



# Enable Cinder Volume Multi-attach to Multiple VNFs

---

- [Feature Summary and Revision History](#), on page 1
- [Feature Description](#), on page 2
- [Configure Multi-attach Cinder Volume](#), on page 3
- [Monitoring and Troubleshooting](#), on page 3

## Feature Summary and Revision History

### Summary Data

Applicable Product(s) or Functional Area	<ul style="list-style-type: none"><li>• P-GW</li><li>• SAEGW</li></ul>
Applicable Platform(s)	<ul style="list-style-type: none"><li>• VPC-DI</li></ul>
Feature Default	Disabled - Configuration Required to Enable
Related Changes in This Release	Not Applicable
Related Documentation	<ul style="list-style-type: none"><li>• <i>P-GW Administration Guide</i></li><li>• <i>SAEGW Administration Guide</i></li><li>• <i>Statistics and Counters Reference</i></li></ul>

**Revision History**

Revision Details	Release
In P-GW and SAEGW, the Cinder Volume Multi-attach functionality is enhanced to: <ul style="list-style-type: none"> <li>• Monitor system volume status through the <b>monitor system volume</b> CLI command</li> <li>• Modify the switchover reasons during detach multi-attach cinder volume function</li> <li>• Raise an SNMP trap notification</li> </ul>	2024.02.0
First introduced.	21.25

## Feature Description

Cinder is the OpenStack Block Storage service for providing volumes to the VNFs. Volumes are block storage devices that is attached to instances to enable persistent storage.

In P-GW, prior to OSP 16.0, operational issues developed when working with Virtual Customer Premises Equipment (VCPE) and RedHat. The Recover VM functionality brings down the VM Control Function (CF) of QvPC-DI and tries to bring it back up on a different compute host due to compute host failures. When a new CF instance comes up and redundant array of independent disks (RAID1) is formed, the active CF instance performs disk synchronization over the internet Small Computer System Interface (iSCSI) channel. This process is done block by block and iterates over the entire disk. Disk synchronization takes place over DI-LAN. When disk sizes are larger than 250GB, it takes time depending on how storage is configured, and DI-LAN network bandwidth, and traffic.

To overcome this issue, OSP16.1 is used to support the Cinder volume multi-attach . You can use this Cinder multi-attach capability to simultaneously attach volumes to multiple VNF instances.

- CF1 (Active) and CF2 (Standby) of QvPC-DI connects to the same multi-attach volume when bringing up the orchestrator.
- StarOS detects if CF1 and CF2 are connected to the same disk volume over the iSCSI channel.
- If a cinder volume multi-attach case is detected, the HD-RAID gets formed using the HD-local disk alone (disk connected to active CF). This process avoids the HD-RAID mirroring to solve the operational issues.

**Disk Failures in Multi-attach**

For disk failure in multi-attach, CF switchover is not possible as both CFs point to the same volume. If a disk failure is detected for Cinder volume multi-attach, it initiates an automatic ICSR switchover. The Interchassis Session Recovery (ICSR) setup is used to handle disk failure scenarios for Cinder volume multi-attach.

**Monitor System Volume Status**

When multi-attach cinder volume fails on the active CF card of vPGW, the monitor system volume functionality under the Service Redundancy Protocol (SRP) global configuration mode allows:

- Monitoring the system volume during volume attach and detach using a CLI command.
- Modification of switchover reason when the multi-attach cinder volume detaches from the active CF card and the SRP switch over happens.
- SNMP traps notification when the standby CF card from the active VNF detects volume detach.

## Configure Multi-attach Cinder Volume

Use the following CLI command to enable the system to monitor multi-attach cinder volume status from the active CF.

```
configure
  context context_name
    service-redundancy-protocol
      [ no ] monitor system volume
    end
```

### NOTES:

- **monitor system volume:** Enables Service Redundancy Protocol (SRP) to monitor volumes.
- **no:** Disables the volume monitoring.

## Monitoring and Troubleshooting

This section provides information on how to monitor and troubleshoot this feature using show commands.

### Show Commands and Outputs

This section provides information about show commands and their outputs for this feature.

#### show hd raid verbose

The following new field is added to the output of this command:

- HD Raid
  - Degraded—No (Multiattach)

The following is the sample output:

```
HD RAID:
State           : Available (clean)
Degraded        : No (Multiattach)
UUID            : 643094ff:2cb03262:6e34e7e0:eddbca63
Size            : 214GB (214000000000 bytes)
Action          : Idle
Disk            : hd-locall
```

```

State                : In-sync component
Created              : Tue Apr  9 09:00:17 2024
Updated             : Wed Apr 24 02:21:18 2024
Events              : 113252
Model               : QEMU QEMU HARDDISK 2.5+
Serial Number       : d6ae9d7e-12b1-4b62-93c3-1946ae882113
Location            : CFC1 A7E5A7CF-561C-48A9-8784-F9580A3A7DAD
Size                : 214.7GB (214748364800 bytes)
Disk                : hd-remotel
State               : Valid image of 643094ff:2cb03262:6e34e7e0:eddbca63
Created              : Tue Apr  9 09:00:17 2024
Updated             : Wed Apr 24 02:21:18 2024
Events              : 113252
Model               : QEMU QEMU HARDDISK 2.5+
Serial Number       : d6ae9d7e-12b1-4b62-93c3-1946ae882113
Location            : CFC2 59AE2254-E3C6-4E8C-8A94-03556C1B2EEB
Size                : 214.7GB (214748364800 bytes)

```

### show srp monitor

The **show srp monitor** command is enhanced to display the volume monitor status for multi-attach cinder volumes.

```

Auth. probe monitor state: Success
Auth. probe monitors up:  0
.....
.....
VFP monitor state:        Success
SX monitor state:         Success
Volume monitor state:     Success

```

The **show srp monitor volume** command is enhanced to display the status of multi-attach cinder volumes.

```

[local]laas-setup# show srp monitor volume
+----- Type: (A) - Auth. probe (B) - BGP (D) - Diameter (F) - BFD (E) - EGQC
|                (C) - Card (V) - VFP (S) - Sx (M) - Multiattach Cinder Volume
|
|+---- State: (I) - Initializing (U) - Up (D) - Down
||
||+---- GroupId
|||
vvv IP Addr          Port Context(VRF Name)                Last Update
-----
MI- -                - - - - - Tue Feb 06 06:04:50 2024
-----

```

### show srp call-loss statistics

The **show srp call-loss statistics** show command displays the switchover reason as **Cinder Volume failure** on volume detach in the active CF.

```

[local]CXTMVPGWVNC-Primary-11# show srp call-loss statistics
Thursday December 21 16:23:28 UTC 2023
Switchover-4 started at : Thu Dec 21 16:20:44 2023, took 1 seconds to finish.
Switchover reason : Cinder Volume failure

```



- 
- Note** When the multi-attach volume is detached from the Standby VNF, the following two situations can occur:
- An event of multi-attach volume detachment on the active VNF can cause service impact. This service impact occurs even if the auto SRP switchover is restricted to happen, because both the VNF are not having their Volumes attached.
  - Manual execution of the SRP switchover command with or without force can cause the switchover to happen. This service impact is because the peer VNF also does not have multi-attach volume attached to it.
- Hence it is recommended to rectify the HD Raid status of Peer as soon as possible. Also verify that peer HD Raid status is proper before issuing any manual SRP Switchover.
- 

## Verify SRP Switchover Reasons through SNMP Traps Notification

The standby CF card raises the following SNMP traps:

- **StorageNotFound**—volume detach
- **StorageFound**—volume attach

When the switchover occurs due to multi-attach volume detach from an active CF card, the SRPSwitchoverOccurred trap displays the reason as **Cinder Volume Failure**.

```
Internal trap notification 1278
(SRPSwitchoverOccured) vpn SRP ipaddr 2002:4888:34:13:386:200:0:11 rtmod
18 Switchover Reason: (18) Cinder Volume Failure
```



- 
- Note** When cinder volume is re-attached to the active card, it is not automatically detected by the system unless a soft reload is done.
-

