



NFVIS Clustering Capability for Cisco Cloud Services Platform

Table 1: Feature History

Feature Name	Release Information	Description
NFVIS Clustering Capability for Cisco CSP	<p>Cisco NFVIS 4.8.1</p> <p>Note This capability is a beta feature in this release and should not be used in production because, Cisco may make changes in the subsequent releases without providing any backward compatibility. The beta is provided as-is and without any warranty of any kind.</p>	This feature enables you to combine three nodes into a single cluster definition so that, all the member nodes display configuration information about all the virtual machines deployed in the cluster. Each member node contributes to a user-specified disk size (same value for all the members) of their total disk size to form a shared and clustered file system. The data is replicated across all the cluster members. This clustering capability is resilient against single node failures.

- [Prerequisites for NFVIS Clustering Capability for Cisco CSP, on page 1](#)
- [Restrictions for NFVIS Clustering Capability on Cisco CSP, on page 2](#)
- [Information About NFVIS Clustering Capability for Cisco CSP, on page 2](#)
- [Configure NFVIS Clustering Capability on Cisco CSP, on page 3](#)
- [Cluster File Operations, on page 6](#)
- [Support Commands for NFVIS Clustering Capability , on page 7](#)
- [Verify NFVIS Clustering Capability, on page 7](#)
- [NFVIS Clustering Capability Notifications and Syslogs, on page 9](#)

Prerequisites for NFVIS Clustering Capability for Cisco CSP

Ensure that you configure NTP, so that the system time on all nodes of the cluster is synchronized.

Restrictions for NFVIS Clustering Capability on Cisco CSP

- To create a cluster for storage virtualization, you can use only three Cisco CSP devices with the same configuration. Configure an IP address on the new bridge. The IP addresses should match one of the three IP addresses used for creating the cluster. We recommend that you configure the 1x10 GigE NIC for data traffic, and Port channel or physical network interface card pNICs.
- When a cluster is created, the Maximum Transmission Unit (MTU) on the cluster bridge is automatically set to 9000. Ensure that the switch present in the network that connects to the pNIC supports an MTU of 9000.
- A cluster cannot be created on NFVIS hosts that have virtual machines deployed on them. If a system has virtual machine deployments, export and delete the virtual machine, before creating the cluster.
- The backup from an old host cannot be restored once the cluster is created, and the backup for a new host cannot be created after a cluster is created.
- When a cluster is created, all the nodes in the cluster share their storage with each other. Other NFVIS configurations such as additional bridges, OVS networks, users, groups, TACACS/RADIUS, and SNMP are not synchronized and must be created on each cluster node, separately.
- This feature does not support storage migration for virtual machines. Cold virtual machine migration between datastores (intdatastore, extdatastore1/2, and shared datastore) is not supported.
- You cannot add, delete, or replace an existing cluster with a new NFVIS host. If you require a new NFVIS host, in case of faulty hardware, delete the existing cluster and create a new cluster with the new host.

Information About NFVIS Clustering Capability for Cisco CSP

NFVIS Clustering Capability

When a device in the network is down, the virtual machines running on the device also go down and become inaccessible. Before Cisco NFVIS 4.8.1, there are two ways to recover the virtual machines:

- Recover the virtual machine from a back-up. The drawback of this method is that the backup could be outdated depending on when the backup was taken.
- Deploy the virtual machine in an active-standby mode. The drawback of this method is when the virtual machine is down, the standby takes over. However, this would require two separate licences, one for the active virtual machine and one for the standby virtual machine.

Starting from Cisco NFVIS 4.8.1, the NFVIS clustering capability is introduced. This feature allows you to form a cluster of 3 nodes. The disks of all virtual machines deployed on the cluster are replicated on all the cluster nodes. This enables you to perform a cold migration of one virtual machine to a different node in the cluster when the virtual machine's node goes down.

Guidelines to Enable Cluster Data Traffic


Use a dedicated pNIC to enable the flow of cluster data traffic between the nodes, and this pNIC port should be attached to a newly created bridge for the cluster traffic to flow. Configure an IP address on the new bridge. The same IP address is one of the three IP addresses used for creating the cluster.

We recommend that you configure the 1×10 GigE NIC for data traffic.

Enable Cluster Data Traffic

The following diagram illustrates a dedicated pNIC that enables the flow of cluster data traffic between the nodes. The dedicated pNIC port is attached to a newly created bridge for the cluster traffic to flow.

Figure 1: Cluster Data Traffic Flow



* On CSP, lan-br is used for management by default. There is no dedicated management port like the ENCS

466143

Configure NFVIS Clustering Capability on Cisco CSP

Before creating the cluster, ensure to complete the following steps on all the participating nodes:

1. Configure the time server.

```
nfvis# config terminal
Entering configuration mode terminal
nfvis(config)# system time ntp preferred_server YOUR-NTPD-IP
nfvis(config)# end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
nfvis#
```

2. Configure the cluster bridge with IP.

```
nfvis# config t
Entering configuration mode terminal
nfvis(config)# bridges bridge cluster-br
nfvis(config-bridge-cluster-br)# ip address 198.51.100.1 255.255.255.0
nfvis(config-bridge-cluster-br)# port eth1-1
nfvis(config-port-eth1-1)# end
Uncommitted changes found, commit them? [yes/no/CANCEL] yes
Commit complete.
nfvis#
```



Note Ensure that the chosen PNIC and IP can communicate with each other.

Create a Cluster

To create a cluster:

1. Send the create cluster API call to all the Cisco NFVIS that must be added to a cluster for storage virtualization.



Note

- A cluster for storage virtualization can contain only three devices.
- The clustering API calls should be issued to all the three nodes within three minutes.
- The API call contains the name of the cluster, the IPv4 address of the nodes to be added to the cluster, and the disk size to be contributed for clustered storage. The API call also has the datastore parameter to indicate which datastore (internal or external datastore) contributes towards the storage pool.

2. After receiving the call, each server validates whether sufficient disk space is allocated for storage clustering. The server also ensures that the space utilization of the datastore is less than 80% . After the disk size validation, a reply is sent indicating that the server has received the API call. When the cluster formation is successful, a notification is pushed from the leader node to indicate that storage virtualization is complete.

The following example shows how to create a cluster:

```
nfvis# config terminal
Entering configuration mode terminal
nfvis(config)# cluster cluster1 datastore intdatastore size 10
nfvis(config-cluster-1)# node 209.165.201.1 address-type ipv4
nfvis(config-node-209.165.201.1)# exit
nfvis(config-cluster-2)# node 209.165.201.2 address-type ipv4
nfvis(config-node-209.165.201.2)# exit
nfvis(config-cluster-3)# node 209.165.201.3 address-type ipv4
nfvis(config-node-209.165.201.3)# commit
```

The cluster distributed file system is created, and it serves as a storage location for virtual machine registration and deployment. For information on verifying cluster creation, see [Verify Storage Virtualization](#).

Register and Deploy the VM

To register and deploy VMs on the cluster, specify **gluster** as the placement option. The VM's disk details are replicated across all the other nodes on the cluster.

The following information is available to all the nodes and can be used while migrating the VM:

- The running configuration of image registration
- Flavor
- Deployment

The following parameters should be unique across the cluster (for both cluster VMs and local VMs):

- Flavor name
- Deployment names
- SRIOV networks
- int-mgmt IPs
- Port numbers used for port forwarding

During VM registration and deployment, if any of the above mentioned parameters aren't unique, the deployment is rejected.

The following example shows how to register the image:

```
nfvis(config)# vm_lifecycle images image centos src
file:///data/intdatastore/uploads/Centos_7.tar.gz properties property placement value gluster
```

```
nfvis(config-property-placement)# commit
```

Commit complete.

The following example shows how to deploy the VM:

```
nfvis(config)# vm_lifecycle tenants tenant admin deployments deployment centosvm vm_group
depl image centos flavor centos-small bootup_time -1 vim_vm_name centosvm placement zone_host
host gluster
nfvis(config-placement-zone_host)# commit
```

Commit complete.

Migrate the Virtual Machine

virtual machines deployed on the cluster can be migrated to a different node on the cluster, either after they are shut down, or while the node on which they were deployed is down. Specify the API for the source and destination nodes for migration. You can choose to migrate all virtual machines from a source node or specify a list of virtual machines to migrate.

The virtual machine migration API is sent to any of the working nodes. If the host is up during migration, stop the virtual machine.

The virtual machine disk is replicated in the destination node, the virtual machine comes up in the Day-N state. After the migration, the source node is active, and created on the destination node. If the source node is down during the migration, the clean-up on the source node happens when the node comes back up.

The following example shows how to migrate all deployments from the source node to the destination node:

```
nfvis# cluster cluster1 migrate-deployment source-node 209.165.201.1 destination-node
209.165.201.2 all-deployments
```

The following example shows how to migrate specific deployments from the source node to the destination node:

```
nfvis# cluster cluster1 migrate-deployment source-node 209.165.201.1 destination-node
209.165.201.2 deployment-list [ centosvirtual machine4 centosvirtual machine5 ]
```

For information about verifying virtual machine migration, see [Verify Storage Virtualization](#).

Virtual Machine Export and Import

Virtual Machines can be exported to any **local/external datastore/NFS** mount path, and can be imported from the **/mnt/gluster/uploads** folder.



Note You cannot export Virtual Machines to the **/mnt/gluster/uploads** folder.

The following example shows how to export a virtual machine:

```
nfvis# vmExportAction exportName asaexport exportPath intdatastore:/uploads/ vmName
ASAvDep.ASAvmGrp
```

The following example shows how to import a virtual machine:

```
nfvis# vmImportAction importPath /mnt/gluster/uploads/backup_router.vmbkp
(or)
nfvis# vmImportAction importPath gluster:/uploads/backup_router.vmbkp
```

Delete the Cluster

Deleting a cluster is similar to cluster creation. The API for cluster deletion has a “no” prefix followed by **cluster-name**. The delete cluster API should be sent to all the nodes in the cluster to delete the cluster.

Delete the virtual machine images and deployments on the cluster before deleting the cluster.



Note After a delete cluster API is sent to a node, sending another create cluster API to the same node does not add the node back to the cluster.

To delete a cluster, use the following command:

```
nfvis(config)# no cluster cluster1
```

Cluster File Operations

Use the following commands to manage the files in the glusterFS storage.

To copy files using the file-copy command:

```
nfvis# system file-copy source /data/intdatastore/uploads/test_file_in_intdatastore.txt
destination /mnt/gluster/uploads
```

To copy files using the scp command:

```
scp into /mnt/gluster/uploads/
nfvis# scp user@172.16.0.1:/nobackup/userdir/test.py gluster:test.py

scp from /mnt/gluster/uploads/
nfvis# scp gluster:test.py user@172.16.0.1:/nobackup/userdir/

scp from outside
nfvis# scp -P 22222 test.txt admin@172.16.0.2:/mnt/gluster/uploads/
```

To view the files present in glusterFS:

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

SI	NO	NAME	PATH	SIZE	TYPE	DATE	MODIFIED
1		TinyLinux.tar.gz	/mnt/gluster/uploads	17M	VM Package	2022-01-18	16:21:03

To delete files from /mnt/gluster/uploads:

```
nfvis# system file-delete file name /mnt/gluster/uploads/test.txt
```



Note The **file-copy** and **file-delete** commands only work for the following folders:

- /data/intdatastore/
- /data/intdatastore/uploads/
- /mnt/extdatastore1/
- /mnt/extdatastore1/uploads/
- /mnt/extdatastore2/
- /mnt/extdatastore2/uploads/
- /mnt/gluster/uploads/ folder

Support Commands for NFVIS Clustering Capability

The following support commands can be used to debug NFVIS clustering:

The following example shows how to check the etcd status:

```
nfvis# support show etcd
```

Possible completions:

all-information	Display all etcd information
cluster-health	Display cluster health
member-list	Display members of the cluster
service-status	Display status of etcd services

The following example shows how to check the glusterFS status:

```
nfvis# support show gluster
```

Possible completions:

all-information	Display all glusterfs information
mount	Display the glusterfs mount information
peer-status	Display the glusterfs peer connectivity information
service-status	Display status of glusterfs services
volume-info	Display the glusterfs volume information

Verify NFVIS Clustering Capability

The following example shows the success scenario of NFVIS clustering capability:

```

nfvis# show cluster
cluster cluster1
node-address 209.165.201.1
cluster-state ok
creation-time 2022-01-18T15:32:23-00:00
details "Gluster peers Healthy"
disk-usage total 10.0
disk-usage available 10.0

```

ADDRESS	DEPLOYMENT	STATE
209.165.201.1		
209.165.201.2		
209.165.201.3		

The following example shows the cluster degraded scenario:

```

nfvis# show cluster
cluster cluster1
node-address 209.165.201.1
cluster-state degraded
creation-time 2022-01-18T15:32:23-00:00
details "One Node:
209.165.201.3 in the cluster is down"
disk-usage total 10.0
disk-usage available 10.0

```

ADDRESS	DEPLOYMENT	STATE
209.165.201.1		
209.165.201.2		
209.165.201.3		

The following example shows the cluster down scenario:

```

nfvis# show cluster
cluster cluster1
node-address 209.165.201.1
cluster-state down
creation-time 2022-01-18T15:32:23-00:00
details "Etcd is inactive"
disk-usage total 10.0
disk-usage available 10.0

```

ADDRESS	DEPLOYMENT	STATE
209.165.201.1		
209.165.201.2		
209.165.201.3		

The following example shows a successful VM migration:

```

nfvis# show cluster migration-status
cluster cluster1
migration-status start-time 2022-01-19T14:48:28
migration-status source-node 209.165.201.1
migration-status destination-node 209.165.201.2
migration-status status MIGRATION-SUCCESS
migration-status details ""
DEPLOYMENT STATUS    LAST UPDATE DETAILS
-----
CentosVM    MIGRATED 2022-01-19  22:48:56.897 Migrated VM[CentosVM] to new host[209.165.201.1]

```


OTHER56 MIGRATED 2022-01-19 22:48:57.214 Migrated VM[OTHER56] to new host[209.165.201.2]



Note In NFVIS 4.8, the **show cluster** command only displays the latest migration status.

NFVIS Clustering Capability Notifications and Syslogs

Table 2: Notifications

Description	Notification	Seen on
Cluster creation in progress	notification eventTime 2022-01-27T12:48:34.440959+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status SUCCESS status_code 0 status_message Cluster creation in progress details NA event_type CLUSTER_CREATION_IN_PROGRESS severity INFO host_name nfvis ! !	All nodes
Cluster creation success	notification eventTime 2022-01-27T12:49:24.068755+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status SUCCESS status_code 0 status_message Cluster creation succeeded details Storage virtualization successful event_type CLUSTER_CREATION_SUCCESS severity INFO host_name nfvis ! !	Leader node

Description	Notification	Seen on
Cluster creation failure	<pre>notification eventTime 2022-02-02T19:43:38.281773+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status FAILURE status_code 0 status_message Cluster creation failed details etcd cluster creation failed ETCD cluster health is not healthy event_type CLUSTER_CREATION_FAILURE severity INFO host_name nfvis !</pre>	All nodes
Cluster deletion in progress	<pre>notification eventTime 2022-01-27T04:42:08.654507+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status SUCCESS status_code 0 status_message Cluster deletion in progress details NA event_type CLUSTER_DELETION_IN_PROGRESS severity INFO host_name nfvis !</pre>	All nodes
Cluster deletion success	<pre>notification eventTime 2022-01-27T04:42:12.768702+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status SUCCESS status_code 0 status_message Cluster deletion succeeded details Cluster deletion successful event_type CLUSTER_DELETION_SUCCESS severity INFO host_name nfvis !</pre>	Node where deletion succeeded

Description	Notification	Seen on
Cluster deletion failure	<pre> notification eventTime 2021-12-03T11:35:55.477465+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status FAILURE status_code 0 status_message Cluster deletion failed details ['Internal error'] event_type CLUSTER_DELETION_FAILURE severity INFO host_name nfvis !</pre>	Node where deletion failed
VM migration initiated	<pre> notification eventTime 2022-01-26T19:49:14.06+00:00 vmlcEvent status SUCCESS status_code 200 status_message Migration request accepted vmlcEvent migration vmlcEvent migration vm name Centos_101 status MIGRATING ! vmlcEvent migration vm name ASAv_101 status MIGRATING ! ! vmlcEvent event type VM_MIGRATE_INIT ! ! !</pre>	Destination node for migration
VM migration completed	<pre> notification eventTime 2022-01-26T19:49:46.139+00:00 vmlcEvent status SUCCESS status_code 200 status_message Migration processing complete vmlcEvent migration vmlcEvent migration vm name Centos_101 status MIGRATION-SUCCESS ! vmlcEvent migration vm name ASAv_101 status MIGRATION-SUCCESS ! ! vmlcEvent event type VM_MIGRATE_COMPLETE ! ! !</pre>	Destination node for migration

Description	Notification	Seen on
VM migration failed at source	<pre>notification eventTime 2022-01-26T00:53:19.737912+00:00 nfvisEvent user_id admin config_change false transaction_id 0 status FAILURE status_code 0 status_message Please stop Deployments ['ASAv_209', 'OTHER18', 'c8kv_209'] before migration details NA event_type MIGRATION_FAILURE severity INFO host_name 172-25-221-9 !</pre>	At source if validation fails at source

Syslogs

Description	Syslog	Sent from
Cluster creation in progress	Jan 27 04:48:34 nfvis %SYS-6-CLUSTER_CREATION_IN_PROGRESS: Cluster creation in progress	All nodes
Cluster creation success	Jan 27 04:49:24 nfvis %SYS-6-CLUSTER_CREATION_SUCCESS: Cluster creation succeeded Storage virtualization successful	Leader node
Cluster creation failed	Jan 26 21:12:34 nfvis %SYS-3-CLUSTER_CREATION_FAILURE: Cluster creation failed etcd cluster creation failed ETCD cluster health is not healthy	All nodes
Cluster deletion in progress	Jan 27 04:45:13 nfvis %SYS-6-CLUSTER_DELETION_IN_PROGRESS: Cluster deletion in progress	All nodes
Cluster deletion success	Jan 27 04:45:16 nfvis %SYS-6-CLUSTER_DELETION_SUCCESS: Cluster deletion succeeded Cluster deletion successful	Node where the deletion succeeded
Cluster delete failed	Jan 26 11:36:55 nfvis %SYS-3-CLUSTER_DELETION_FAILURE: Cluster deletion failed ['Internal error']	Node where the deletion failed
VM migration initiated	Jan 26 11:49:14 172-25-221-101 %SYS-6-VM_MIGRATE_INIT: VM migration initiation successful: Centos_101 MIGRATING ASAv_101 MIGRATING	Destination node for migration
VM migration completed	Jan 26 11:49:47 172-25-221-101 %SYS-6-VM_MIGRATE_COMPLETE: VM migration successful: Centos_101 MIGRATION-SUCCESS ASAv_101 MIGRATION-SUCCESS	Destination node for migration

Description	Syslog	Sent from
VM migration failed at source	Jan 25 16:53:19 172-25-221-9 %SYS-3-MIGRATION_FAILURE: Please stop Deployments ['ASAv_209', 'OTHER18', 'c8kv_209'] before migration	At source if validation fails at source

