual machine names (mandatory) and providers (optional, for differentiating virtual machines with identical names), in the following format:

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

If you are migrating a small number of virtual machines, you can use the migration plan to discover the virtual machines. The migration plan displays a list, from which you manually select the virtual machines to migrate.

- Check the target datastore to ensure that there are no virtual machines with the same name as a source virtual machine. A name conflict will cause the migration to fail.
- Unmount any ISO images that are mounted on source virtual machines. A virtual machine with an attached ISO image will cause the migration to fail.
- Check that the source virtual machine disks are unencrypted. The current IMS version does not support the migration of encrypted disks.

Procedure

- Go to **Compute** → **Migration** and click **Create Migration Plan**.
- Select the infrastructure mapping from the drop-down list.

3. Enter the name of your migration plan.
4. Discover the virtual machines:
 - If you are using the migration plan to discover the virtual machines automatically:
 - a. Select **Choose from a list of VMs discovered in the selected infrastructure mapping**.
 - b. Click **Next**.

Discovered Virtual Machines

Migration Plan Wizard

General VMs Advanced Options Schedule Results

1 2 3 4 5

VMs to be Migrated

VM Name Filter by VM

<input type="checkbox"/>	<input checked="" type="checkbox"/>	VM Name ^	Source Cluster	Path	Allocated Size
<input type="checkbox"/>	<input checked="" type="checkbox"/>	vmware-Datacenter	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	vmware-Datacenter	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	vmware-Datacenter	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	vmware-Datacenter	Cluster	vmware/Datacenter	536.87GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	vmware-Datacenter	Cluster	vmware/Datacenter	26.84GB

5 per page 1-5 of 10 1 of 2

Cancel < Back Next >

- If you are importing a CSV file:
 - a. (Optional) To validate your infrastructure mapping, select **Choose from a list of VMs discovered in the selected infrastructure mapping** and click **Next**. If your virtual machines appear in the list, the infrastructure mapping is correct. Click **Back**.
 - b. Select **Import a CSV file with a list of VMs to be migrated** and click **Next**.
 - c. Click **Import**, browse to the CSV file, and click **Open**.
- 5. Select the virtual machines to migrate and click **Next** to go to **Advanced Options**.

Advanced Options

Migration Plan Wizard

General VMs Advanced Options Schedule Results

1 2 3 4 5

Select pre-migration playbook service (optional)

<Choose>

Select post-migration playbook service (optional)

<Choose>

Select VMs on which to run the playbook services

VM Name Filter by VM

Pre-Migration Service	Post-Migration Service	VM Name	Source Cluster
<input type="checkbox"/>	<input type="checkbox"/>		
<input type="checkbox"/>	<input type="checkbox"/>		Cluster

Cancel < Back Next >

Select an Ansible playbook service catalog item if you are automating pre- or post-migration tasks.

6. Select **Start Migration Immediately** and click **Create**.
7. Click **Close**. The migration plan may take some time to complete. **In Progress Plans**, the amount of data being transferred, the number of virtual machines being migrated, and the elapsed time are displayed. To view the migration progress of individual virtual machines, click the migration plan to go to the migration plan details.
When the migration is complete, the migration plan is displayed under **Completed Migrations**, with the number of successfully migrated virtual machines. A **Retry** option enables you to re-migrate any virtual machines in this plan that failed to migrate. See [Retrying a Migration Plan](#).



NOTE

During migration, the counter displaying the migration plan progress (in **Compute** → **Migration**) may be a few seconds ahead of the counter that appears in the migration plan details (**Compute** → **Migration** → **Migration Plan**). This is because the migration plan counter displays the total time to run the migration plan, while the migration plan details counter displays the time to migrate the virtual machines.

4.3. ADVANCED MIGRATION OPTION: AUTOMATING PRE- AND POST-MIGRATION TASKS WITH ANSIBLE

Red Hat Infrastructure Migration Solution enables you to use Ansible playbooks to automate tasks on virtual machines before or after migration. For example, use a playbook to remove web server virtual machines from a load-balancing pool before migration and adding them back to the pool after migration. You create an Ansible service catalog item that contains the playbook and select it while creating the infrastructure migration plan.

Use the [IMS RHEL Pre-migration Ansible Playbook Example](#) to configure the target Red Hat Enterprise Linux virtual machines.

Creating an Ansible Service Catalog Item

The CloudForms inventory must contain at least one Ansible repository, one playbook, and one credential.

Procedure

1. Configure the [Ansible repository and credentials](#).
2. Create a [catalog bundle](#).
3. Create an [Ansible playbook service catalog item](#).

IMS RHEL Pre-migration Ansible Playbook Example

The `ims.rhel_pre_migration` role configures a target Red Hat Enterprise Linux virtual machine for migration by performing the following tasks:

- Preserving the static IP address configuration
- Installing the Red Hat Virtualization guest agent

This example registers a virtual machine with a Satellite 6 server and enables the **rhel-7-server-rpms** repository. If no variables are provided, the [playbook](#) assumes that the repositories are already configured and installs the guest agent while maintaining the static IP address.

```
---
- hosts: all
  vars:
    rhsm_config:
      server_hostname: "satellite.example.com"
      server_username: "admin"
      server_password: "password"
      org_id: "organization"
    rhsm_repositories:
      - "rhel-7-server-rpms"
  roles:
    - role: ims.rhel_pre_migration
```

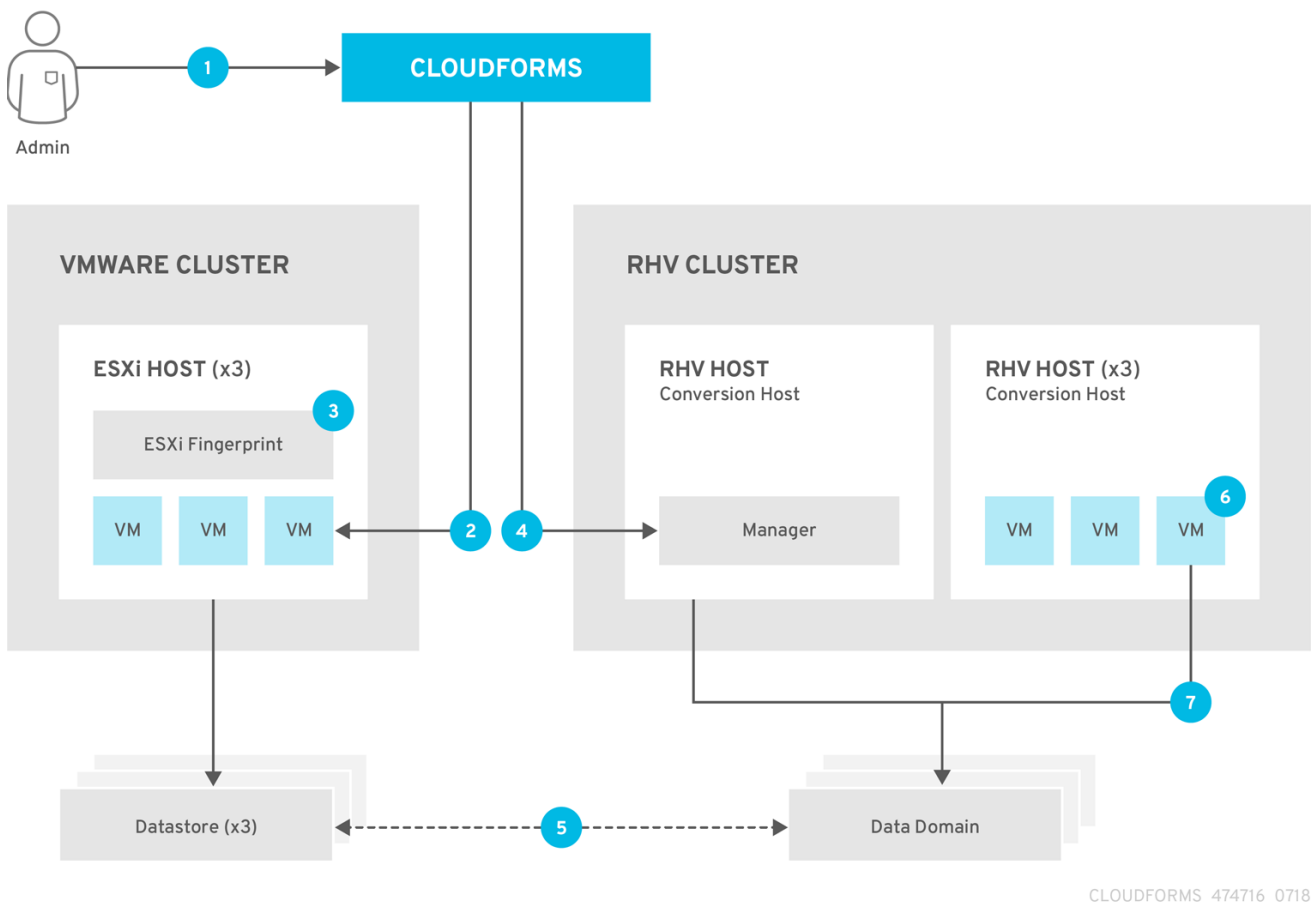

CHAPTER 5. MIGRATION

This section describes the actual migration process, while you are running the migration plan.

5.1. WORKFLOW

The diagram below provides a high-level view of the infrastructure migration workflow. An understanding of this workflow will greatly simplify troubleshooting.

Migration Workflow



1. The Infrastructure Admin creates an infrastructure mapping and a virtual machine migration plan in CloudForms and runs the migration plan.
2. CloudForms locates the virtual machines to be migrated based on the infrastructure mapping.
3. The ESXi host fingerprint is captured for authentication during the conversion process if the VDDK transport method is used. If SSH is used, a shared SSH key is used to connect to the ESX host where the virtual machine resides.
4. Using the RHV attributes for the target environment, CloudForms initiates communication with the RHV conversion host.
5. The RHV conversion host connects to the source datastore through the ESX host, using **virt-v2v-wrapper.py**, and streams the disk to be converted to the target data domain chosen in the infrastructure mapping using **virt-v2v**.
6. After the disk is converted, the target virtual machine is created in RHV. During creation, 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. After the virtual machine is created, the disk is attached to the target virtual machine.

The virtual machine migration is complete and its status is displayed in CloudForms.



NOTE


Depending on the source virtual machine's power state before migration, the virtual machine is either left powered off after migration or is powered on.

5.2. TROUBLESHOOTING

Remigrating Virtual Machines

You can remigrate previously migrated virtual machines:

1. Delete the target virtual machines corresponding to the source virtual machines that you are remigrating. Otherwise, the remigration will fail because a target machine with the same name already exists.
2. Delete the disks belonging to these virtual machines, which were created in the target datastore during the earlier migration, to free up space.
3. Create a [CSV file](#) with the names and, optionally, the providers, of the source virtual machines.



NOTE

The virtual machines cannot be discovered automatically by the migration plan because they are marked in the CloudForms VMDB as **migrated**.

4. Create and run a new migration plan, using the CSV file to import the virtual machines.

Retrying a Migration Plan

If one or more virtual machines in a migration plan fails to migrate, the plan's migration status is **Failed**. You can retry a migration plan to migrate the virtual machines that failed to migrate:

1. Click the migration plan to go to the details page to identify the virtual machines that failed to migrate.
2. Delete the target virtual machines corresponding to the source virtual machines that failed to migrate. Otherwise, the retry will fail because a target machine with the same name already exists.
3. Go to **Compute** → **Migration**, select **Failed Migrations**, select the failed migration plan, and click **Retry**.

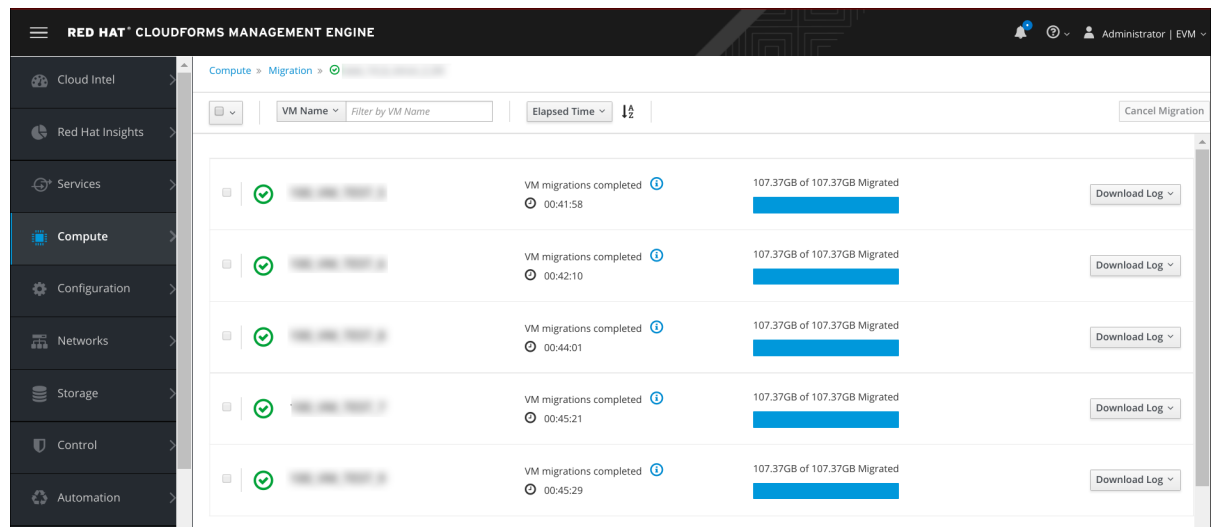
Log Files

If a migration plan does not complete successfully, you can check the logs for errors:

- CloudForms migration log: `/var/www/miq/vmdb/log/automation.log`
- Conversion host logs: `/var/log/vdsm/import/`

- Virtual machine migration log: Go to **Compute** → **Migrate**, click the migration plan to go to the details page, and click **Download Log**.

Migrated Virtual Machine Logs



Common Issues and Mistakes

Incorrect infrastructure mapping

If you create an infrastructure mapping and subsequently add or remove providers or provider objects, the infrastructure mapping will not work because the object IDs of the providers and their objects have changed. You must delete the infrastructure mapping and re-create it.

RHV host credentials not set in CloudForms

If the RHV host credentials are not set, CloudForms cannot communicate with the hosts. To authenticate the RHV hosts with CloudForms, see [Authenticating Red Hat Virtualization Hosts](#) in *CloudForms: Managing Providers*.

Incorrect attributes in the virtual machine CSV import file

Create a new migration plan, import the correct CSV file, and run the migration plan.

Migration fails immediately

If a migration fails immediately, before the progress bar is displayed, this is normally caused by a configuration error in the environment. Check the **automation.log** file on the CloudForms appliance.

Migration fails at virtual machine conversion

If a migration fails when the virtual machines are being converted (the progress bar updates), click the migration plan to go to the details page and download the conversion log for the the virtual machine(s) that failed to migrate.

SSH transformation fails

If you are using SSSD with single sign-on for RHV, SSH fails for the **vdsm** account. Reinstall **ipa-client** without configuring the OpenSSH client:

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

Changing the configuration file is not recommended because it is restored during upgrades. See https://bugzilla.redhat.com/show_bug.cgi?id=1544379 for details.

