



Cinder Volume Multi-Attach

- [Feature Summary and Revision History, on page 1](#)
- [Feature Description, on page 1](#)
- [Monitoring and Troubleshooting, on page 2](#)

Feature Summary and Revision History

Summary Data

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

Revision History

Revision Details	Release
First introduced.	21.25

Feature Description

In the existing infrastructure, 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 RAID1 is formed, the active CF instance performs disk synchronization over the 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. CF1 (Active) and CF2 (Standby) of QvPC-DI connect 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 is formed using the HD-local disk alone (disk connected to active CF). Therefore, it avoids the HD-RAID mirroring to solve the operational issues.

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.

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 regarding 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 a sample output:

```

HD RAID:
State                : Available (clean)
Degraded              : No (Multiattach)
UUID                 : 3b545d34:e59553c8:a5d0ac8b:78736ca3
Size                  : 230GB (230000000000 bytes)
Action                : Idle
Disk                  : hd-locall
  State               : In-sync component
  Created              : Wed Aug 18 14:44:04 2021
  Updated              : Thu Aug 19 02:23:36 2021
  Events               : 32678
  Model                : QEMU QEMU HARDDISK 2.5+
  Serial Number        : 2b4ea4e4-948e-4641-a342-c82941f9fc98
  Location              : CFC1 A875D844-0E1E-4CA9-A4DB-FB52CB871359
  Size                  : 230.8GB (230854492160 bytes)
Disk                  : hd-remotel
  State               : Valid image of 3b545d34:e59553c8:a5d0ac8b:78736ca3
  Created              : Wed Aug 18 14:44:04 2021
  Updated              : Thu Aug 19 02:23:36 2021
  Events               : 32678
  Model                : QEMU QEMU HARDDISK 2.5+
  Serial Number        : 2b4ea4e4-948e-4641-a342-c82941f9fc98

```

```
Location          : CFC2 7CEB1976-97FE-4E75-A711-9FDC51EF7617
Size              : 230.8GB (230854492160 bytes)
```

