



# VM Life Cycle Management

---

VM life cycle management refers to the entire process of registering, deploying, updating, monitoring VMs, and getting them service chained as per your requirements. You can perform these tasks and more using a set of REST APIs or NETCONF commands or the Cisco Enterprise NFVIS portal.

## VM Packaging Format

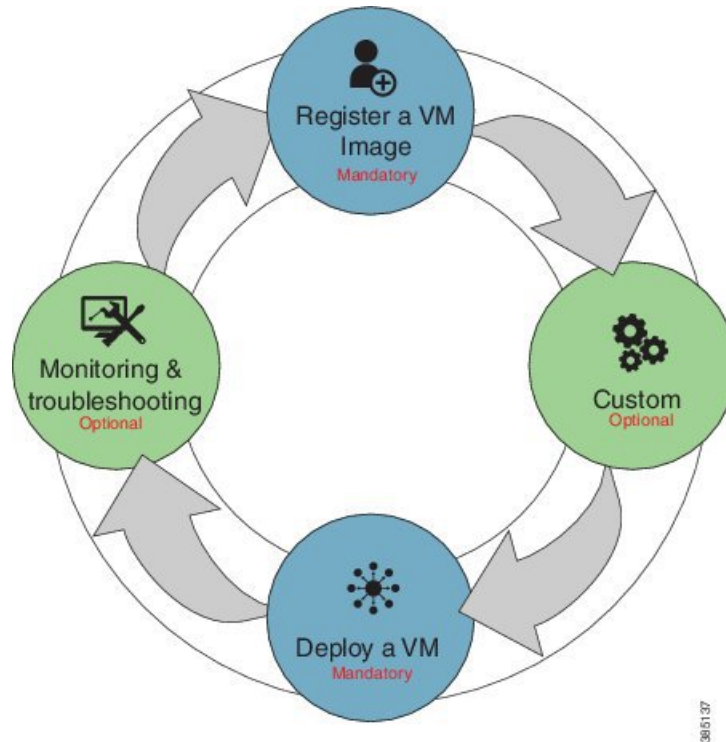
All VM images are available in the `.tar.gz/qcow2/vmdk/img/iso` format. All Cisco supplied VMs are available in the prescribed format. Vendors are responsible for packaging all third party VMs in the prescribed format.

- [Workflow of VM Life Cycle Management, on page 1](#)
- [Uploading VM Images to an NFVIS Server, on page 3](#)
- [VM Bootstrap Configuration Options with a VM Deployment, on page 4](#)
- [OpenStack Configuration Drive Support for Third Party VMs, on page 5](#)
- [Performing Resource Verification, on page 6](#)
- [Configuring Management IP Address, on page 7](#)
- [VM States, on page 7](#)

## Workflow of VM Life Cycle Management

The following diagram depicts the basic workflow of the VM life cycle management using REST APIs:

Figure 1: VM Life Cycle Management



1. **Register a VM Image**—To register a VM image, you must first copy or download the relevant VM image to the NFVIS server, or host the image on a http or https server. Once you have downloaded the file, you can register the image using the registration API. The registration API allows you to specify the file path to the location (on the http/https server) where the tar.gz file is hosted. Registering the image is a one-time activity. Once an image is registered on the http or https server, and is in active state, you can perform multiple VM deployments using the registered image.
2. **Customizing the Setup**—After registering a VM image, you can optionally create a custom profile or flavor for the VM image if the profiles defined in the image file do not match your requirement. The flavor creation option lets you provide specific profiling details for a VM image, such as the virtual CPU on which the VM will run, and the amount of virtual memory the VM will consume.

Depending on the topology requirement, you can create additional networks and bridges to attach the VM to during deployment.

3. **Deploy a VM**— A VM can be deployed using the deployment API. The deployment API allows you to provide values to the parameters that are passed to the system during deployment. Depending on the VM you are deploying, some parameters are mandatory and others optional.
4. **Manage and Monitor a VM**—You can monitor a VM using APIs and commands that enable you to get the VM status and debug logs. Using VM management APIs, you can start, stop, or reboot a VM, and view statistics for a VM such as CPU usage.

A VM can also be managed by changing or updating its profile. You can change a VM's profile to one of the existing profiles in the image file; alternatively, you can create a new custom profile for the VM.

The vNICs on a VM can also be added or updated.



**Note** Before performing the VM life cycle management tasks, you will have to upload the VM images to the NFVIS server or http/s server.

For details on APIs, see the [VM Lifecycle Management APIs](#) chapter in the *API Reference for Cisco Enterprise Network Function Virtualization Infrastructure Software*.

## Uploading VM Images to an NFVIS Server

You can upload VM images to an NFVIS server in the following ways. The files are copied to the default location (/data/intdatastore/uploads) on the host server.

- Copy the images from your local system to the NFVIS server—Use the **Image Upload** option from the Cisco Enterprise NFVIS portal.
- Copy the images using the USB drive—Ensure that you have plugged the USB drive that contains the required images into the server before mounting the USB drive.
- Copy using the **scp** command (scp username@external\_server:/path/image.tar.gz intdatastore:image.tar.gz).



**Note** From 3.8.1 release, NFVIS supports deleting images while the download is in progress. NFVIS also supports resuming image download after a power outage or lost connectivity.

To copy an image using the USB device:

```
configure terminal
system usb-mount mount ACTIVE
system file-copy usb file name usb1/package/isrv-universalk9.16.03.01.tar.gz
commit
```



**Note** Use the **show system file-list disk usb** command in privileged EXEC mode to view a list of files available with the mounted USB drive. To save space, you can delete all unwanted text and TAR files from the default location using the **system file-delete** command in global configuration mode.

### Verifying the Image Copied from the USB Drive

After copying the file from the USB drive to the host server, you can verify the file using the **show system file-list disk local** command:

```
nfvis# show system file-list disk local
```

SI	NO	NAME	PATH	SIZE	TYPE	DATE	MODIFIED
1		lastlog-20170314.gz	/data/intdatastore/logs/2017-03/14/10-00	337	Other	2017-03-14	21:55:42
2		escmanager-tagged-log.log-20170314.gz	/data/intdatastore/logs/2017-03/14/10-00	167K	Other		

```

2017-01-18 05:58:26
3 confd_audit.log-20170317.gz /data/intdatastore/logs/2017-03/17/09-30 4.6K Other 2017-03-17
  21:29:59
4 esc_postinit.log-20170317.gz /data/intdatastore/logs/2017-03/17/05-00 605K Other 2017-03-17
  16:40:19
5 error.log-20170317.gz /data/intdatastore/logs/2017-03/17/05-00 1.3K Other 2017-03-17
  16:40:15
6 ovs-ctl.log-20170317.gz /data/intdatastore/logs/2017-03/17/12-00 20 Other 2017-03-16
  00:00:01 4:01
!
!
!
62 ovs-ctl.log-20170323.gz /data/intdatastore/logs/2017-03/23/12-00 20 Other 2017-03-22
  00:00:01
63 CentOS-7-x86_64-Everything-1511.ova /data/intdatastore/uploads 1.1G VM 2017-03-15 19:20:03
  Package
64 TinyLinux.tar.gz /data/intdatastore/uploads 17M VM 2017-03-15 18:25:00 Package
65 Cisco-KVM-vWAAS-1300-6.3.0-b98.tar.gz /data/intdatastore/uploads 979M VM 2017-03-15
  19:19:11 Package
66 ubuntu_14.04.3-server-amd64-disk1.tar /data/intdatastore/uploads 527M VM 2017-03-15
  19:20:17.gz Package
67 asav961.tar.gz /data/intdatastore/uploads 164M VM 2017-03-15 18:24:57 Package
68 isrv-universalk9.16.03.01.tar.gz /data/intdatastore/uploads 1.3G VM 2017-03-15 19:19:53

```

### Related APIs and Commands

APIs	Commands
<ul style="list-style-type: none"> <li>• /api/operations/system/file-copy/usb/file</li> <li>• /api/config/system/usb-mount</li> </ul>	<ul style="list-style-type: none"> <li>• system file-copy usb file name</li> <li>• system usb-mount mount ACTIVE</li> <li>• system file-delete</li> <li>• show system file-list disk usb</li> <li>• show system file-list disk local</li> </ul>

## VM Bootstrap Configuration Options with a VM Deployment

You can include the bootstrap configuration (day zero configuration) of a VM in the VM deployment payload in the following three ways:

- Bundle bootstrap configuration files into the VM package—In this method, the bootstrap configuration variables can be tokenized. Token names must be in bold text. For each tokenized variable, key-value pairs must be provided during deployment in the deployment payload.
- Bootstrap configuration as part of the deployment payload—The entire bootstrap configuration is copied to the payload without tokens.
- Bootstrap configuration file in the NFVIS server—In this method, the configuration file is copied or downloaded to the NFVIS server, and referenced from the deployment payload with the filename including full path.

For examples on how to use bootstrap configuration options in the deployment payload, see the [API Reference for Cisco Enterprise Network Function Virtualization Infrastructure Software](#).

## OpenStack Configuration Drive Support for Third Party VMs

To enable staging of bootstrap configuration files at the time of a third party VM deployment as per OpenStack standards, the following cloud init format is supported:

```
openstack/content
openstack/content/0000
openstack/content/0001
openstack/latest/meta_data.json
```

In the above sample, the "0000" and "0001" files are the actual bootstrap files from the deployment payload. A third party VM can use the init file to fetch its configuration files.

The following metadata file is used to provide the file path on the configuration drive and reference to the actual bootstrap configuration files.

```
{
  "files": [
    {
      "content_path": "/content/0000",
      "path": "/config/day-0.txt"
    },
    {
      "content_path": "/content/0001",
      "path": "/sample/path/iosxe_config.txt"
    }
  ]
}
```

With this implementation, two copies of the same bootstrap configuration file will be present on the virtual CD-ROM package. The first version at the root (iosxe\_config.txt) and the second inside the "openstack/content" folder.

The admin will also have to specify the bootstrap configuration file in the image properties file before packaging the VM.

### Example for the Bootstrap Configuration File in the Image Properties File

```
--optimize=OPTIMIZE [REQUIRED] optimized VM: --optimize=true/false;
--root_file_disk_bus=ROOT_FILE_DISK_BUS root disk file type:
--root_file_disk_bus=virtio/ide; default is virtio
--virtual_interface_model=VIRTUAL_INTERFACE_MODEL
--virtual_interface_model=rtl8139; default is none
--thick_disk_provisioning=THICK_DISK_PROVISIONING
--thick_disk_provisioning=true; default is false
--bootstrap_cloud_init_bus_type=BOOTSTRAP_CLOUD_INIT_BUS_TYPE
--bootstrap_cloud_init_bus_type=virtio; default is ide
--bootstrap_cloud_init_drive_type=BOOTSTRAP_CLOUD_INIT_DRIVE_TYPE
--bootstrap_cloud_init_drive_type=disk; default is cdrom
--bootstrap=BOOTSTRAP bootstrap file/s for VM (two parameters required in the format of
dst:src; dst filename including path has to match exactly to what the VM expects;
upto 20 bootstrap files are accepted.)
examples:
--bootstrap ovf-env.xml:file1,ios-xe.txt:file2 for ISRV; both files get mounted at the
```

```

root level on the VM.
--bootstrap day0-config:filename1 for ASAv
--bootstrap
/:bootstrap.xml,/license/lic.txt:license.txt
bootstrap.xml get mounted as bootstrap.xml at root, and license.txt get mounted as
/license/lic.txt.

```



**Note** If any of the strings in the configuration file has wild characters, wrap the string with this `#[ ]#` so that the token/key replacement engine does not consider wild characters as key or token, and looks for key value pairs to replace during a VM deployment.

For details on the OpenStack standards, visit <http://docs.openstack.org>.

## Performing Resource Verification

Given below are the APIs and commands to perform different types of resource verification:

Task	API	Command
To display CPU information for each CPU or the user specified CPU, and the VMs pinned to the CPU	<ul style="list-style-type: none"> <li>• <code>api/operational/resources/cpu-info/cpus</code></li> <li>• <code>/api/operational/resources/cpu-info/cpus/cpu</code></li> <li>• <code>/api/operational/resources</code> <code>/cpu-info/cpus/cpu/&lt;cpu-id&gt;</code></li> </ul>	<code>show resources cpu-info cpus</code>
To display information on the VMs running in all the physical CPUs or a specific physical CPU in the system	<ul style="list-style-type: none"> <li>• <code>/api/operational/resources/cpu-info/vnfs</code></li> <li>• <code>/api/operational/resources/cpu-info/vnfs/vnf</code></li> <li>• <code>/api/operational/resources/cpu-info/vnfs/vnf/</code> <code>&lt;deployment_name&gt;.&lt;vm_group_name&gt;</code></li> </ul>	<code>show resources cpu-info vnfs</code>
To get information on the number of CPUs allocated to VMs and the CPUs that are already used by the VMs	<code>/api/operational/resources/cpu-info/allocation</code>	<code>show resources cpu-info allocation</code>



**Note** To display information on all CPUs, VMs pinned to the CPUs, and VMs allocated to the CPUs, use the **show resources cpu-info** command.

### CPU Over-Subscription

Cisco Enterprise NFVIS does not allow CPU over-subscription for low-latency network appliance VMs (for example, Cisco ISRV and Cisco ASAv). However, the CPU over-subscription is allowed for non low-latency VMs (for example, Linux Server VM and Windows Server VM).

## Configuring Management IP Address

The following commands need to be executed in a sequence to first delete an existing subnet and then add a new subnet in the network. For these commands to work, ensure there is no managed VNF's in the system before you change management network address.

To delete an existing subnet use **no vm\_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet** command.

To create a new subnet:

```
configure terminal
vm_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet address 105.20.0.0
gateway 105.20.0.1 netmask 255.255.255.0 dhcp false
commit
```

## VM States

VM States	Description
VM_UNDEF_STATE	The initial state of VM or VNF before deployment of this VM.
VM_DEPLOYING_STATE	VM or VNF is being deployed on to the NFVIS.
VM_MONITOR_UNSET_STATE	VM or VNF is deployed in the NFVIS but the monitoring rules are not applied.
VM_MONITOR_DISABLED_STATE	Due to a VM action request or recovery workflow, the monitoring or KPI rules applied on the VM or VNFs were not enabled.
VM_STOPPING_STATE	VM or VNF is being stopped.
VM_SHUTOFF_STATE	VM or VNF is in stopped or shutoff state.
VM_STARTING_STATE	VM or VNF is being started.
VM_REBOOTING_STAT	VM or VNF is being rebooted.
VM_INERT_STATE	VM or VNF is deployed but not alive. The KPI monitor is applied and waiting for the VM to become alive.
VM_ALIVE_STATE	VM or VNF is deployed and successfully booted up or alive as per the monitor or kpi metric.
VM_UNDEPLOYING_STATE	VM or VNF is being undeployed or terminated.
VM_ERROR_STATE	VM or VNF will be in error state if deployment or any other operation is failed.

