



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 © 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 migrate your virtual machines from VMware 5.5 to Red Hat Virtualization 4.2.7.

Table of Contents

CHAPTER 1. INFRASTRUCTURE MIGRATION SOLUTION OVERVIEW	3
CHAPTER 2. PLANNING THE MIGRATION	5
2.1. PLANNING QUESTIONS	5
2.2. BEST PRACTICES FOR MIGRATION	6
CHAPTER 3. PREPARING THE MIGRATION ENVIRONMENT	8
3.1. PREPARING THE SOURCE AND TARGET ENVIRONMENTS	8
3.2. PREPARING THE RHV CONVERSION HOSTS	9
3.2.1. Creating and Configuring RHV Conversion Hosts for VDDK Transformation	10
3.2.2. Creating and Configuring RHV Conversion Hosts for SSH Transformation	11
3.2.3. Configuring RHV Conversion Hosts in CloudForms	12
3.2.4. Configuring the Maximum Number of Concurrent Migrations	13
CHAPTER 4. MIGRATING THE INFRASTRUCTURE	15
4.1. CREATING AN INFRASTRUCTURE MAPPING	15
4.2. CREATING AND RUNNING A MIGRATION PLAN	16
4.3. CANCELING A MIGRATION PLAN	20
4.4. ADVANCED OPTION: AUTOMATING PRE- AND POST-MIGRATION TASKS WITH ANSIBLE	20
CHAPTER 5. TROUBLESHOOTING	22
5.1. RETRYING A MIGRATION PLAN	22
5.2. LOGS	22
5.3. COMMON ISSUES AND MISTAKES	23

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.

[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 IMS user interface, [Red Hat CloudForms](#), provides the control and automation required to manage challenging enterprise environments.

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

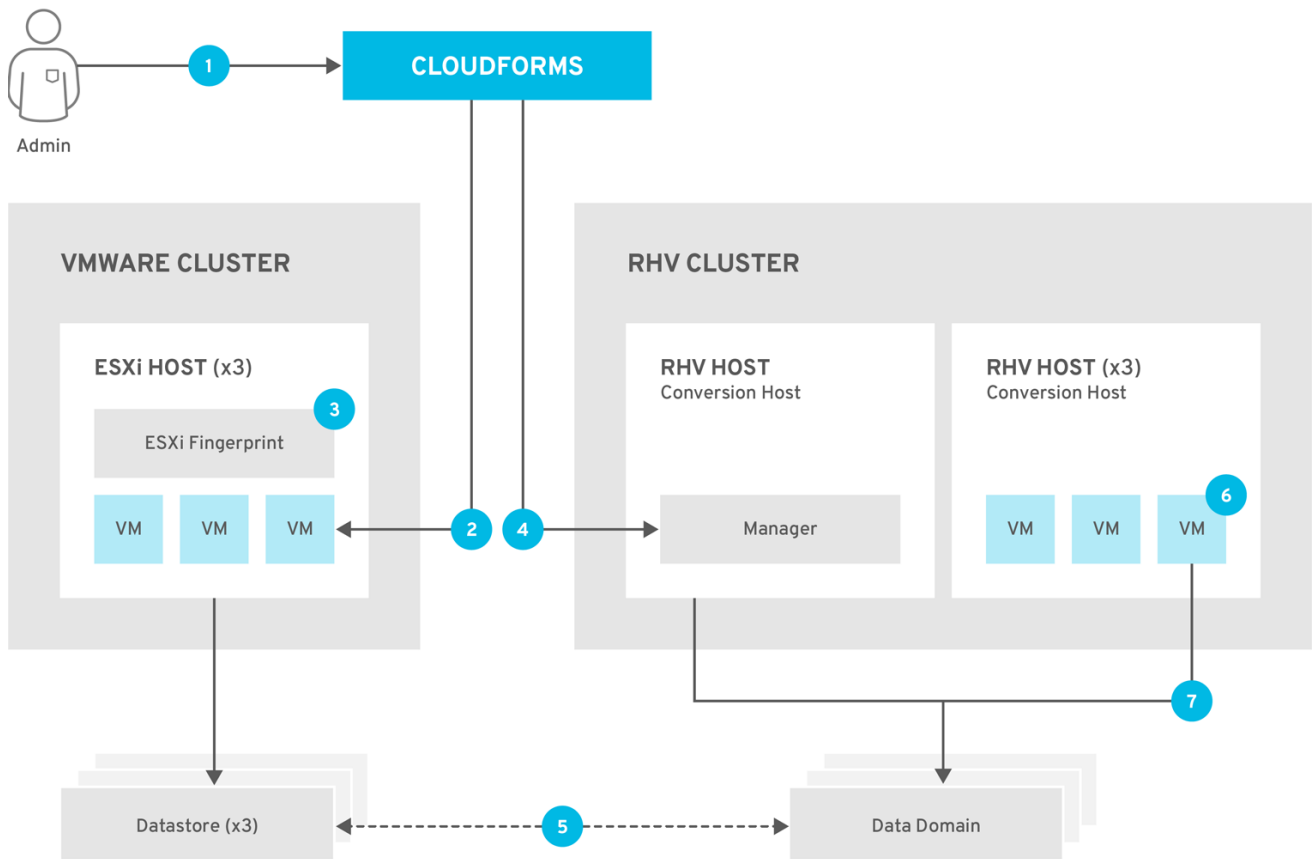
IMS Phases

The Infrastructure Migration Solution has 4 key phases:

1. Planning the migration:
 - a. Consult the [planning questions](#) to ensure that you have sufficient resources for the migration.
 - b. Follow [best practices](#) to minimize the impact on your users. You may choose to schedule the migration in phases or to create groups of virtual machines that should migrate together.
2. Preparing the migration environment:
 - a. Create and configure the [source and target environments](#).
 - b. Create and configure the [RHV conversion host\(s\)](#).
3. Migrating the infrastructure:
 - a. [Map](#) your source environment to the target environment, selecting the virtual machines to migrate.
 - b. Create and run a [migration plan](#), with optional automated pre- and post-migration tasks.
4. Verifying the migration:
 - a. Check the completed migration plan for successfully migrated virtual machines.
 - b. If errors appear, see [Chapter 5, Troubleshooting](#).

Workflow: VMware to Red Hat Virtualization

This workflow occurs when you run a migration plan to migrate your infrastructure from VMware to Red Hat Virtualization.



CLOUDFORMS_474716_0718

1. After creating an infrastructure mapping and a migration plan in CloudForms, you run 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.

CHAPTER 2. PLANNING THE MIGRATION

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

The [planning questions](#) provide guidance for 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

These questions will ensure that you have sufficient resources to perform the migration:

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. 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

- Run the [IMS RHEL pre-migration Ansible playbook](#) to preserve static IP address configuration and to install the RHV guest agent. See [Section 4.4, “Advanced 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 3. PREPARING THE MIGRATION ENVIRONMENT

Preparing your environment for migration comprises the following tasks:

- Configuring your current infrastructure as the [source environment](#)
- Creating and configuring Red Hat Virtualization as the [target environment](#)
- Creating and configuring the [RHV conversion host\(s\)](#)

3.1. PREPARING THE SOURCE AND TARGET ENVIRONMENTS

Before migration, you must prepare your source and target environments:

1. Plan and set up the [RHV target environment](#).
2. Ensure that the target datastore has sufficient space for the virtual machines you are migrating.
3. 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.
4. Install [Red Hat Virtualization Manager 4.2.4](#) or later.
5. Install [Red Hat Virtualization Hosts 4.2.4](#) or later and [RHEL 7.5 hosts](#).
6. Extend the VMware network 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.

7. Enable the following [ports](#) in the target environment:
 - 22 - SSH
 - 443 - CloudForms, Red Hat Virtualization Manager, and VDDK
 - 902 - CloudForms to VMware
 - 5480 - `virt-v2v` to vCenter
8. Attach the [Data and ISO storage domains](#) to the target RHV data center.



NOTE

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

9. 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`.
10. Install [Red Hat CloudForms 4.6.4](#) or later.

11. Add the [source](#) and [target](#) environments to CloudForms.
12. Validate the source virtual machine names. Virtual machine names may contain only upper- or lower-case letters, numbers, underscores (`_`), hyphens (`-`), or periods (`.`). Special characters and spaces are not permitted.
13. Check that no source virtual machine has the same name as a virtual machine in the target environment. A name conflict will cause the migration to fail.
14. Unmount ISO images from source virtual machines. A virtual machine with an attached ISO image will cause the migration to fail.
15. Authenticate the [RHV hosts](#).
16. Create and configure the [RHV conversion hosts](#).

3.2. PREPARING THE RHV CONVERSION HOSTS

Before you create the RHV conversion hosts, you must decide [how many conversion hosts](#) you need and which datapath transformation option, [VDDK or SSH](#), to use.

Number of RHV Conversion Hosts

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

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.

You must configure 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, however, applies only to RHV 4.2.5, not to later versions of RHV.

Datapath Transformation Options: VDDK or SSH

You must choose the datapath transformation option, VDDK or SSH, that is appropriate for your migration size and setup:

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
- 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
- Better suited for large-scale migrations than VDDK

3.2.1. Creating and Configuring RHV Conversion Hosts for VDDK Transformation

1. Go to [VMware Documentation](#), 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 SDK package (for example, **VMware-vix-disklib-6.5.3-8315684.x86_64.tar.gz**) to your local machine.
4. Using SCP, copy the VDDK package to **/var/www/html** on the Manager machine.
5. On the Manager machine, create **/usr/share/ovirt-ansible-v2v-conversion-host-version/playbooks/conversion_hosts_inventory.yml** 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
```

6. Enable the conversion hosts:

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



NOTE

The **ovirt-ansible-v2v-conversion-host** 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.

7. 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 can [configure them](#) in CloudForms.

3.2.2. Creating and Configuring RHV Conversion Hosts for SSH Transformation

1. On the Manager machine, create `/usr/share/ovirt-ansible-v2v-conversion-host-version/playbooks/conversion_hosts_inventory.yml` 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
```

2. Enable the conversion hosts:

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



NOTE

The **ovirt-ansible-v2v-conversion-host** 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.

3. Verify that the conversion hosts are enabled:

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

4. Enable SSH on your VMware hypervisors. (For details, see [VMware vSphere Documentation](#). In the left 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.**)
5. (Optional) If you created the RHV conversion hosts manually, without using the **ovirt-ansible-v2v-conversion-host** role, create the key pairs on each conversion host for the **vds** user:

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

```
/var/lib/vdsm/.ssh/id_rsa
# chown vds:m /var/lib/vdsm/.ssh/id_rsa
```

- Copy the SSH keys to all VMware hypervisors:

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

- 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 vds:m 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 vds:m SSH_AUTH_SOCKET=/tmp/ssh-Gi2oSn44DHNL/agent.65904 ssh-
add
# sudo -u vds:m \
    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 can [configure them](#) in CloudForms.

3.2.3. Configuring RHV Conversion Hosts in CloudForms

- Click **Compute** → **Infrastructure** → **Hosts** and select a Red Hat Virtualization Host.
- Click the **Policy** drop-down button and select **Edit Tags**.
- Select **V2V - Transformation Host** from the tag drop-down list and **t** as the assigned value.
- Select **V2V - Transformation Method** from the tag drop-down list and either **VDDK** or **SSH** as the assigned value.
- Click **Save**.
- 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 navigation options like Cloud Intel, Red Hat Insights, Services, Compute, Configuration, Networks, Storage, Control, Automation, Optimize, and Monitor. The main content area shows the 'rhel 7.5' host details. The 'Properties' section includes fields for Hostname, IP Address (10.12.69.31), IPMI IP Address, Hypervisor Hostname, VMM Information (RedHat), Manufacturer / Model, Operating System (RHEL), Power State (on), Lockdown Mode (Disabled), Maintenance Mode (Disabled), Devices (376), Network (Available), Storage Adapters (0), Number of CPUs (2), Number of CPU Cores (12), CPU Cores Per Socket (6), Memory (126 GB), 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 (V2V - Transformation Host: t, V2V - Transformation Method: VDDK).

3.2.4. 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 RHV 4.2.5, not to later versions of RHV.

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> $vm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $vm.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> $vm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $vm.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> $vm =
MiqAeMethodService::MiqAeService.new(MiqAeEngine::MiqAeWorkspaceRuntime.ne
w)
irb(main):002:0> $vm.vmdb(:ext_management_system).find_by(:name => "My
RHV").custom_get("Max Transformation Runners")
```

CHAPTER 4. MIGRATING THE INFRASTRUCTURE


Migrating your infrastructure comprises 2 tasks:

- [Mapping](#) the resources (cluster, datastore, and networks) of your source environment to your target environment and selecting virtual machines to migrate
- Creating a [migration plan](#) with optional [automated tasks](#) that run before or after each virtual machine is migrated
You can [cancel](#) a migration plan that is in progress. This option allows you to retry the migration plan immediately or to schedule it in the future.

4.1. CREATING AN INFRASTRUCTURE MAPPING

The infrastructure mapping maps your VMware resources to RHV resources.

Migration Screen




NOTE

Before creating an infrastructure mapping, you must [authenticate](#) the RHV conversion hosts. If the hosts are not authenticated, the mapping will display a [missing networks error](#).

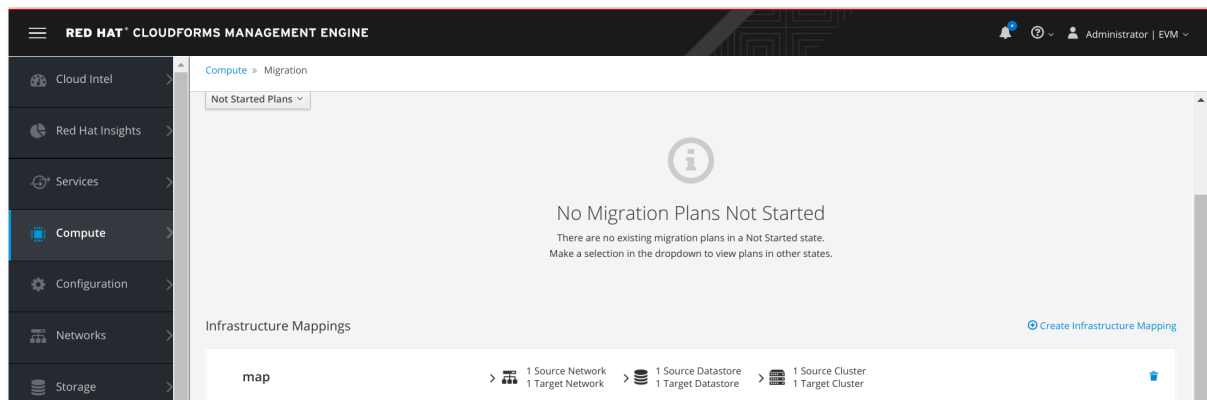
If you add or remove providers or provider objects from an infrastructure mapping, the mapping will display [missing resource errors](#).

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

**NOTE**

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

5. Click **Add Mapping** and **Next**.
6. 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**.
7. 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**.
8. Click **Create**.
9. If a successful mapping message is displayed, click **Close**. If error messages appear, see [Chapter 5, Troubleshooting](#).
Your infrastructure mapping is displayed in **Compute** → **Migration**.



After you have created a successful infrastructure mapping, you are ready to create and run a [migration plan](#).


4.2. CREATING AND RUNNING A MIGRATION PLAN

Prerequisites

- Create a CSV file listing the source virtual machines if you are migrating a large number of virtual machines or if you are [remigrating previously migrated virtual machines](#). The CSV file contains the virtual machine names (mandatory) and providers (optional, to differentiate virtual machines with identical names) in the following format:

Name, Provider

```
vm01, vSphere3  
vm02, vSphere3  
vm03, vSphere3  
...
```

**NOTE**

If you are migrating a small number of virtual machines, you do not need to create a CSV file. The migration plan can discover and display a list from which you manually select the virtual machines to migrate.

- Check the target datastore to ensure that no virtual machine has the same name as a source virtual machine.
- Unmount ISO images from source virtual machines.

**NOTE**

A virtual machine will be in the same power state (**on** or **off**) after migration as it was before migration.

Procedure

1. Click **Compute** → **Migration**.
2. Click **Create Migration Plan**.
3. Select the infrastructure mapping from the drop-down list.
4. Enter the name of your migration plan.
5. Discover the virtual machines:
 - If you are migrating a small number of virtual machines, select **Choose from a list of VMs discovered in the selected infrastructure mapping** and 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"/>	[REDACTED]	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	[REDACTED]	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	[REDACTED]	Cluster	vmware/Datacenter	107.37GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	[REDACTED]	Cluster	vmware/Datacenter	536.87GB
<input type="checkbox"/>	<input checked="" type="checkbox"/>	[REDACTED]	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) You can 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**.
 1. Select the virtual machines to migrate and click **Next** to go to **Advanced Options**.

Advanced Options

(Optional) You can select [Ansible playbook services](#) to run on the virtual machines in the migration plan before or after the migration. For example, the [RHEL pre-migration playbook](#) performs several tasks to configure a target RHEL virtual machine. You must create the [Ansible environment and service catalog item](#) before the migration plan.

2. Click **Next** to go to **Schedule**.
 3. Select **Save migration plan to run later** or **Start Migration Immediately** and click **Create**.
- If you selected **Save migration plan to run later**, the migration plan is saved in **Migration Plans Not Started**. To schedule the migration plan, click the **Schedule** button, select a date and time, and click **Schedule**. The plan's status is **Migration Scheduled** with the date and time.



NOTE

If you wish to run a scheduled plan immediately, click **Migrate**.

To edit or delete a schedule or plan, click the **More Actions** icon.

- If you selected **Start Migration Immediately**, the migration plan runs immediately. The migration plan may take some time to complete. The migration plan displays progress bars showing the amount of data being transferred, the number of virtual machines being migrated, and the elapsed time.

To view the migration progress of individual virtual machines, click **Migration Plans in Progress** → **Migration Plan** to display the migration plan details.

**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.

1. When the migration plan has finished, check the status of the migration plan under **Completed Migrations**. The completed migration plan shows successfully migrated virtual machines. If error messages appear, see [Chapter 5, Troubleshooting](#).

Remigrating Virtual Machines

If your migration goals change, you can remigrate previously migrated virtual machines:

1. Delete the target virtual machines corresponding to the source virtual machines that you are remigrating.
2. Delete the disks that were created in the target datastore during the earlier migration to free up space.
3. Create a [CSV file](#) to import the source virtual machines.

**NOTE**

The source 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.

4.3. CANCELING A MIGRATION PLAN

You can cancel a migration plan that has started and displays the progress bars:

1. Click **Migration Plans in Progress** → **Migration Plan** to display the migration plan details.
2. Select the migration plan check box 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.

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

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](#) when you create the migration plan.

The Ansible service catalog item must exist before you create the migration plan. The CloudForms inventory must contain at least one repository, one playbook, and one credential.

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

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

IMS RHEL Pre-migration Ansible Playbook Example

The `ims.rhel_pre_migration` role performs the following tasks to configure a target Red Hat Enterprise Linux virtual machine:

- 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. TROUBLESHOOTING

If errors occur for reasons unrelated to the infrastructure mapping or migration plan, you can [retry the migration plan](#).

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 A MIGRATION PLAN

If one or more virtual machines in a migration plan fails to migrate, the migration plan’s status is **Failed**. After resolving the issue that caused the migration failure, you can retry the migration plan:

1. Click the migration plan to go to the details page to identify the virtual machines that failed to migrate.
2. If the migration failed while virtual machines were being converted:
 - a. Delete newly created target virtual machines to avoid name conflicts with source virtual machines.
 - b. Delete newly created disks in the target datastore to free up space.
3. Click **Compute** → **Migration**, select **Failed Migrations**, select the failed migration plan, and click **Retry**.

5.2. LOGS

If a migration plan fails and you are not sure of the cause, check the following logs:

- CloudForms migration log: `/var/www/miq/vmdb/log/automation.log`
- Conversion host logs: `/var/log/vdsm/import/`
- Virtual machine migration log: Click **Compute** → **Migrate**, click the migration plan name, and click **Download Log**.

Migrated Virtual Machine Logs

VM Name	Elapsed Time	Migration Status	Action
[VM Name]	00:41:58	VM migrations completed	Download Log
[VM Name]	00:42:10	VM migrations completed	Download Log
[VM Name]	00:44:01	VM migrations completed	Download Log
[VM Name]	00:45:21	VM migrations completed	Download Log
[VM Name]	00:45:29	VM migrations completed	Download Log


5.3. COMMON ISSUES AND MISTAKES

Infrastructure mapping is missing datastores or clusters

If you create an infrastructure mapping and subsequently add or remove providers or provider objects, the infrastructure mapping displays **Datastores missing** or **Clusters missing** error messages because the object IDs of the providers and their objects have changed. You must delete the infrastructure mapping and re-create it.

Infrastructure mapping is missing networks

If the infrastructure mapping displays a **Networks missing** error message, you must [authenticate the RHV conversion hosts](#) and create a new infrastructure mapping.



Migration plan does not discover virtual machines

If you are migrating previously migrated virtual machines, the migration plan cannot discover them because they are marked in the CloudForms VMDB as **migrated**. You must [remigrate the virtual machines](#).

Incorrect attributes in the virtual machine CSV import file

Correct the [CSV file](#), create a new migration plan, import the updated CSV file, and run the migration plan.

Migration fails immediately

If a migration fails immediately, before the progress bar is displayed, this is often caused by a configuration error in the environment. Check the **automation.log** file on the CloudForms appliance to determine the cause. When the cause has been resolved, you can [retry the migration plan](#).

Migration fails during virtual machine conversion

If a migration fails while 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 virtual machine(s) that failed to migrate. When the cause of the failed migration has been resolved, you can [retry the migration plan](#).

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 [BZ#1544379](#) for details.