

6. Now that the database has been migrated, start the **ovirt-engine** service:

```
# systemctl start ovirt-engine.service
```

12.2. BACKING UP AND RESTORING VIRTUAL MACHINES USING THE BACKUP AND RESTORE API

12.2.1. The Backup and Restore API

The backup and restore API is a collection of functions that allows you to perform full or file-level backup and restoration of virtual machines. The API combines several components of Red Hat Virtualization, such as live snapshots and the REST API, to create and work with temporary volumes that can be attached to a virtual machine containing backup software provided by an independent software provider.

For supported third-party backup vendors, consult the [Red Hat Virtualization Ecosystem](#).

12.2.2. Backing Up a Virtual Machine

Use the backup and restore API to back up a virtual machine. This procedure assumes you have two virtual machines: the virtual machine to back up, and a virtual machine on which the software for managing the backup is installed.

Backing Up a Virtual Machine

1. Using the REST API, create a snapshot of the virtual machine to back up:

```
POST /api/vms/{vm:id}/snapshots/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<snapshot>
  <description>BACKUP</description>
</snapshot>
```

Note

- » Here, replace **{vm:id}** with the ID of the virtual machine whose snapshot you are making.
- » When you take a snapshot of a virtual machine, a copy of the configuration data of the virtual machine as at the time the snapshot was taken is stored in the **data** attribute of the **configuration** attribute in **initialization** under the snapshot.

Important

You cannot take snapshots of disks that are marked as shareable or that are based on direct LUN disks.

2. Retrieve the configuration data of the virtual machine from the **data** attribute under the snapshot:

```
GET /api/vms/{vm:id}/snapshots/{snapshot:id} HTTP/1.1
All-Content: true
Accept: application/xml
Content-type: application/xml
```

Note

- » Here, replace **{vm:id}** with the ID of the virtual machine whose snapshot you made earlier. Replace **{snapshot:id}** with the snapshot ID.
- » Add the **All-Content: true** header to retrieve additional OVF data in the response.

3. Identify the disk ID and snapshot ID of the snapshot:

```
GET /api/vms/{vm:id}/snapshots/{snapshot:id}/disks HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

4. Attach the snapshot to a backup virtual machine as an active disk attachment, with the correct interface type (for example, **virtio_scsi**):

```
POST /api/vms/{vm:id}/diskattachments/ HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

```
<disk_attachment>
  <active>true</active>
  <interface>_virtio_scsi_</interface>
  <disk id="{disk:id}">
  <snapshot id="{snapshot:id}"/>
</disk>
</disk_attachment>
```

Note

Here, replace **{vm:id}** with the ID of the *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace **{disk:id}** with the disk ID. Replace **{snapshot:id}** with the snapshot ID.

5. Use the backup software on the backup virtual machine to back up the data on the snapshot disk.
6. Remove the snapshot disk attachment from the backup virtual machine:

```
DELETE /api/vms/{vm:id}/diskattachments/{snapshot:id} HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

Note

Here, replace **{vm:id}** with the ID of the *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace **{snapshot:id}** with the snapshot ID.

7. Optionally, delete the snapshot:

```
DELETE /api/vms/{vm:id}/snapshots/{snapshot:id} HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

Note

Here, replace **{vm:id}** with the ID of the virtual machine whose snapshot you made earlier. Replace **{snapshot:id}** with the snapshot ID.

You have backed up the state of a virtual machine at a fixed point in time using backup software installed on a separate virtual machine.

12.2.3. Restoring a Virtual Machine

Restore a virtual machine that has been backed up using the backup and restore API. This procedure assumes you have a backup virtual machine on which the software used to manage the previous backup is installed.

Restoring a Virtual Machine

1. In the Administration Portal, create a floating disk on which to restore the backup. See [Section 10.6.1, "Creating a Virtual Disk"](#) for details on how to create a floating disk.
2. Attach the disk to the backup virtual machine:

```
POST /api/vms/{vm:id}/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<disk id="{disk:id}">
</disk>
```

Note

Here, replace **{vm:id}** with the ID of this *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace **{disk:id}** with the disk ID.

3. Use the backup software to restore the backup to the disk.
4. Detach the disk from the backup virtual machine:

```
DELETE /api/vms/{vm:id}/disks/{disk:id} HTTP/1.1
Accept: application/xml
Content-type: application/xml

<action>
  <detach>true</detach>
</action>
```

+ NOTE: Here, replace **{vm:id}** with the ID of this *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace **{disk:id}** with the disk ID.

5. Create a new virtual machine using the configuration data of the virtual machine being restored:

```
POST /api/vms/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<vm>
  <cluster>
    <name>cluster_name</name>
  </cluster>
  <name>_NAME_</name>
  <initialization>
  <configuration>
<data>
< -- omitting long ovf data -->
</data>
  <type>ovf</type>
  </configuration>
  </initialization>
  ...
</vm>
```

Note

To override any of the values in the ovf while creating the virtual machine, redefine the element within the **initialization** element. For example, to change the name of the virtual machine, use **<name>OVERRIDING_NAME</name>**.

6. Attach the disk to the new virtual machine:

```
POST /api/vms/{vm:id}/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<disk id="{disk:id}">
</disk>
```

Note

Here, replace `{vm:id}` with the ID of the *new* virtual machine, not the virtual machine whose snapshot you made earlier. Replace `{disk:id}` with the disk ID.

You have restored a virtual machine using a backup that was created using the backup and restore API.

CHAPTER 13. ERRATA MANAGEMENT WITH RED HAT SATELLITE

Red Hat Virtualization can be configured to view errata from Red Hat Satellite in the Red Hat Virtualization Manager. This enables the administrator to receive updates about available errata, and their importance, for hosts, virtual machines, and the Manager once they have been associated with a Red Hat Satellite provider. Administrators can then choose to apply the updates by running an update on the required host, virtual machine, or on the Manager. For more information about Red Hat Satellite see the [Red Hat Satellite User Guide](#).

Red Hat Virtualization 4.2 supports errata management with Red Hat Satellite 6.1.

Important

The Manager, hosts, and virtual machines are identified in the Satellite server by their FQDN. This ensures that external content host IDs do not need to be maintained in Red Hat Virtualization.

The Satellite account used to manage the Manager, hosts and virtual machines must have Administrator permissions and a default organization set.

Configuring Red Hat Virtualization Errata

To associate a Manager, host, and virtual machine with a Red Hat Satellite provider first the Manager must be associated with a provider. Then the host is associated with the same provider and configured. Finally, the virtual machine is associated with the same provider and configured.

1. Associate the Manager by adding the required Satellite server as an external provider. See [Section 11.2.1, “Adding a Red Hat Satellite Instance for Host Provisioning”](#) for more information.

Note