
ETERNUS OpenStack Volume Driver Release Note

February 2023

Contents

ABOUT THIS DOCUMENT	3
NOTES	3

Trademarks

Red Hat is a trademark of Red Hat, Inc., registered in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

The company names, product names and service names mentioned in this document are registered trademarks or trademarks of their respective companies.

About this document

This document provides complementary information not described on the ETERNUS OpenStack Volume Driver

Notes

1. Notice about Volume Driver feature (Except for notice related to only ETERNUS OpenStack Volume Driver)

1. Notice about detaching volume soon after attaching volume

When detaching volume soon after attaching volume, detach volume sometimes failed by following bug. Please wait for a while after attaching volume and detach volume.

<https://bugs.launchpad.net/nova/+bug/1430553>

2. Notice about multiple execution of Migrate Volume

If you execute Migrate Volume in multiple, the tasks sometimes fail. Executable multiplicity depends on OpenStack version, system configuration, and server specs.

In case of configuring system which executes multiple Migrated Volume, please confirm executable multiplicity before starting system operation.

3. Notice about Retype Volume

If you execute Retype Volume for volumes without volume type, the operation may fail. Although we confirmed such failure in Liberty, it may happens in other version.

<https://bugs.launchpad.net/cinder/+bug/1547546>

When backend for destination of retype failed to be launched because of incorrect configuration and so on, Retype Volume failed. Since in this timing, cinder does not output log message which show retype failure clearly, you need check update of host information in the volume using command such as cinder show.

4. Notice about Migrate Volume/Retype Volume for the volume created by using "Create Volume from Snapshot" or "Clone Volume"

In OpenStack Liberty version / Mitaka version, when executing Migrate Volume / Retype Volume for the volume created by using "Create Volume from Snapshot" or "Clone Volume", the operation does not work correctly. As for later version of OpenStack, same problem may happen. Please check latest community information.

<https://bugs.launchpad.net/cinder/+bug/1580013>

5. Notice about Extend Volume

When executing Extend Volume on Raid Group, ETERNUS OpenStack Volume Driver will use LUN Concatenation function of ETERNUS AF/DX. Execute Extend Volume repeatedly may fail due to LUN concatenation limits. If it fails, use Raid Migration of volume on ETERNUS GUI to resolve the LUN Concatenation state.

Other functions using "Extend Volume" also have the above limitations.

6. Notice about Copy volume to image

When set glance backend to cinder and create an image from cinder backend, if the parameter `image_upload_use_cinder_backend` is set to True, the image will be created by clone volume, so the disk format must be raw and bare. If the parameter `image_upload_use_cinder_backend` is set to False, please specify a Thin Provisioning Pool for StoragePool of the backend that image volume will be created on. When specify a Raid Group, glance will create a 1GB volume and extend it by 1GB repeatedly with Extend Volume, which may fail because of the LUN concatenation limit. Please refer to the following for details.

<https://docs.openstack.org/cinder/latest/admin/volume-backed-image.html>

2. Notice in Red Hat Enterprise Linux 7 with iSCSI

1. Notice about the reboot of compute node

When shutting down or restarting Compute Node running on Red Hat Enterprise Linux 7 with iSCSI, please do so after all VMs have been migrated to other Compute Node. If not so, the shutdown of the Compute Node sometimes does not complete normally.

2. Message about failure of data transfer via iSCSI connection

When data transfer via iSCSI connection failed in the operation accessing volumes on ETERNUS AF/DX from Controller node/Compute node such as Attach Volume, Copy Image to Volume, and Migrate Volume, please read system log messages in node where the trouble occurred. (In case of RHEL7.0, /var/log/messages)
If following messages are output in the log, iSCSI initiator service in the node may have a problem. Please solve the problem manually or reboot the node.

```
iscsid: Kernel reported iSCSI connection 23:0 error (1022 - Invalid or unknown error code)
state (3)
kernel: sd 29:0:0:0: rejecting I/O to offline device
kernel: sd 29:0:0:0: [sdxx] killing request
```

3. Backend name

Please do not use backend name which consists of same characters but different upper/lower case because some Cinder processing ignores the case difference. And, in case of using backend name which consists of same characters as the name used before but different upper/lower case, volume creation usually fails.

Example:

1. Edit /etc/cinder/cinder.conf as below and start Cinder

```
[DEFAULT]
enabled_backends=BACKEND1
```

...

```
[BACKEND1]
```

```
volume_driver = cinder.volume.drivers.fujitsu.eternus_dx_iscsi.FJDXISCSIDriver
cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml
```

2. Edit /etc/cinder/cinder.conf as below and restart Cinder

```
[DEFAULT]
enabled_backends=backend1
```

...

```
[backend1]
```

```
volume_driver = cinder.volume.drivers.fujitsu.eternus_dx_iscsi.FJDXISCSIDriver
cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml
```

3. Create Volume ⇒ Volume create is failed.

4. Notice about device-mapper multipath environment

When attaching/detaching volume is done repeatedly in device-mapper multipath environment, the relation between multipath device also known as mpathXX and SCSI device also known as /dev/sdX is rarely inconsistent. Once this inconsistency occurs, attaching volume can fail.

Example) Multipath information without /dev/sdX is displayed.

```
# multipath -ll
3600000e00d2b0000002b010800000000 dm-0
size=1.0G features='0' hwhandler='0' wp=rw
```

Example) Message “alua not supported” is displayed.

```
# multipath -ll
Sep 11 11:37:27 | sdx: alua not supported
Sep 11 11:37:27 | sdy: alua not supported
```

In this case, please recover the consistency by using any of following way

- a) Reboot the service where consistency of multipath occurred
 - b) Solve the consistency with following operations
 - 1) Check the construction of multipath

```
# multipath -ll
```
 - 2) Delete multipath device having consistency

```
# multipath -f <device>
```
 - 3) When /dev/sdX consisting multipath device remains, delete it

```
# echo "1" > /sys/block/sdX/device/delete
```
- (*) sdX: the name of device consisting multipath

5. Notice about shutdown down or reboot of ETERNUS AF/DX

When ETERNUS AF/DX shutdowns or reboots by trouble such as hardware failure, hang-up may happen in the connection between ETERNUS DX and ETERNUS OpenStack Volume Driver. After this hang-up occurs, ETERNUS OpenStack Volume Driver does not receive any operation. If this happening occurs, please reboot OpenStack Cinder.

6. Notice about migration suspend when Cinder reboots during migrating volume

If Cinder reboots when volume is being migrated by Migrate Volume, the migration is sometime suspended. Then, please delete volume as destination of migration and retry the migration.

Please note that delete operation may fail because of remaining migration status in volume information. Then, please ask system administrator to clear migration status of volumes as source and destination of migration. When Cinder or other OpenStack components which you use provide abort feature for migration, you can use it instead of clearing migration status.

Example of failure of delete operation)

```
# cinder delete <volume>
```

```
Delete for volume <volume> failed: Invalid volume: Volume cannot be deleted while migrating (HTTP 400) (Request-ID: req-xxx)
```

```
ERROR: Unable to delete any of the specified volumes.
```

Example of operation for clearing migration status)

```
# mysql -u root
```

```
MariaDB [(none)]> use cinder
```

```
MariaDB [cinder]> UPDATE volumes SET migration_status=NULL WHERE  
id='<volume id>';
```

* Commands are different depending on DB and system which you use

7. Notice about Retype Volume operation in RDO Mitaka

When using older OpenStack Volume Driver, we found that Retype Volume do not work correctly in following packages about RDO.

Mitaka:

- rdo-release-mitaka-1.noarch.rpm
- rdo-release-mitaka-2.noarch.rpm
- rdo-release-mitaka-3.noarch.rpm

* The trouble may happen when you use later version RDO or other distribution as well

The problem is below, and please consider this problem when you use Retype Volume.

- (1) Although Retype Volume completes successfully, volume type is not changed.
 - (2) When you execute retype with "--migration-policy on-demand" option, volume is successfully migrated to specified ETERNUS AF/DX, although volume type is not changed.
-

8. Notice in Red Hat Openstack Platform 10

When using OpenStack Volume Driver 1.4 in the Redhat Openstack Platform 10 environment, the installation destination, cinder.conf contents are different.

①installation directory

`/usr/lib/python2.7/site-packages/cinder/volume/drivers/fujitsu_140`

②cinder.conf

- `volume_driver`

- For FC connection

```
volume_driver = cinder.volume.drivers.fujitsu_140.eternus_dx_fc.FJDXFCDriver
```

- For iSCSI connection

```
volume_driver = cinder.volume.drivers.fujitsu_140.eternus_dx_iscsi.FJDXISCSIDriver
```

9. Notice about update to OpenStack Volume Driver 1.5

The following procedure is required when updating from the old version OpenStack Volume Driver to OpenStack Volume Driver 1.5.

- It is necessary to change the host affinity setting on ETERNUS AF/DX so that only one LUN group is assigned to the same host.

-important notice

- ① Make sure that the volumes in the LUN group assigned to the same host are the same before changing the host affinity setting is the same.
- ② Do not operate the Attach volume and Detach volume when changing the host affinity setting.
- ③ When changing the host affinity setting, do not perform Host IO on the VM

Specific examples are shown below.

■ Before the change

For one HOST, four LUN groups are assigned to each of four ports.

HOST: HOST_NAME # 1,

LUN groups: AFIN_GRP_ # 0, AFIN_GRP_ # 1, AFIN_GRP_ # 2, AFIN_GRP_ # 3

Ports: CM#0 CA#1 Port#0, CM # 0 CA # 1 Port # 1, CM # 1 CA # 1 Port # 0, CM # 1 CA # 1 Port # 1

■ After the change

For one HOST, one LUN group Assigned to 4 ports.

●HOST:HOST_NAME # 1

●LUN group: AFIN_GRP_ # 0

●Ports : CM # 0 CA # 1 Port # 0, CM # 0 CA # 1 Port # 1, CM # 1 CA # 1 Port # 0, CM # 1 CA # 1 Port # 1

■ Change procedure

1. Check the Volume of LUN group corresponding to HOST.

```
CLI> show lun-groups-lg-name AFIN_GRP_ # 0
```

```
CLI> show lun-groups-lg-name AFIN_GRP_ # 1
```

```
CLI> show lun-groups-lg-name AFIN_GRP_ # 2
```

```
CLI> show lun-groups-lg-name AFIN_GRP_ # 3
```

Confirm that the volume of each LUN group is same from each ETERNUS CLI command result.

2. Delete host affinity other than CM # 0 CA # 1 Port # 0 with the following ETERNUS CLI command.

```
CLI> release host-affinity-host-name HOST_NAME # 1 -port 011, 110, 111
```

3. Set host affinity of AFIN_GRP_ # 0 for PORT other than CM # 0 CA # 1 Port # 0.

```
CLI> set host-affinity -host-name HOST_NAME#1 -port 011,110,111 -lg-name AFIN_GRP_#0
```

4. Confirm whether the LUN group assigned to the host is one by the following ETERNUS CLI command.

```
CLI> show host-affinity -host-name HOST_NAME#1
```

10. Notice about failover function in OpenStack Volume Driver version 1.5

Please do not execute Detach Volume of failover function in over 25 multiples.

If problems with more 25 multiple operations occur on the connection with ETERNUS AF/DX, inconsistency occurs in relationship between TFO group, and thereafter, the Attach Volume of failover function may not operate properly.

In this case, please recover the consistency by using following way.

1)Please remove all TFO volumes once from the LUN group and add them again to the LUN group on Secondary Storage. When adding again, please match the corresponding volume of the primary storage with the LUN number.

For more details on this procedure,
refer to the ETERNUS SF manuals or the ETERNUS CLI manuals.

11. Notice about update to OpenStack Volume Driver version 1.6.1 or later

OpenStack Volume Driver 1.6.1 support Queens, Rocky, Stein and Train. When updating from OpenStack Volume Driver 1.5 in Queens to OpenStack Volume Driver 1.6.1 or later, the host information of the volume that was created in the Cinder database must be updated. For more details on this procedure, refer to the Installation section of the ETERNUS OpenStack Volume Driver 1.6 User's Guide.

"Copy Image to Volume" uses Cinder caching function and does not use any additional features. Image volumes created with OpenStack Volume Driver 1.5 cannot use with OpenStack Volume Driver 1.6.

12. Notice about creating and deleting volumes in a short time

If the amount of processing of requests to ETERNUS AF/DX increases as a result of executing "Create Volume" and "Delete Volume" repeatedly and consecutively, the processing response will be delayed and the system will become unresponsive. For improvement, please update to OpenStack Volume Driver 1.6 or higher, edit cinder.conf as below and restart Cinder.

```
fujitsu_concurrent_create_num = 100
```

This value is the max number of simultaneously creation volumes with OpenStack Volume Driver. If this limit is exceeded, the create process is rejected and an error response is issued.

13. Notice in Red Hat Openstack Platform 13 / 15/ 16

When using Openstack Volume Driver in a Red Hat Openstack Platform 13, 15 or 16 environment, the Fujitsu ETERNUS Back End must be set.

When using Openstack Volume Driver 1.6.1 in Red Hat Openstack Platform 13, refer to the following guide for installation.

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/13/html/fujitsu_eternus_back_end_guide/index

When using Openstack Volume Driver 1.6.1 in Red Hat Openstack Platform 15, refer to the following guide for installation.

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/15/html/fujitsu_eternus_back_end_guide/index

When using Openstack Volume Driver 1.6.2 in Red Hat Openstack Platform 16.0, refer to the following guide for installation.

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/16.0/html/fujitsu_eternus_back_end_guide/index

When using Openstack Volume Driver 1.7.2 in Red Hat Openstack Platform 16.2, refer to the following guide for installation.

https://www.fujitsu.com/downloads/JP/jstrg/download/openstack-vd/deployguide_osvd17for_rhosp162_en.pdf

14. Notice on configuring DriverFilter

When using DriverFilter for Cinder scheduler, edit the following options for each backend in Cinder.conf.

- filter_function
- fujitsu_use_eternus_snap_pool
- fujitsu_create_retry_wait_time

① filter_function

Specifies driver filter conditions for the schedule. The condition specified by this parameter is valid only when DriverFilter is set in scheduler_default_filters. For example, when creating a volume with a RAID group, the volume cannot be created in the following situations.

- 1) Maximum continuous physical free space is smaller than the capacity of the created volume.
- 2) 128 volumes have already been created in the RAID group.

For the above reasons, when updating pool information with Get Volume Stats, also update "total_volumes" (the number of volumes already created in the RAID group) and "fragment_capacity_gb" (maximum continuous physical capacity). If the user sets the filter_function as shown below, the Cinder schedule will select the pool according to the set driver filter conditions.

```
filter_function = "capabilities.total_volumes <128 and
capabilities.fragment_capacity_mb/1024> = volume.size"
```

For details, refer to the following page.

<https://docs.openstack.org/cinder/latest/admin/blockstorage-driver-filter-weighing.html>

Note:

When using DriverFilter, the following restrictions apply:

- When creating a volume in Pike, you need to specify the volume type with the --volume-type option.
 - When expanding the volume capacity, filter_function is also applied. Since volume.size of filter_function is the capacity before expansion of the specified volume, False may be returned by filter_function even if there is continuous free space in the RAID group. For example, if the volume capacity before expansion is 100GB and the maximum continuous free space of the RAID group is 80GB, when expanding the volume capacity with the command "cinder extend <volume-id> 150", the maximum continuous free space of the RAID group is 80GB. Since is less than the volume capacity before expansion (100GB), the Cinder schedule cannot select the pool and the capacity expansion will fail.
 - "Total_volumes" and "fragment_capacity_gb" are not updated except for regular Get Volume Stats. Therefore, if you create volumes / snapshots continuously before updating the pool information with Get Volume Stats, "total_volumes" and "fragment_capacity_gb" will not be updated, and the filter_function judgment may be invalid.
-

- "Total_volumes" and "fragment_capacity_gb" are attribute information of RAID group.

Therefore, if the backend has a thin provisioning pool (TPP), filter_function cannot be specified.

② fujitsu_use_eternus_snap_pool (After Train)

When taking a snapshot, the Cinder scheduler selects the pool where the original volume was created. This parameter specifies the pool in which to take the snapshot. The default value is True before OpenStack Volume Driver 1.7.2. The default value is False after OpenStack Volume Driver 1.7.2. If True, take a snapshot in the pool specified by EternusSnapPool. If False is specified, the snapshot will be taken with the pool of the original volume.

③ fujitsu_create_retry_wait_time

Specify the waiting time until retry is executed when an error occurs due to insufficient capacity after creating a volume or reaching the upper limit of the number of volumes that can be created in a RAID group. When creating Volume continuously, "total_volumes" and "fragment_capacity_mb" may not be updated. After waiting for a while and performing another retry, the above two values were updated with regular "Get Volume Stats", so the Cinder schedule side can select an appropriate pool. The default value is 40 (unit: seconds). Possible values are 0 to 60.

15. Notice about mapping setting in OpenStack Volume Driver 1.5/1.6.1/1.6.2/1.7.0

For each backend in cinder.conf, remove the setting for "fujitsu_port_host_mapping_file" and set it to "fujitsu_target_port_use_number". The value of " fujitsu_target_port_use_number" must be greater than or equal to the number of CA ports installed in the ETRENUS DX. ETERNUS OpenStack Volume Driver set the Host Affinity for all CA ports, cannot set only specific CA ports.

16. Notice when using mapping file

When attaching a volume using mapping file, do not set up different host and different LUN Group with the same CA port when setting multiple LUN groups with multiple hosts on the same node and CA ports in host affinity and mapping file.

For example, the mapping file is set as the following.

```
PortX=WWN1(host1 on computeA)=LG1
```

```
PortY=WWN2(host1 on computeB)=LG1
```

In this case, the following paths not listed in mapping file will be scanned, which may cause Nova to attach the wrong LUN.

```
WWN1-PortY
```

```
WWN2-PortX
```

To avoid Nova to attach the wrong LUN, do not set the following host affinity on ETERNUS AF/DX.

```
WWN1-PortY-LG2
```

```
WWN2-PortX-LG2
```

This document is devoted to providing technical information. The contents of this document may be modified without any prior notice.

Please contact FUJITSU LIMITED if you find any error in descriptions.

FUJITSU LIMITED is not responsible for indemnity that might be caused by the contents in this documentation or any damage related to contents in this documentation.

FUJITSU LIMITED

<https://www.fujitsu.com/global/products/computing/storage/>
