

ETERNUS OpenStack VolumeDriver 1.6 User's Guide

This page is intentionally left blank.

Preface

This manual provides an overview and explains how to use the ETERNUS OpenStack VolumeDriver. This manual is designed for users with a knowledge of the functions and the terms for OpenStack and ETERNUS AF/DX storage systems.

For information about the basic configuration and operation, refer to documents such as "OpenStack Cloud Administrator Guide".

Third Edition
May 2022

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 Manual

Organization

This manual is composed of the following ten chapters and an appendix:



- Chapter 1 Overview
This chapter provides an overview of the ETERNUS OpenStack VolumeDriver.
- Chapter 2 Operating Environment
This chapter provides support information, required resources, and preliminary settings for the ETERNUS OpenStack VolumeDriver.
- Chapter 3 Installation
This chapter provides the installation procedure for the ETERNUS OpenStack VolumeDriver.
- Chapter 4 Uninstallation
This chapter provides the uninstallation procedure for the ETERNUS OpenStack VolumeDriver.
- Chapter 5 How the ETERNUS OpenStack VolumeDriver Works
This chapter explains how the ETERNUS OpenStack VolumeDriver works.
- Chapter 6 Using the Periodic Monitoring Tool
This chapter explains how to configure, start, and stop the periodic monitoring tool of the ETERNUS OpenStack VolumeDriver.
- Chapter 7 Restrictions and Cautions
This chapter provides restrictions and cautions for the ETERNUS OpenStack VolumeDriver.
- Chapter 8 Message
This chapter provides ETERNUS OpenStack VolumeDriver messages, their meanings, and troubleshooting.
- Chapter 9 Obtaining the Required Information for Investigating Failures
This chapter explains how to obtain the required information for investigating a failure when an error occurs with the ETERNUS OpenStack VolumeDriver.
- Chapter 10 ETERNUS OpenStack VolumeDriver License
This chapter explains the ETERNUS OpenStack VolumeDriver license.

An appendix with "Driver Configuration File" is provided.

Naming Conventions

Symbols Used in This Manual

The following symbol is used throughout this manual:

- | | |
|---|---|
|  | This indicates information for the user to note when using the ETERNUS OpenStack VolumeDriver. Please be sure to read this information. |
|  | This indicates supplementary and reference information. |

Other Names

- "ETERNUS AF/DX" refers to the All-Flash Arrays and Hybrid Storage Systems that are supported by this product.
- "ETERNUS Web GUI" refers to the Web GUI of an ETERNUS AF/DX.
- "ETERNUS CLI" refers to the CLI of an ETERNUS AF/DX.
- Trademark symbols such as ™ and ® are omitted in this document.

Table of Contents

Chapter 1	Overview	9
1.1	ETERNUS OpenStack VolumeDriver	9
1.2	Release Information	9
1.3	Incompatible Information	12
Chapter 2	Operating Environment	13
2.1	Support Information	13
2.2	Required Resources	14
2.3	Required Packages	14
2.4	Preliminary Settings	15
Chapter 3	Installation	17
3.1	Installation Procedure	17
3.1.1	Installation Procedure (Non-Interactive Mode)	18
3.1.2	Installation Procedure (Interactive Mode)	22
3.2	cinder.conf	28
Chapter 4	Uninstallation	30
4.1	Uninstallation Procedure	30
Chapter 5	How the ETERNUS OpenStack VolumeDriver Works	31
5.1	Supported Functions of the ETERNUS OpenStack VolumeDriver	32
5.2	Supplementary Information for the Supported Functions	38
5.2.1	QoS Settings	38
5.2.2	The Initialization Function During a Volume Deletion	40
5.2.3	Mapping File	41
5.2.4	Failover Host Function	43
5.2.5	Specification of the Snapshot Creation Destination Pool	45
5.2.6	Using Self-Signed HTTPS Certificates	46

Chapter 6	Using the Periodic Monitoring Tool	47
6.1	Monitoring Tool Configuration File	47
6.2	Starting and Stopping the Monitoring Tool	48
Chapter 7	Restrictions and Cautions	49
7.1	Restriction	49
7.2	Caution	49
Chapter 8	Message	51
Chapter 9	Obtaining the Required Information for Investigating Failures	61
Chapter 10	ETERNUS OpenStack VolumeDriver License	62
Appendix A	Driver Configuration File	63

List of Tables

Table 1.1	Release information for each ETERNUS OpenStack VolumeDriver version	9
Table 1.2	Incompatible information of ETERNUS OpenStack VolumeDriver.....	12
Table 2.1	Disk drive capacity that is required in the installation destination directory	14
Table 5.1	ETERNUS OpenStack VolumeDriver support functions.....	32
Table 5.2	Cases where Extend Volume cannot be executed and the workarounds.....	36
Table 5.3	Parameter set in cinder.conf (1).....	41
Table 5.4	Parameter set in cinder.conf (2).....	43
Table 5.5	Parameters set in cinder.conf (3)	46

Chapter 1

Overview

This chapter provides an overview of the ETERNUS OpenStack VolumeDriver.

1.1 ETERNUS OpenStack VolumeDriver

The ETERNUS OpenStack VolumeDriver enables an ETERNUS AF/DX to be used as OpenStack BlockStorage. The ETERNUS AF/DX supports Fibre Channel (FC) and iSCSI as connection protocols.

1.2 Release Information

Table 1.1 Release information for each ETERNUS OpenStack VolumeDriver version

Version	Contents
1.0.0 (Released in March 2014)	First Edition
1.0.1 (Released in April 2014)	Modified copyright information
1.1.0 (Released in July 2014)	<ul style="list-style-type: none">• Added four functions Create Volume from Volume (Clone) Create Volume from Image Create Image from Volume Extend Volume• Added support for DeviceMapper Multipath• Added support for Thin Provisioning Pool (TPP)• Modified the SMI-S method for "Create Volume from Snapshot" function Modified from "CreateReplica" to "CreateElementReplica"
1.3.0 (Released in September 2015)	<ul style="list-style-type: none">• Modified the installer (to support the non-interactive mode and multiple selections of iSCSI ports)• Added support for the Create/Delete Snapshot function of the Thin Provisioning Pool (TPP)• Added a script to obtain data for investigations• Modified the driver so that a single backend can handle multiple storage pools

Version	Contents
<p>1.4.0 (Released in June 2016)</p>	<ul style="list-style-type: none"> • Added support for Migrate Volume • Added support for Retype Volume under certain conditions (When Migrate Volume is used by specifying the "--migration-policy on-demand" option) • Added support for Manage Volume • Added support for Unmanage Volume • Eliminated the "thin_provisioning_support" option and added support for the "max_over_subscription_ratio" option • Added support for overwrite installations
<p>1.5.0 (Released in July 2018)</p>	<ul style="list-style-type: none"> • Added support for the mapping file function • Added support for the monitoring tool function • Modified the installer (Added support for the settings related to HTTPS and eliminated the multiple iSCSI port option) • Added support for the QoS function • Added support for the initialization function during a volume deletion • Expanded the "Copy Image to Volume" function • Added support for a volume capacity expansion on RAID groups • Added support for the timeout function during communication with storage systems • Improved the capacity returned to the cinder scheduler when TPPs are used • Corrected the failure related to the upper limit of the volumes in snapshot pools when the Create Snapshot function is used • Added support for the Failover Host function (for the OpenStack Mitaka version or later) • Added support for the Multiattach function (only for the OpenStack Queens version) • Added support for connections to IPv6 environments • Added support for the use of self-signed HTTPS certificates

Version	Contents
1.6.0 (Released in July 2019)	<ul style="list-style-type: none">• Added support for Manage Snapshot• Added support for Unmanage Snapshot• Added support for Get Manageable Volume• Added support for Get Manageable Snapshot• Added support for Revert Snapshot to Volume• Modified the Manage Volume function• Deleted some options of the expanded "Copy Image to Volume" function• Modified the multiple pool function
1.6.1 (Released in January 2020)	<ul style="list-style-type: none">• Modified the QoS function• Added support for the following parameters<ul style="list-style-type: none">- fujitsu_concurrent_create_num- fujitsu_copy_wait_time• Modified the following parameter name<ul style="list-style-type: none">- fujitsu_ssl_verification- fujitsu_ca_certs_path• Added support for simultaneously registering a host as an IPv4 and IPv6 host, and for allocating volumes

1.3 Incompatible Information

Table 1.2 Incompatible information of ETERNUS OpenStack VolumeDriver

Version	Contents
1.0.0	First Edition
1.0.1	None
1.1.0	None
1.3.0	Deleted the create pool function. A RAID group or a Thin Provisioning Pool (TPP) must be created in advance to be specified as a storage pool.
1.4.0	Eliminated the "thin_provisioning_support" option and added support for the "max_over_subscription_ratio" option.
1.5.0	<p>Eliminated the multiple iSCSI port option, "ETERNUSISCSIIP". Because the module storage directory is changed, the backend information in cinder.conf must be modified as shown below when upgrading from earlier versions of the ETERNUS OpenStack VolumeDriver to version 1.5.0.</p> <ul style="list-style-type: none"> • volume_driver=cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_fc.FJDXFCDriver • volume_driver=cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_iscsi.FJDXISCSIDriver
1.6.0	<p>Obsoleted the following options of the "Copy Image to Volume" function.</p> <ul style="list-style-type: none"> • fujitsu_image_management_dir • use_fujitsu_image_volume • fujitsu_min_image_volume_per_storage <p>These functions are supported as OpenStack standard functions.</p> <p>Obsoleted the following option as a result of modifying the multiple pool function.</p> <ul style="list-style-type: none"> • fujitsu_scheduler
1.6.1	<p>When ETERNUS OpenStack VolumeDriver 1.5.0 is upgraded to 1.6.0 or later, the following parameter names in cinder.conf must be renamed:</p> <ul style="list-style-type: none"> • "verification" to "fujitsu_ssl_verification" • "ca_certs_file" to "fujitsu_ca_certs_path"

Chapter 2

Operating Environment

This chapter provides support information, required resources, and preliminary settings for the ETERNUS OpenStack VolumeDriver.

2.1 Support Information

For the supported versions of OpenStack and the latest information about the ETERNUS AF/DX, refer to the following download website.

<https://www.fujitsu.com/global/support/products/computing/storage/download/openstack-vd.html>

■ Supported functions

Supported functions when a RAID group or a Thin Provisioning Pool (TPP) is specified for the storage pool are shown below.

- Create Volume
- Delete Volume
- Attach Volume
- Detach Volume
- Create Snapshot
- Delete Snapshot
- Create Volume from Snapshot (*1)
- Get Volume Stats
- Clone Volume (*1)
- Copy Image to Volume
- Copy Volume to Image
- Extend Volume
- Migrate Volume
- Retype Volume (*2)
- Manage Volume
- Unmanage Volume
- Manage Snapshot
- Unmanage Snapshot
- Get Manageable Volume
- Get Manageable Snapshot
- Revert Snapshot to Volume
- Failover Host (*1)
- Multiattach

*1: Not supported in the ETERNUS DX S2 series.

*2: For retypes that are accompanied by changes in Extra Specs, Encryption, QoS Specs, or volume_backend_name, the "--migration-policy on-demand" option must be specified.

2.2 Required Resources

■ Disk drive capacity

To install the ETERNUS OpenStack VolumeDriver, the following disk drive capacity is required in the installation destination directory.

Table 2.1 Disk drive capacity that is required in the installation destination directory

No.	Directory	Disk drive capacity	Contents
1	Cinder installation directory /volume/drivers	400KByte	Driver
2	Storage directory of the driver configuration file (e.g. /etc/cinder)	300Byte	Driver configuration file

2.3 Required Packages

- To operate the ETERNUS OpenStack VolumeDriver, the Python library that is described below is required. Make sure to install the required package in advance.
 - PyWBEM
For details about the supported PyWBEM versions, refer to the download site.
- For details about the required packages when an FC connection is used, refer to "Fibre Channel support in Compute" in "OpenStack Configuration Reference".

2.4 Preliminary Settings

■ Licenses

The Advanced Copy Feature license is required to create snapshots, create volumes from snapshots, or clone volumes. Purchase this license separately and register the license in the ETERNUS AF/DX. Note that the Advanced Copy table size setting is also required. For details on registering and configuring the Advanced Copy function, refer to the ETERNUS Web GUI manuals.

In addition to the Advanced Copy Feature license, ETERNUS SF Storage Cruiser (the ETERNUS SF Storage Cruiser Storage Cluster option) is required to use the Failover Host function. Purchase the Storage Cluster option license separately and register the license in the ETERNUS AF/DX. For details on the settings, refer to the ETERNUS SF manuals.

■ SMI-S settings for the ETERNUS AF/DX

The ETERNUS OpenStack VolumeDriver configures an ETERNUS AF/DX via SMI-S. The SMI-S of the ETERNUS AF/DX must be enabled in advance.

Enable the SMI-S of the ETERNUS AF/DX by using ETERNUS Web GUI or ETERNUS CLI. For more details on this procedure, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

■ SSH server key settings for the ETERNUS AF/DX

To use the ETERNUS OpenStack VolumeDriver, the SSH server key must be set in the ETERNUS AF/DX.

Use ETERNUS Web GUI to set the SSH server key. For details on the settings, refer to the ETERNUS Web GUI manuals.

■ Account used for the ETERNUS AF/DX

For connections via SMI-S and ETERNUS CLI, register a user account in advance. To access the ETERNUS AF/DX via SMI-S, a user account with "Software" as a user level is required. In addition, the SSH public key must be uploaded for the user account.

Use ETERNUS Web GUI or ETERNUS CLI to register ETERNUS AF/DX user accounts. For more details on the registration procedure, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

Caution

Account information is entered in the driver configuration file. Setting a password that is not used for other systems is recommended since the driver configuration file is not encrypted.

■ Creating the storage pool

Create a storage pool for creating volumes in advance.

A RAID group or a Thin Provisioning Pool (TPP) can be specified for the storage pool.

Use ETERNUS Web GUI or ETERNUS CLI to create a RAID group or a Thin Provisioning Pool (TPP) in the ETERNUS AF/DX. For more details on the creation procedure, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

Create a storage pool for creating snapshots in advance.

RAID groups or TPPs can be specified as a storage pool.

■ Connecting the LAN for operation management

A LAN connection with an operation management LAN is required to access the ETERNUS AF/DX from the ETERNUS OpenStack VolumeDriver. In the same way as ETERNUS Web GUI or ETERNUS CLI, connect the controller node (Cinder) of OpenStack to the ETERNUS AF/DX via the operation management LAN. For more details, refer to the site planning manuals for the ETERNUS AF/DX that is to be used.

■ SAN connection

FC and iSCSI can be used as a host interface.

The compute node of OpenStack and the ETERNUS AF/DX must be connected to the SAN and be available for communication in advance. To use Fibre Channel switches, zoning settings for the Fibre Channel switches are also required. To use the iSCSI connections, logging in to the iSCSI target is required.

The host affinity mode for all of the host interface ports of the ETERNUS AF/DX must also be enabled in advance. For more details, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

■ IPv6 connection

Configure IPv6 IP addresses for the ETERNUS AF/DX management port and CA port for connections from the ETERNUS OpenStack VolumeDriver to the ETERNUS AF/DX via IPv6. For more details, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

■ Enabling the Snap Data Pool

The snapshot function of the ETERNUS OpenStack VolumeDriver requires the Snap Data Pool (SDP) when specifying a RAID group as a storage pool. Therefore, the SDPV must be created and the SDP must be enabled.

For more details on SDPVs, refer to the ETERNUS Web GUI manuals or the ETERNUS CLI manuals.

Chapter 3

Installation

This chapter provides the installation procedure for the ETERNUS OpenStack VolumeDriver.

3.1 Installation Procedure

Note

- If the installer terminates with an error, uninstall the ETERNUS OpenStack VolumeDriver and install it again.
- An "&" cannot be used as a parameter in the driver configuration file. In addition, characters that are not supported by the ETERNUS AF/DX cannot be used. For details, refer to the ETERNUS Web GUI manuals.
- The package file name varies for the ETERNUS OpenStack VolumeDriver depending on the version used. The version of ETERNUS OpenStack VolumeDriver is indicated by "*nnn*" and the version of OpenStack is indicated by "*xxx*".
`osvdnnn_xxx.tar.gz`
- If the ETERNUS OpenStack VolumeDriver has already been installed in the installation destination directory, an installation is performed according to the following mode options:
Non-interactive mode: An overwrite installation is performed.
Interactive mode: A confirmation message is displayed asking whether or not to overwrite the current installation. If "yes" is selected, an overwrite installation is performed.

Caution

- Note the following points when upgrading from versions earlier than 1.5.0 of the ETERNUS OpenStack VolumeDriver to version 1.5.0 or later.
For versions 1.5.0 and later of the ETERNUS OpenStack VolumeDriver, the Cinder installation directory has changed.
 - Versions earlier than 1.5.0
Cinder installation directory: `/volume/drivers/fujitsu`
 - Versions 1.5.0 and later
Cinder installation directory: `/volume/drivers/fujitsu/eternus_dx`
- After the installation, folders other than the `eternus_dx` folder under the Cinder installation directory `/volume/drivers/fujitsu` must be deleted manually.

3.1.1 Installation Procedure (Non-Interactive Mode)

The following procedure shows how to install the ETERNUS OpenStack VolumeDriver with the non-interactive mode.

Procedure

- 1 Download the ETERNUS OpenStack VolumeDriver package from the download site and transfer the obtained package to the working directory in the controller node.
- 2 Log in to the controller node and set the working directory as the current directory. Make sure to log in with the account that was used when installing OpenStack.
- 3 Extract the ETERNUS OpenStack VolumeDriver package by using the following command.
Example: For ETERNUS OpenStack VolumeDriver version 1.6.1 of the Queens release

```
$ /bin/tar xvzf ./osvd161_queens.tar.gz
```

A folder is created in the current directory. The folder name is the same as the package file name without ".tar.gz".

- 4 Go into the extracted directory and execute "osvd_install.sh install" to install the ETERNUS OpenStack VolumeDriver.

```
$ ./osvd_install.sh install <config file>
```

For <config file>, specify the file that describes the parameters.
The "cinder_driver_sample.ini" in the extracted directory is a template of the input file.

- Input file format

```
[DEFAULT] (1)
enabled_backends=FJFC,FJISCSI,FJISCSI_MAPFILE (2)
db_username=Cinder (3)
db_password=cinder (4)
#FibreChannel:fc,iSCSI:iscsi
[FJFC] (5)
fujitsu_volume_driver=true (6)
protocol=fc (7)
driver_config_file=/etc/cinder/xxx.xml (8)
ip=10.11.12.13 (9)
port=5989 (10)
user=osvd (11)
password=osvd (12)
poolname=raid1,raid2 (13)
snappoolname=raid3,raid4 (14)
volume_backend_name=FJFC (15)
fujitsu_ssl_verification=true (16)
fujitsu_ca_certs_path=/etc/pki/tls/certs/ca-bundle.trust.crt (17)

[FJISCSI] (5)
fujitsu_volume_driver=true (6)
protocol=iscsi (7)
driver_config_file=/etc/cinder/yyy.xml (8)
ip=10.11.12.14 (9)
port=5988 (10)
user=osvd1 (11)
password=osvd1 (12)
poolname=raid1 (13)
snappoolname=TPP1 (14)
volume_backend_name=FJISCSI (15)

[FJISCSI_MAPFILE] (5)
fujitsu_volume_driver=true (6)
protocol=iscsi (7)
driver_config_file=/etc/cinder/zzz.xml (8)
ip=10.11.12.15 (9)
port=5989 (10)
user= osvd2 (11)
password= osvd2 (12)
poolname= raid1 (13)
snappoolname= TPP1 (14)
volume_backend_name= FJISCSI_MAPFILE (15)
fujitsu_port_host_mapping_file=/etc/cinder/fujitsu_port_host_mapping.conf (18)
```

- (1) [DEFAULT] section that affects the entire Cinder
- (2) Section name of the backend to be used
When using multiple backends, separate them using commas.
- (3) Cinder database username
- (4) Cinder database password
If the ETERNUS OpenStack VolumeDriver is upgraded from version 1.5.0 or earlier, the host information of the volume that was created in the Cinder database must be updated. Enter the user-name and password for the Cinder database.

If the username and password are not configured, the update is skipped. If the update is skipped, make sure to update the host information after the installation is completed using the "update_hostname.sh" tool in the working directory.

```
$ sh update_hostname.sh
```

- (5) Section name
- (6) Flag that indicates the Fujitsu volume driver section
When this parameter is "true", the settings in this section are reflected in /etc/cinder/cinder.conf during installation.
- (7) Connection protocol of the ETERNUS AF/DX
- (8) Driver configuration file
- (9) IP address for the SMI-S connection
- (10) Port for the SMI-S connection
- (11) Username for the SMI-S connection
- (12) Password for the SMI-S connection
- (13) Storage pool (RAID group or Thin Provisioning Pool (TPP)) name in which volumes are created
Make sure to enter the name of the RAID group or Thin Provisioning Pool (TPP) that was created in the ETERNUS AF/DX in advance.
When specifying multiple storage pool names, separate them using commas.
Example: pool1, pool2
- (14) Storage pool (RAID group or Thin Provisioning Pool (TPP)) name in which snapshot volumes are created
When specifying multiple snapshot pool names, separate them using commas.
Optional. However, if omitted, snapshot volumes cannot be created or managed.
- (15) through (18) Other
For items other than (1) through (14), the contents are copied directly to cinder.conf.
However, other items in the [DEFAULT] section are not reflected in cinder.conf.

5 After the installation, use ETERNUS Web GUI to register the created public key (eternus.ietf) to the ETERNUS AF/DX.

For information about how to set the public key, refer to the ETERNUS Web GUI manuals.

 **Caution**

If the public key (eternus.ietf) that was created is deleted by mistake, use the following command to recreate the key.

```
$ ssh-keygen -e -f /root/.ssh/eternus.pub > ./eternus.ietf
```

6 Edit /etc/cinder/cinder.conf.

- 6-1** To collect DEBUG log level information, under the [DEFAULT] section, set the debug value to "True".

- 6-2 If the storage pool to be specified for "poolname" is a RAID group, add DriverFilter to "scheduler_default_filters" under the [DEFAULT] section. The consecutive free space or the number of the already created volumes is checked and then an available pool is selected from the multiple RAID groups that are specified.

7 Restart Cinder.

```
$ service openstack-cinder-volume restart
```

The ETERNUS OpenStack VolumeDriver is enabled.

End of procedure

Note

- Create a volume type and then associate it with the backend. An execution example is provided below.

```
cinder type-create DXFC  
cinder type-key DXFC set volume_backend_name=FJFC  
cinder type-create DXISCSI  
cinder type-key DXISCSI set volume_backend_name=FJISCSI
```

- After the installation, when updating the configuration of ETERNUS OpenStack VolumeDriver, use the following command.

```
$ ./osvd_install.sh update <config file>
```

The "cinder.conf" file is updated and a driver configuration file is created without the module being updated. The format of the input file is the same as [Step 4](#).

- If an installation or an update is performed when sections with the same name already exist in cinder.conf, the following message appears and the setting is not added to cinder.conf. Even if the section is commented out, the addition is skipped.

```
=====  
Update Cinder Configuration  
=====  
SECTION: FJFC already exists, Skip update  
SECTION: FJISCSI already exists, Skip update
```

- If the ETERNUS OpenStack VolumeDriver is upgraded from version 1.5.0 or earlier and the host information of the volume that was created in the Cinder database is not updated, clone volumes cannot be created from the volumes created with the previous version (restriction). In addition, downgrading to version 1.5.0 or earlier is not possible after volumes are created with the ETERNUS OpenStack VolumeDriver 1.6.0 or later.

3.1.2 Installation Procedure (Interactive Mode)

The following procedure shows how to install the ETERNUS OpenStack VolumeDriver with the interactive mode.

Procedure

- 1 Download the ETERNUS OpenStack VolumeDriver package from the download site and transfer the obtained package to the working directory in the controller node.
- 2 Log in to the controller node and set the working directory as the current directory. Make sure to log in with the account that was used when installing OpenStack.
- 3 Extract the ETERNUS OpenStack VolumeDriver package by using the following command.
Example: For ETERNUS OpenStack VolumeDriver version 1.6.1 of the Queens release

```
$ /bin/tar xvzf ./osvd161_queens.tar.gz
```

A folder is created in the current directory. The folder name is the same as the package file name without ".tar.gz".

- 4 Go into the extracted directory and execute "osvd_install.sh install" to install the ETERNUS OpenStack VolumeDriver.

```
$ ./osvd_install.sh install
```

The installation is performed interactively.

If the ETERNUS OpenStack VolumeDriver has already been installed, the message below is displayed. If "yes" is selected, an installation is performed after the existing ETERNUS OpenStack VolumeDriver is deleted. If "no" is selected, an installation is not performed.

```
=====
  check volume driver
=====
osvd_install.sh: INFO: This driver has already been installed.
Overwrite the driver? [yes/no]:
```

5 Enter the following items that are displayed in the screen.

```
=====
Update Cinder Configuration
=====
input enabled backends, if you don't want to change please enter to skip
enabled backends:dx1,dx2 (1)
ENABLED_BACKENDS:dx1,dx2
Are you sure? [yes/no]: yes

=====
input driver configuration
=====
section: dx1 (2)
driver configuration path (.xml): /etc/cinder/xxxx.xml (3)
input host affinity configuration file path, if you don't want to set
please enter to skip host affinity configuration file path: /etc/cinder/
fujitsu_port_host_mapping.conf (4)
input volume backend name, if you don't want to set please enter to skip
volume_backend_name:DX1 (5)
```

- (1) Section name of the backend to be used
When using multiple backends, separate them using commas.
- (2) Section name of the first backend
The settings that are described hereinafter are for the first backend.
- (3) Path for the first driver configuration file
If the specified driver configuration file already exists, the following message appears. To overwrite the existing file, select "yes". To specify a different file, select "no".

```
SECTION: dx1, specified driver configuration file (/etc/cinder/xxxx.xml)
has already existed
Overwrite the file (/etc/cinder/xxxx.xml)? [yes/no]:
```

- (4) Path for the first mapping file
Optional. If this path is omitted, a mapping file is not set.
- (5) "volume_backend_name" character string of the backend that is to be used
Optional. If the character string is omitted, "volume_backend_name" is not specified.

```
=====
select protocol
=====
1) FibreChannel
2) iSCSI
Which way of FibreChannel and iSCSI do you use?[1 or 2]: 2      (6)
osvd_install.sh: INFO: protocol is iSCSI
1)ETERNUS ip address for SMIS: 10.11.12.13      (7)
2)ETERNUS port number for SMIS:
=====
select ETERNUS port number for SMIS
=====
1) http port: 5988
2) https port: 5989
Which way of http and https do you use?[1 or 2]: 2      (8)
osvd_install.sh: INFO: https's port 5989 has been selected
Do you want to verify the SSL certificate? [yes/no]:yes      (9)
Please input CA Cert path:(default:/etc/pki/tls/certs/ca-bundle.trust.crt) :
                                                                    (10)
3)ETERNUS username for SMIS: user      (11)
4)ETERNUS password for SMIS: password      (12)
5)ETERNUS poolname (e.g. Pool01,Pool02,Pool03): pool01,pool02      (13)
6)ETERNUS snappoolname (e.g. Snappool01,Snappool02,Snappool03): pool03,pool04      (14)
```

- (6) Connection protocol of the ETERNUS AF/DX
- (7) IP address for the SMI-S connection
- (8) Port for the SMI-S connection
- (9) Confirm whether to perform an SSL authentication if HTTPS communication is specified with the SMI-S connection port
If "yes" is selected, proceed to line (10).
- (10) SSL certificate path
- (11) Username for the SMI-S connection
- (12) Password for the SMI-S connection
- (13) Storage pool (RAID group or Thin Provisioning Pool (TPP)) name in which volumes are created
Make sure to enter the name of the RAID group or Thin Provisioning Pool (TPP) that was created in the ETERNUS AF/DX in advance. When specifying multiple names, separate them using commas.
- (14) Storage pool (RAID group or Thin Provisioning Pool (TPP)) name in which snapshot volumes are created
When specifying multiple names, separate them using commas.
Optional. However, if omitted, snapshot volumes cannot be created or managed.


```
1) IP      : 10.11.12.13
2) PORT   : 5989
3) USER   : user
4) PASS   : password
5) POOL   : pool01,pool02
6) SNAPPOL : pool03,pool04
Are you sure? [yes/no]: yes           (15)
Continue to add configuration? [yes/no]:no (16)
```

(15)Checking the entered contents

Check the entered contents that are displayed and enter "yes" if there are no problems.
Entering "no" enables the contents to be re-entered.

(16)Checking for additional settings

Entering "yes" enables information about the next session to be entered.
If "no" is entered, the settings that were entered are reflected.

```
=====
Upgrade cinder database information
=====
If you do not know the username and password for accessing the database,
please check the [database] option in cinder.conf (eg: /etc/cinder/cinder.conf)
Such as: mysql+pymysql://cinder:abcd@controller/cinder
The username is "cinder" and password is "abcd".
=====
When upgrade from version 1.5.0 or earlier, should update cinder database to
change the volume's host name.
Continue to update cinder database? [yes/no]: yes           (17)
username: cinder                                           (18)
password: cinder                                           (19)
```

(17)Checking the database update

If the ETERNUS OpenStack VolumeDriver is upgraded from version 1.5.0 or earlier, update the host information of the volume that was created in the Cinder database.

Entering "yes" updates the database.

Entering "no" skips the database update.

If the update is skipped, make sure to update the host information after the installation is completed using the "update_hostname.sh" tool in the working directory.

```
$ sh update_hostname.sh
```

(18)Cinder database username

(19)Cinder database password

The following message is displayed when the installation is complete.

```
osvd_install.sh: INFO: Install is FINISHED
osvd_install.sh: INFO: Please reboot cinder services to reflect configuration done by
this installer.
```

- 6** After the installation, use ETERNUS Web GUI to register the created public key (eternus.ietf) to the ETERNUS AF/DX.

For information about how to set the public key, refer to the ETERNUS Web GUI manuals.

 **Caution**

If the public key (eternus.ietf) that was created is deleted by mistake, use the following command to recreate the key.

```
$ ssh-keygen -e -f /root/.ssh/eternus.pub > ./eternus.ietf"
```

7 Edit /etc/cinder/cinder.conf.

7-1 To collect DEBUG log level information, under the [DEFAULT] section, set the debug value to "True".

7-2 If the storage pool to be specified for "poolname" is a RAID group, add DriverFilter to "scheduler_default_filters" under the [DEFAULT] section. The consecutive free space or the number of the already created volumes is checked and then an available pool is selected from the multiple RAID groups that are specified.

8 Restart Cinder.

```
$ service openstack-cinder-volume restart
```

The ETERNUS OpenStack VolumeDriver is enabled.

End of procedure

Note

- Create and associate a volume type with the backend.
An execution example is provided below.

```
cinder type-create FJDX1
cinder type-key FJDX1 set volume_backend_name=DX1
```

- After the installation, when updating the configuration of ETERNUS OpenStack VolumeDriver, use the following command.

```
$ ./osvd_install.sh update
```

The "cinder.conf" file is updated and a driver configuration file is created without the module being updated. The format of the input file is the same as [Step 5](#).

- If an installation or an update is performed when sections with the same name already exist in cinder.conf, the following message appears and the setting is not added to cinder.conf. Even if the section is commented out, the addition is skipped.

```
=====
  Update Cinder Configuration
=====
SECTION: FJFC already exists, Skip update
SECTION: FJISCSI already exists, Skip update
```

- If the ETERNUS OpenStack VolumeDriver is upgraded from version 1.5.0 or earlier and the host information of the volume that was created in the Cinder database is not updated, clone volumes cannot be created from the volumes created with the previous version (restriction).
In addition, downgrading to version 1.5.0 or earlier is not possible after volumes are created with the ETERNUS OpenStack VolumeDriver 1.6.0 or later.

3.2 cinder.conf

The configuration file for Cinder is "cinder.conf". This section describes the parameters that are related to when the ETERNUS OpenStack VolumeDriver is used.

This configuration file is read when Cinder is started. Restart Cinder if the settings are changed.

- **volume_driver**
This item enables the ETERNUS OpenStack VolumeDriver to be read.
The specification for the driver varies depending on the connection protocol.
This item is set by the installer.
 - For FC connection
volume_driver = cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_fc.FJDXFCDriver
 - For iSCSI connection
volume_driver = cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_iscsi.FJDXISCSIDriver

- **cinder_eternus_config_file**
Location for the driver configuration file
This item is set by the installer.

Example: cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml

- **max_over_subscription_ratio**
This item specifies the ratio that is used to calculate the maximum logical capacity for the target Thin Provisioning Pool (TPP).
The logical free space can be calculated with the following formula. Capacities larger than the physical free space in the TPP can be used.
If multiple pools are defined in the driver configuration file, the logical free space is calculated per pool.

Logical free space (GB) =
"max_over_subscription_ratio" × physical capacity (GB) - used logical capacity (GB)

Specify a numerical value of 1 or larger for this item. Fractional values can also be used.
If a value is not specified, the default value that is defined internally in Cinder is used.

Example: max_over_subscription_ratio = 5.5

If 5.5 is specified as the ratio, the physical capacity of the TPP is 500GB, and the used logical capacity is 100GB, the following logical space is available for the relevant TPP.
 $500\text{GB} \times 5.5 - 100\text{GB} = 2,650\text{GB}$

- **fujitsu_target_port_use_number**
This item specifies the maximum number of iSCSI ports in the storage system that can be used by Attach Volume. If this parameter is omitted, 4 is set. If host affinity is not set for the number of ports specified, select ports for the missing number of ports, and create host affinity in the same host and LUN group as the created host affinity.

If the total number of available storage system iSCSI ports is less than or equal to the value specified for fujitsu_target_port_use_number, the host affinities are set for all available storage system iSCSI ports.

If the total number of available storage system iSCSI ports is larger than the value specified for fujitsu_target_port_use_number, the host affinities are set for the specified number of ports starting with the port with the least number of host affinities.

- **fujitsu_query_func_timeout**
This item specifies the timeout value for reference functions such as Get Volume Stats.
The default value is 60 (in seconds). The values that can be specified are 0 to 3600.
- **fujitsu_set_func_timeout**

This item specifies the timeout value for setting functions such as Create Volume. The default value is 0 (in seconds). The values that can be specified are 0 to 3600.

- `fujitsu_eternus_req_timeout`
This item specifies the receive timeout value for responses to data requests (SMI-S/CLI) from the storage system.
The default value is 1800 (in seconds). The values that can be specified are 0 to 3600.

 **Note**

This function prevents the ETERNUS OpenStack VolumeDriver from becoming unable to continue the operation when a response is not received for SMI-S/CLI requests. This function determines that a network error has occurred if a certain period of time elapses and triggers a timeout so that the system can be recovered.

 **Caution**

Note the following points when specifying timeout values.

- If the receive timeout value (`fujitsu_query_func_timeout/ fujitsu_set_func_timeout/ fujitsu_eternus_req_timeout`) is invalid, Cinder Volume Service cannot be started.
 - After the specified timeout period elapses, the connection is automatically disconnected from the Socket and an error is output to the log file.
 - If "0" is specified for the receive timeout value, a timeout does not occur.
-
- `fujitsu_timeout_retry_times`
This item specifies the number of retries that is attempted when a timeout occurs. The default value is 3 (retries). The values that can be specified are 3 to 100.
 - `fujitsu_concurrent_create_num`
This item specifies the maximum number of ETERNUS OpenStack VolumeDriver volumes that can be created concurrently. If the number of created volumes exceeds the maximum number, the creation process is rejected and an error is returned.
The default value is 0. The values that can be specified are 0 to 500.
If a non-specifiable value is specified, it operates with no limitation.
 - `fujitsu_copy_wait_time`
When the copy source is a volume or a Snapshot which is set as the copy destination of a 2-stage cascade copy, the wait time must be specified before starting the copy process. Specify the time it takes for 2-stage cascade copy to be completed for the wait time. If the copy process cannot be started even after the wait time has exceeded, an error is returned.
The default value is 0 (seconds). The values that can be specified are 0 to 3600.
If a non-specifiable value is specified, it operates with no wait time.

Chapter 4

Uninstallation

This chapter provides the uninstallation procedure for the ETERNUS OpenStack VolumeDriver.

4.1 Uninstallation Procedure

Procedure

- 1 Log in to the controller node.
- 2 Edit `cinder.conf` so that the ETERNUS OpenStack VolumeDriver is not used.
Delete the entry for the ETERNUS OpenStack VolumeDriver from `/etc/cinder/cinder.conf`.
- 3 Restart Cinder.

```
$ service openstack-cinder-volume restart
```

After restarting Cinder, the setting performed in [Step 2](#) is enabled and the ETERNUS OpenStack VolumeDriver is disabled.

- 4 Delete the ETERNUS OpenStack VolumeDriver.
Specify the "uninstall" option and execute `osvd_install.sh`.

```
$ ./osvd_install.sh uninstall
```

The following message is displayed when the uninstallation is complete.

```
osvd_install.sh: INFO: Uninstall is FINISHED  
osvd_install.sh: INFO: Please reboot cinder services to reflect configuration done by  
this installer.
```

Note

The driver configuration file is not deleted. Delete the file as necessary.

End of procedure

Chapter 5

How the ETERNUS OpenStack VolumeDriver Works

This chapter describes how the ETERNUS OpenStack VolumeDriver works.

■ Settings for the ETERNUS AF/DX

The ETERNUS OpenStack VolumeDriver configures an ETERNUS AF/DX via SMI-S. For details about SMI-S, refer to "ETERNUS SMI-S Server SMI-S API Reference".

■ Volume name

The volume name is 6 characters long with a prefix of "FJosv_". For the ETERNUS DX S2 series, it is fixed to 16 characters. For the ETERNUS AF series and the ETERNUS DX S5/S4/S3 series, it is fixed to 30 characters.

Example: FJosv_2PLuIWcWg49djrPkHC-jRw==

■ Creating RAID groups

Note that an OpenStack user interface cannot be used to create RAID groups in the ETERNUS AF/DX. Create RAID groups by using an ETERNUS AF/DX user interface (such as ETERNUS Web GUI) in advance.

5.1 Supported Functions of the ETERNUS OpenStack VolumeDriver


Table 5.1 ETERNUS OpenStack VolumeDriver support functions

Function name	Description
Create Volume	<p>Creates volumes</p> <p>Note</p> <p>The QoS function of the ETERNUS AF/DX can be used. For details, refer to "5.2.1 QoS Settings" (page 38).</p>
Delete Volume	<p>Deletes volumes</p> <p>Note</p> <p>The data initialization function can be used. For details, refer to "5.2.2 The Initialization Function During a Volume Deletion" (page 40).</p>
Attach Volume	<p>Allocates volumes to the specified host</p> <ul style="list-style-type: none"> • ETERNUS AF/DX functions <ul style="list-style-type: none"> Sets access paths to volumes <ul style="list-style-type: none"> - Adds FC hosts / iSCSI hosts - Creates / modifies LUN groups - Sets host affinities <p>Note</p> <ul style="list-style-type: none"> • The mapping file function can be used. For details, refer to "5.2.3 Mapping File" (page 41). • A host can be simultaneously registered to a storage system as an IPv4 and IPv6 host, and volumes can be allocated to that host. This function is available from the following firmware versions: <ul style="list-style-type: none"> - ETERNUS AF S3 series, ETERNUS DX S5 series, and ETERNUS DX8900 S4 V11L10 (if the firmware version is earlier than V11L30, this function cannot be set from ETERNUS Web GUI) - ETERNUS AF S2/AF series and ETERNUS DX S4/S3/S2 series (excluding ETERNUS DX8900 S4) V10L88
Detach Volume	<p>Deletes access paths to volumes</p> <ul style="list-style-type: none"> • Modifies / deletes LUN groups • Deletes host affinities

Function name	Description
Create Snapshot	<p>Creates snapshots</p> <ul style="list-style-type: none"> • Creates snapshot volumes • Starts SnapOPC <p>Note</p> <p>The destinations where snapshots are created can be specified. For details, refer to "5.2.5 Specification of the Snapshot Creation Destination Pool" (page 45).</p>
Delete Snapshot	<p>Deletes snapshots</p> <ul style="list-style-type: none"> • Stops SnapOPC • Deletes snapshot volumes
Create Volume from Snapshot	<p>Creates volumes from snapshots</p> <ul style="list-style-type: none"> • Starts OPC (from snapshot volumes to the newly created volumes)
Get Volume Stats	<p>Obtains free space information in the storage pool (or a RAID group/Thin Provisioning Pool (TPP))</p>
Clone Volume	<p>Clones volumes</p> <ul style="list-style-type: none"> • Start OPC (from the specified volumes to the newly created volumes)
Copy Image to Volume	<p>Creates a volume and copy the image managed in Glance to the volume</p> <ul style="list-style-type: none"> • ETERNUS AF/DX functions <ul style="list-style-type: none"> - Creates volumes, sets / deletes access paths - Uses Create Volume, Attach Volume, and Detach Volume • Cinder operation <ul style="list-style-type: none"> - Downloads an image from Glance and write it in the volume <p>Note</p> <p>The QoS function of the ETERNUS AF/DX can be used. For details, refer to "5.2.1 QoS Settings" (page 38).</p>
Copy Volume to Image	<p>Uploads volume data as an image to Glance</p> <ul style="list-style-type: none"> • ETERNUS AF/DX functions <ul style="list-style-type: none"> - Sets / deletes access paths - Uses Attach Volume and Detach Volume • Cinder operation <ul style="list-style-type: none"> - Reads data from a volume and upload the image to Glance
Extend Volume	<p>Changes the volume size</p>

Function name	Description
Migrate Volume	<p>Moves volumes to a different storage pool</p> <ul style="list-style-type: none"> • ETERNUS AF/DX functions <ul style="list-style-type: none"> - Creates migration destination volumes / deletes migration source volumes - Sets access paths to migration destination volumes / deletes access paths to migration source volumes - Uses Create Volume, Delete Volume, Attach Volume, and Detach Volume • Cinder operation <ul style="list-style-type: none"> - Copies data in the migration source volume to the migration destination volume
Retype Volume	<p>Changes the volume type that is set to the volume</p> <ul style="list-style-type: none"> • When the "--migration-policy on-demand" option is specified, Migrate Volume is used • No action is performed on the ETERNUS AF/DX side for cases other than above <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>Note</p> <p>For retypes that change parameters (such as Extra Specs, Encryption, QoS Specs, and volume_backend_name), the "--migration-policy on-demand" option must be specified.</p> </div>
Manage Volume	<p>Registers information for ETERNUS AF/DX volumes in OpenStack environments</p> <ul style="list-style-type: none"> • Searches volumes <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>Note</p> <ul style="list-style-type: none"> • Volumes that are to be imported can be specified by "id" and "name". <ul style="list-style-type: none"> - Specify a volume number when using "id". - Specify a volume name when using "name". • Specifying "id" is recommended since it requires a shorter processing time than "name". • Only the volumes existing in the pools that are defined in the driver configuration file can be specified. • The QoS function of the ETERNUS AF/DX can be used for the volume management. For details, refer to "5.2.1 QoS Settings" (page 38). • Volumes that satisfy all the following conditions can be specified. <ul style="list-style-type: none"> - They exist in the pools that are defined in the driver configuration file - The type is Standard/TPV - They are not snapshot volumes - They are not registered in an OpenStack environment - They do not exist in the host affinity configuration </div>

Function name	Description
Manage Snapshot	<p>Registers information for ETERNUS AF/DX snapshots in OpenStack environments</p> <ul style="list-style-type: none"> • Searches snapshots <p>Note</p> <ul style="list-style-type: none"> • Snapshot volumes that are to be imported can be specified by "id" and "name". <ul style="list-style-type: none"> - Specify a snapshot volume number when using "id". - Specify a snapshot volume name when using "name". • Specifying "id" is recommended since it requires a shorter processing time than "name". • Only the snapshot volumes existing in the pools that are defined in the driver configuration file can be specified. • Snapshot volumes that satisfy all the following conditions can be specified. <ul style="list-style-type: none"> - They exist in the pools that are defined in the driver configuration file - Snapshots are performed with SnapOPC - The snapshot source volumes are registered in an OpenStack environment - They are not registered in an OpenStack environment
Unmanage Snapshot	<p>Deletes only snapshot information that is registered in OpenStack environments without deleting snapshots from the ETERNUS AF/DX</p> <ul style="list-style-type: none"> • No action is performed on the ETERNUS AF/DX side
Get Manageable Volume	<p>Acquires the volume information that can be registered in OpenStack environments</p> <p>Note</p> <p>Displays the volumes that satisfy all the following conditions.</p> <ul style="list-style-type: none"> • They exist in the pools that are defined in the driver configuration file • They are not snapshot volumes • The type is Standard/TPV • They are not registered in an OpenStack environment • They do not exist in the host affinity configuration
Get Manageable Snapshot	<p>Acquires the snapshot information that can be registered in OpenStack environments</p> <p>Note</p> <p>Displays the snapshot volumes that satisfy all the following conditions.</p> <ul style="list-style-type: none"> • They exist in the pools that are defined in the driver configuration file • Snapshots are performed with SnapOPC • They are registered in an OpenStack environment including snapshot source volumes • The snapshot information is not registered in an OpenStack environment

Function name	Description
Volume Revert to Snapshot	Restores volumes from snapshots
Failover Host	<p>Switches the I/O destination to an available storage system with the ETERNUS AF/DX Storage Cluster function if the storage system becomes unavailable. Only manual Failover and manual Failback are supported.</p> <div style="background-color: #f0f0f0; padding: 5px;"> <p> Note</p> <p>For details on the Failover Host function, refer to "5.2.4 Failover Host Function" (page 43).</p> </div>
Multiattach	Enables multiple volume connections simultaneously to hosts or servers

● Cases where Extend Volume cannot be executed and the workarounds

The following table shows the cases where Extend Volume cannot be executed and the workarounds.

Table 5.2 Cases where Extend Volume cannot be executed and the workarounds

Case	Workaround
If Extend Volume is executed for the volume that was specified as the copy source with Clone Volume or the volume that was created with Clone Volume, the process fails if the copy operation has not been completed.	Wait until the session disappears after the copy operation is completed and then execute Extend Volume.
If Extend Volume is executed for the volume that was specified as the copy source with Create Snapshot, the process fails.	The copy session must be deleted. Delete all snapshots that were acquired from the target volume by using Delete Snapshot.
If Extend Volume is executed while a copy is set manually with ETERNUS SF Storage Cruiser or ETERNUS Web GUI, the process fails.	<ul style="list-style-type: none"> • For OPC, wait until the copy operation has been completed. • For copy operations other than OPC, delete all copy sessions that use the target volume as the copy source or the copy destination before starting Extend Volume. Create the copy sessions again after a volume expansion.

- Specification of host information

Host information must be specified in Migrate Volume and Manage Volume.

The input format is as follows:

Host-Name@Backend-Name#Pool-Name

For the following environment or settings, specify "test.localhost@Backend1#PoolA" for the host.

- PoolA is a pool specified in "/etc/cinder/cinder_fujitsu_eternus_dx.xml".

Example

```
$ hostname
test.localhost

$ cat /etc/cinder/cinder.conf

(snip)
[Backend1]
volume_driver=cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_fc.FJDXFCDriver
cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml
volume_backend_name=volume_backend_name1
```

5.2 Supplementary Information for the Supported Functions

5.2.1 QoS Settings

The QoS settings that are linked with the volume QoS function of the ETERNUS AF/DX are available. This function is not supported in the ETERNUS DX S2 series.

When the firmware version of the ETERNUS AF/DX is V11L30 or later, IOPS/Throughput of Total/Read/Write for the volume can be specified separately.

Only an upper limit value of the bandwidth (BWS) for each volume can be set. A lower limit value cannot be set.

The following functions are managed by the QoS settings:

- Functions that create volumes such as Create Volume and Copy Image to Volume
- Functions that register volumes to OpenStack environments such as Manage Volume

The following procedure shows how to set the QoS.

Procedure

1 Create a QoS definition.

- The firmware version of the ETERNUS AF/DX is earlier than V11L30

```
$ cinder qos-create <qos_name> maxBWS=xx
```

For *<qos_name>*, specify the name of the definition that is to be created.
For *maxBWS*, specify a value in MB.

- The firmware version of the ETERNUS AF/DX is V11L30 or later

```
$ cinder qos-create <qos_name> read_iops_sec=15000 write_iops_sec=12600  
total_iops_sec=15000 read_bytes_sec=800 write_bytes_sec=700 total_bytes_sec=800
```

For *<qos_name>*, specify the name of the definition that is to be created.

2 When not using the existing volume type, create a new volume type.

```
$ cinder type-create <volume_type_name>
```

For *<volume_type_name>*, specify the name of the volume type that is to be created.

3 Associate the QoS definition with the volume type.

```
$ cinder qos-associate <qos_specs> <volume_type_id>
```

For *<qos_specs>*, specify the ID of the QoS definition that was created.
For *<volume_type_id>*, specify the ID of the volume type that was created.

End of procedure

■ Cautions

- For the procedure to cancel the QoS settings, refer to "OpenStack Command-Line Interface Reference".
- The QoS mode of the ETERNUS AF/DX must be enabled in advance. For details, refer to the ETERNUS Web GUI manuals.
- When the firmware version of the ETERNUS AF/DX is earlier than V11L30, for the volume QoS settings of the ETERNUS AF/DX, upper limits are set using the predefined options. Therefore, set the upper limit of the ETERNUS AF/DX side to a maximum value that does not exceed the specified maxBWS. The following table shows the upper limits that can be set on the ETERNUS AF/DX side and example settings. For details about the volume QoS settings of the ETERNUS AF/DX, refer to the ETERNUS Web GUI manuals.

Settings for the ETERNUS AF/DX
Unlimited
15000 IOPS (800 MB/s)
12600 IOPS (700 MB/s)
10020 IOPS (600 MB/s)
7500 IOPS (500 MB/s)
5040 IOPS (400 MB/s)
3000 IOPS (300 MB/s)
1020 IOPS (200 MB/s)
780 IOPS (100 MB/s)
600 IOPS (70 MB/s)
420 IOPS (40 MB/s)
300 IOPS (25 MB/s)
240 IOPS (20 MB/s)
180 IOPS (15 MB/s)
120 IOPS (10 MB/s)
60 IOPS (5 MB/s)

- When specified maxBWS=750
 "12600 IOPS (700 MB/s)" is set on the ETERNUS AF/DX side.
- When specified maxBWS=900
 "15000 IOPS (800 MB/s)" is set on the ETERNUS AF/DX side.
- While a QoS definition is being created, if an option other than maxBWS/read_iops_sec/write_iops_sec/total_iops_sec/read_bytes_sec/write_bytes_sec/total_bytes_sec is specified, a warning log is output and the QoS information setting is continued.
- For an ETERNUS AF/DX with a firmware version of V11L30 or later, if a QoS definition volume type that is set with read_iops_sec/write_iops_sec/total_iops_sec/read_bytes_sec/write_bytes_sec/total_bytes_sec is specified for Create Volume, Migrate Volume, or Retype Volume (specified with the "--migration-policy on-demand" option), a warning log is output and the process is terminated.
- For an ETERNUS AF/DX with a firmware version of V11L30 or later, if a QoS definition volume type that is set with maxBWS is specified for Create Volume, Migrate Volume, or Retype Volume (specified with the "--migration-policy on-demand" option), a warning log is output and the process is terminated.
- After the firmware of the ETERNUS AF/DX is upgraded from V11L10/V11L2x to a newer version, the volume types related to the QoS definition created before the firmware upgrade can no longer be used. Set a QoS definition and create a new volume type.

- When the firmware of the ETERNUS AF/DX is downgraded to V11L10/V11L2x, do not use a volume type linked to a pre-firmware downgrade QoS definition, because the QoS definition may work differently from ones post-firmware downgrade. For the volume, create and link a volume type not associated with any QoS definition and after the downgrade, create and link a volume type associated with a QoS definition.
- If Create Volume terminates with an error, Cinder may not invoke Delete Volume. If volumes are created but the QoS settings fail, the ETERNUS OpenStack VolumeDriver ends the process to prevent the created volumes from being left in the ETERNUS AF/DX. If volumes fail to be created, the process terminates with an error.

5.2.2 The Initialization Function During a Volume Deletion

When deleting volumes, this function can overwrite used areas with zeros.

This function is not supported in the ETERNUS DX S2 series.

Note

For the required time, use the following examples as a reference for initialization.

- Examples of the required time to initialize a RAID1+0 volume
50GB: 10 minutes
100GB: 22 minutes

The following procedure shows how to initialize areas during a volume deletion.

Procedure

- 1 Specify a value for the volume type of the deletion target volume.

```
$ cinder type-key <volume_type_id> set type:delete_with_volume_format=True
```

For *<volume_type_id>*, specify the ID of the volume type that is to be deleted.

- 2 Delete the volume.

Overwrite the areas that are being used with zeros and delete the volume.

```
$ cinder delete <volume_id>
```

For *<volume_id>*, specify the volume ID that is to be deleted.

End of procedure

Cautions

- For the procedure to cancel the settings, refer to "OpenStack Command-Line Interface Reference".
- To use this function, the firmware version of the ETERNUS AF/DX must be V10L30 or later.
- If a volume creation and a volume deletion are performed simultaneously, it may take time to format.
- Snapshot volumes cannot be overwritten with zeros.

5.2.3 Mapping File

The mapping file sets the mapping information that specifies the port of the ETERNUS AF/DX to be connected. Make sure it is set to limit the connection port.

5.2.3.1 Parameters

Table 5.3 Parameter set in cinder.conf (1)

Parameter	Description
fujitsu_port_host_mapping_file	Specify the absolute path to the host affinity configuration file of the VolumeDriver. If this parameter is omitted, host affinity is set to any available port in the ETERNUS AF/DX.

5.2.3.2 Setup Procedure

The setup procedure is as follows.

Procedure

- 1 Create an entry in the host affinity configuration file of the VolumeDriver with the following format:

```
<port name>, <port name>.....=<host info>=<lungroup name>
```

port name: Specify the ETERNUS AF/DX port information. Multiple names can be specified.

- For iSCSI

host info: Specify the iSCSI name of the host.

- For FC

host info: Specify the HBA WWPN of the host.

lungroup name: Specify the host affinity LUN group name of the ETERNUS AF/DX.

Example

- For iSCSI

```
#port1,port2,..=host iscsi name=lungroup name  
#port can be set to be xyz or wxyz format.  
#xyz: Fixed 3 digits, CM number + CA number + Port number;  
#wxyz: Fixed 4 digits, CE number(hex) + CM number + CA number + Port  
  
001,010=iqn.1994-05.com.redhat:2b96a70bdd8=AFIN_GRP_#01  
001,010=iqn.1994-05.com.redhat:2b96a70bdd9=AFIN_GRP_#02  
001,010=iqn.1994-05.com.redhat:d9f7bd7fe5aa=AFIN_GRP_#03  
001,010=iqn.1994-05.com.redhat:d9f7bd7fe5a8=AFIN_GRP_#04  
101,011=iqn.1994-05.com.redhat:bc48bdaae32=AFIN_GRP_#02
```

- For FC

```
#port1,port2,...=host HBA WWPN=lungroup name
#port can be set to be xyz or wxyz format.
#xyz: Fixed 3 digits, CM number + CA number + Port number;
#wxyz: Fixed 4 digits, CE number(hex) + CM number + CA number + Port

001,010=10000000c986888b=AFIN_GRP_#01
1101,1011=10000000c986888c=AFIN_GRP_#02
```

2 Set host affinity related entries defined in the host affinity configuration file of the VolumeDriver in the ETERNUS AF/DX.

Use ETERNUS CLI. For details on the settings, refer to the ETERNUS CLI manuals.

3 Change the cinder.conf settings.

Enter the path of the host affinity configuration file of the VolumeDriver in the "fujitsu_port_host_mapping_file =" line. Add the line if it does not exist.

Example

```
fujitsu_port_host_mapping_file=/etc/cinder/fujitsu_port_host_mapping.conf
```

End of procedure

■ Cautions

- Cinder must be restarted after the /etc/cinder/cinder.conf file is updated.
- Cinder must be restarted after the host affinity configuration file of the VolumeDriver is updated.
- If the host affinity relationship that is set in the host affinity configuration file of the VolumeDriver is not set in the ETERNUS AF/DX, a failure may occur in the attach process.
 - If a storage system set with host affinity has a LUN group name that is different from the one set in the host affinity configuration file of the VolumeDriver, and a LUN group that is set with the same name as the one set in the host affinity configuration file of the VolumeDriver exists in the storage system, volume to attach is added to the LUN group that is defined in the file, but LUN access from Nova fails to find that volume and the attach process fails.
 - If the LUN group name that is set in the host affinity configuration file of the VolumeDriver does not exist in the storage system, the failure occurs during the ETERNUS OpenStack VolumeDriver attach process.
- If the "fujitsu_port_host_mapping_file" parameter is not specified, the following settings are used:
 - For FC
Host affinity is set for all available FC ports in the ETERNUS AF/DX that are connected to hosts. It is not set in ports such as those that are not set with zoning or those that are not connected.
 - For iSCSI
Host affinity for the number set in "fujitsu_target_port_use_number" is set for the port starting from the least number of host affinities among the available iSCSI ports in the ETERNUS AF/DX that are connected to hosts. If host affinity has been set, the LUN is added to the LUN group.
- If the "fujitsu_port_host_mapping_file" parameter is set but the host affinity configuration file of the VolumeDriver is not specified or does not exist, the host affinity setting will fail when the ETERNUS OpenStack VolumeDriver is started.

- If the specified host is not defined in the host affinity configuration file of the VolumeDriver when a volume is attached, the following settings are used:
 - For FC
Host affinity is set for all available FC ports in the ETERNUS AF/DX that are connected to hosts. It is not set in ports such as those that are not set with zoning or those that are not connected.
 - For iSCSI
Host affinity for the number set in "fujitsu_target_port_use_number" is set for the port starting from the least number of host affinities among the available iSCSI ports in the ETERNUS AF/DX that are connected to hosts. If host affinity has been set, the LUN is added to the LUN group.

5.2.4 Failover Host Function

For the ETERNUS OpenStack VolumeDriver, the Failover Host function is realized by using the ETERNUS AF/DX Storage Cluster function.

The Storage Cluster function provides business continuity by switching I/O access destinations when one storage system fails while maintaining the data equivalency with the other storage system. The ETERNUS OpenStack VolumeDriver supports only manual failover and manual failback.

5.2.4.1 Preliminary Settings

To introduce Storage Cluster, the settings must be configured using ETERNUS SF Web Console of ETERNUS SF Storage Cruiser. For the Storage Cluster settings, refer to "FUJITSU Storage ETERNUS SF Storage Cruiser Operation Guide".

5.2.4.2 Parameters

Table 5.4 Parameter set in cinder.conf (2)

Parameter	Description
replication_device	Specify the definition of the back-end storage to create a replica.

5.2.4.3 Setup Procedure

The setup procedure is as follows.

Procedure

- 1 Create a local driver configuration file using the local storage pool of the TFO group.
- 2 Create a remote driver configuration file using the remote storage pool of the TFO group.
- 3 Change the cinder.conf settings.
In the "replication_device=" line of the local back-end, enter the back-end storage to create a replica. Add the line if it does not exist.

Example

```
replication_device=
backend_id:hostname@DX_Remote#remote,cinder_eternus_config_file:/etc/cinder/
remote.xml, local_eternus_boxid:
00ETERNUSDXLS3ET203ACU####NL4601417434## ,tfog_name: tfogl
```

- backend_id*: The remote back-end name
- cinder_eternus_config_file*: The location where the remote driver configuration file is stored
- local_eternus_boxid*: The boxid of the local device
- tfog_name*: The TFO group name that is set between the ETERNUS AF/DX storage systems that belong to the local and remote back-ends

End of procedure

5.2.4.4 Usage Procedure

The procedure for using the function is as follows.

Procedure

1 Create a Replication Volume for the Failover Host.

1-1 Create a "Replication" volume type.

```
$ cinder type-create replica
```

1-2 Associate the back-end with the created volume type.

```
$ cinder type-key replica set volume_backend_name= dx_L
```

1-3 Associate the replication definition with the created volume type.

```
$ cinder type-key replica set type:replication_enabled='<is> True'
```

1-4 Perform [Step 1-1](#) through [Step 1-3](#) to create a volume type for the Failover Host.

```
$ cinder extra-specs-list
```

ID	Name	extra_specs
xxxx	replica	{u'volume_backend_name': u'dx_L', u'type:replication_enabled': u'<is> True'}

1-5 Create a volume for the Failover Host.

```
$ cinder create --volume-type replica -display-name <name> <volume size>
```

2 Set an access path to the volume for the Failover Host.

```
$ nova volume-attach <server> <volume>
```

When the above procedure is completed, the Failover Host function becomes available.

End of procedure

■ Cautions

- The local back-end and remote back-end must be set to a TPP or a RAID group at the same time.
- The Own Port and Pair Port, which are used for setting a TFO group, must not be used in the mapping file.
- A TFO group can only exist in one driver configuration file. A driver configuration file can only be used for one TFO group.
- Normally, a REC copy between TPVs is faster than a REC copy between Standard Volumes. Therefore, TPPs are recommended as the local back-end and the remote back-end.

5.2.5 Specification of the Snapshot Creation Destination Pool

A RAID group or a TPP can be specified as the snapshot creation destination pool. In an ETERNUS AF/DX with a firmware version earlier than or equal to V10L60 or in the ETERNUS DX S2 series, TPPs cannot be used as the snapshot creation destination pool.

Multiple snapshot creation destination pools can be specified.

If the creation destination pool is a RAID group, more than 128 snapshots cannot be created. Therefore, to create more than 128 snapshots in a RAID group, multiple RAID groups must be specified as snapshot creation destination pools.

In addition, unlike volume creation destination pools, scheduling is not performed for snapshot creation destination pools.

A pool where snapshots can be created is searched in the order written in the driver configuration file and if one is found, snapshots are created in that pool.

Example of a driver configuration file:

```
<?xml version='1.0' encoding='UTF-8'?>
<FUJITSU>
<EternusIP>0.0.0.0</EternusIP>
<EternusPort>5988</EternusPort>
<EternusUser>smisuser</EternusUser>
<EternusPassword>smispassword</EternusPassword>
<EternusPool>pool1</EternusPool>
<EternusSnapPool>raid1</EternusSnapPool>
<EternusSnapPool>raid2</EternusSnapPool>
</FUJITSU>
```

The above "raid1" and "raid2" are RAID groups.

Example operation when one snapshot is created:

raid1	raid2	Operation
SDV × 128	SDV × 10	A snapshot is created in "raid2" where space is available
SDV × 28, Standard Volume × 100	SDV × 10	A snapshot is created in "raid2" where space is available (cannot be created in "raid1" because there are 128 volumes regardless of volume types)
SDV × 128	SDV × 128	The operation fails because both "raid1" and "raid2" have reached the upper limit of 128
SDV × 10	SDV × 10	A snapshot is created in "raid1" because snapshots are created in the order written in the driver configuration file

 **Caution**

If multiple snapshot creation destination pools are specified, a different pool must be specified for the volume creation destination pool (EternusPool and EternusSnapPool can be specified multiple times but the same pool name cannot be specified). If the same pool name is specified and instructions to create multiple volumes and multiple snapshots are issued at the same time, the number of logical volumes in a RAID group will reach 128 and the operation may fail.

5.2.6 Using Self-Signed HTTPS Certificates

The ETERNUS OpenStack VolumeDriver uses SMI-S to configure the ETERNUS AF/DX. HTTP communication (port 5988) or HTTPS communication (port 5989) is used for accessing SMI-S. The following two types of certificates are supported for HTTPS communication.

- (1) SMI-S self-signed certificate (default)
- (2) SSL certificate for Web GUI

Perform one of the following steps to use HTTPS self-signed certificates.

- Use ETERNUS Web GUI to register self-signed SSL certificates.
For details on how to set certificates, refer to ETERNUS Web GUI manuals.
- Use ETERNUS CLI to specify the SSL certificate for Web GUI as the certificate used with SMI-S.

```
set smi-s -ssl-certificate shared-with-webgui
```

For HTTPS communication, changing the parameter setting can enable or disable the SSL authentication check.

5.2.6.1 Parameters

Table 5.5 Parameters set in cinder.conf (3)

Parameter	Description
fujitsu_ssl_verification	Specify whether to enable or disable the SSL authentication check. If this parameter is omitted, False is specified. <ul style="list-style-type: none">• True Enables the SSL authentication check.• False Disables the SSL authentication check.
fujitsu_ca_certs_path	Specify the absolute path to the SSL certificate file. If this parameter is omitted, /etc/pki/tls/certs/ca-bundle.trust.crt is set.

Chapter 6

Using the Periodic Monitoring Tool

ETERNUS OpenStack VolumeDriver has an auxiliary tool for the Attach/Detach processes. This tool is also used for periodic monitoring.

This monitoring tool periodically checks and sets the host affinity settings based on the definition file to reduce the process time for Attach Volume and Detach Volume.

- For iSCSI
Define the host affinity information (host, ports, and lun group mapping information) in the host affinity configuration file of the VolumeDriver. The monitoring tool executes and checks the host affinity settings based on the definition file.
The Attach/Detach processes of the ETERNUS OpenStack VolumeDriver set only the LUN mappings.
- For FC
Define the Host information (host and host HBA port information) in the VolumeDriver host information configuration file. When the monitoring tool is started, it creates the host affinity configuration file of the VolumeDriver after obtaining the host port and storage port information from the host information configuration file and the ETERNUS AF/DX, and after executing the host affinity settings. Once the host affinity configuration file of the VolumeDriver is created, the host affinity settings are periodically checked based on the configuration file.
The Attach/Detach processes of the ETERNUS OpenStack VolumeDriver set only the LUN mappings based on the host affinity configuration file of the VolumeDriver.

6.1 Monitoring Tool Configuration File

Set the necessary parameters for the periodic monitoring tool in `osvd_host_affinity_monitor.conf`. The following parameters can be specified.

This file must be stored in the same directory as `osvd_host_affinity_monitor.py`.

Parameter	Description
<code>protocol</code>	Optional. Specify the protocol. <ul style="list-style-type: none">• For iSCSI <code>iscsi</code>• For FC <code>fc</code>
<code>osvd_conf</code>	Required. The value of the <code>cinder_eternus_config_file</code> parameter in <code>/etc/cinder/cinder.conf</code> .
<code>host_info_conf</code>	The path name of the VolumeDriver host information configuration file. <ul style="list-style-type: none">• For iSCSI Optional.• For FC Required.

Parameter	Description
host_affinity_mapping_file	Required. The path name of the host affinity configuration file of the VolumeDriver.
temp_volume_name	Optional. The name of the Temp Volume (24MB minimum). <ul style="list-style-type: none">• Default value "!@#_OSVD_temp#@!"
pool_name	Required. The Pool name where the Temp Volume is created.
polling_time	Optional. The interval in which periodic checks are performed (one hour<=POLLING_TIME<=24 hours), (unit: seconds). <ul style="list-style-type: none">• Default value 3600

Setting example of the monitoring tool configuration file

```
protocol=fc
osvd_conf = cinder_fujitsu_eternus_dx.xml
host_info_conf = /etc/cinder/host_info.conf
host_affinity_mapping_file = /etc/cinder/fujitsu_port_host_mapping.conf
temp_volume_name = "!@#_OSVD_temp#@!"
pool_name = OSVD_TPP
polling_time = 3600
```

Setting example of the VolumeDriver host information configuration file

```
#host name="host HBA WWPN","host HBA WWPN",.....

nova_computer_1=10000000c986888a, 10000000c986888b
nova_computer_2=10000000c986888c, 10000000c986888d
```

For details on the host affinity configuration file of the VolumeDriver, refer to ["5.2.3 Mapping File" \(page 41\)](#).

6.2 Starting and Stopping the Monitoring Tool

■ Starting the monitoring tool

Execute the following command to start the monitoring tool.

```
python osvd_host_affinity_monitor.py &
```

■ Stopping the monitoring tool

Execute the following commands to stop the monitoring tool.

```
ps -ef | grep osvd_host_affinity_monitor.py | grep -v grep | cut -c 9-15 | xargs kill -9
```


Chapter 7

Restrictions and Cautions

This chapter provides restrictions and cautions for the ETERNUS OpenStack VolumeDriver.

7.1 Restriction

- CHAP authentication for an iSCSI host interface is not supported.
- When executing the "Create Volume from Volume (Clone)" function or the "Create Volume from Snapshot" function, make sure to specify the same volume type as the source volume. Even when a different volume type has been specified, volumes are created in the backend that is related to the source volume type.
- Execution instructions may fail while the firmware is replaced without stopping the system. Do not create volumes, set access paths, or perform other operations in the storage system where a firmware replacement is being performed. However, VMs that are running do not need to be stopped.

7.2 Caution

- Do not modify the ETERNUS AF/DX resources that are created by the OpenStack user interface (CLI or Dashboard) by using other user interfaces (such as ETERNUS Web GUI or ETERNUS CLI). Doing so results in inconsistency between the ETERNUS AF/DX status and the management information for the OpenStack, causing malfunction.
- Using a single storage pool for different backends is not recommended. The use of the target storage pool becomes inconsistent and a volume creation may fail.
- The following events may occur when a volume migration or a retype fails, and the resulting process that deletes the migration destination volume also fails:

(1) Because the volume is not displayed in the OpenStack screen, it appears to have been successfully deleted.

(2) The volume remains in the ETERNUS AF/DX without being deleted.

For example, the events above may occur in cases where the controller firmware is being loaded or a RAID group fails during a migration.

If these events occur, delete the target volume directly from the ETERNUS Web GUI screen.

- If the following error occurs during an OpenStack version upgrade, uninstall the ETERNUS OpenStack VolumeDriver and then try again.
After the OpenStack version upgrade is completed, install the ETERNUS OpenStack VolumeDriver again.

```
file /usr/lib/python2.7/site-packages/cinder/volume/drivers/fujitsu/
__init__.pyc from install of python-cinder-XXX.el7ost.noarch conflicts with
file from package newton_osvd-XXX.el7.centos.x86_64
file /usr/lib/python2.7/site-packages/cinder/volume/drivers/fujitsu/
__init__.pyo from install of python-cinder-XXX.el7ost.noarch conflicts with
file from package newton_osvd-XXX.el7.centos.x86_64
```

For details about the installation and uninstallation of the ETERNUS OpenStack VolumeDriver, refer to ["Chapter 3 Installation" \(page 17\)](#) and ["Chapter 4 Uninstallation" \(page 30\)](#).

- When using the HTTPS function in OpenStack Rocky, use a version of pyOpenSSL other than 17.3.0.

Chapter 8

Message

This chapter provides ETERNUS OpenStack VolumeDriver messages, their meanings, and troubleshooting.

Messages are output in the Cinder log. Multiple strings of information divided by a "," (comma) are included in the messages. The string of information after the last comma (called "primal keyword") indicates the context of the message.

The following section describes the meanings and actions for each message according to the primal keyword.

■ Error: The CA port setting is invalid

Message example

```
yyyy-mm-dd hh:mm:ss WARNING [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_map_lun,lun_name:[u'600000E00D2A000002A011500020000'],Initiator:[u'10000000c978c574',
u'10000000c978c575'],target:TARGET-ID,Return code:32801,Error: The CA port setting is invalid
```

Meaning

The target port setting for the ETERNUS AF/DX is incorrect.
The WWN or the IQN for the host interface port of the ETERNUS AF/DX is displayed in "TARGET-ID".

Action

When the CA port is used, change the following host interface port parameters by using the relevant commands from the ETERNUS CLI.

- Change the port mode to "CA".
Use the "set host-port-mode" command to change the port mode.
- Enable the host affinity setting.
Use the ETERNUS CLI commands to change the host affinity setting.

Example: For FC ports

```
$ set fc-parameters -port all -host-affinity enable
```

Example: For iSCSI ports

```
$ set iscsi-parameters -port all -host-affinity enable
```

- `import pywbem failed!!pywbem is necessary for this volume driver.`

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common] import
pywbem failed!!pywbem is necessary for this volume driver.
```

Meaning

pywbem is not installed.

Action

Install pywbem in the controller node in which the ETERNUS OpenStack VolumeDriver is installed and then restart Cinder.

- `Error:Maximum number of multi-hop has been reached`

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
create_snapshot,snapshotname:snapshot-ad3bd7a7-a773-458d-8afd-a5bd957b7022,
source volume name:volume-d9a233bb-5305-4674-9195-177ec844d3eb,vol_instance.
path:root/eternus:FUJITSU_StorageVolume.CreationClassName="FUJITSU_Storage
Volume",SystemName="systemname",DeviceID="deviceid",SystemCreationClassName
="FUJITSU_StorageComputerSystem",dest_volumename:FJosv_volumename,pool:root/eternus:
FUJITSU_RAIDStoragePool.InstanceID="FUJITSU:poolname",Return code:35318,
Error:Maximum number of multi-hop has been reached
```

Meaning

- The number of multi-stage copy sessions exceeds the maximum number.
- A copy process is being performed for the volume in which snapshots are created.
- This message may be displayed when creating a snapshot of the volume that was created from the snapshot.

Action

Wait for a while after creating the volume with "Create Volume from Snapshot" and then obtain the snapshot of the created volume.

The status of the created snapshot is "Error". Delete this volume.

■ target_port does not found.

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_get_target_port,protocol:protocol,target_port does not found
```

Meaning

The Target port of the ETERNUS AF/DX could not be obtained.

Action

Confirm that the same protocol is used for the connection protocol of the ETERNUS OpenStack VolumeDriver and the host interface port protocol of the ETERNUS AF/DX.
This message appears when performing "Attach Volume" with the ETERNUS OpenStack VolumeDriver for an iSCSI connection that is connected to the ETERNUS AF/DX by an FC port (an iSCSI port does not exist).

■ Error:Maximum number of session has been reached

Message example

```
yyyy-mm-dd hh:mm:ss.110 6639 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-2f0f97dc-6d8c-4aaa-beb2-658cf8a2da48 eebf3f70b0b84307b2f6d1f6e90ee3c8 a6f7055cfed942e
b9ef8283778a04561] create_snapshot,snapshotname:snapshot-6d16802e-0c10-4efa-ac70-08ea2e0e
150a,source volume name:volume-8212c7b6-025f-4037-8120-08e00f434370,vol_instance.path:root/
eternus:FUJITSU_StorageVolume.CreationClassName="FUJITSU_StorageVolume",SystemName=
"fe80::200:e50:da80:0",DeviceID="600000E00D2A0000002A011500020000",SystemCreationClass
Name="FUJITSU_StorageComputerSystem",dest_volumename:FJosv_Kwe9RfwgZY-NU-
rYoNIUoA==,pool:root/eternus:FUJITSU_RAIDStoragePool.InstanceID="FUJITSU:RSP0008",Return
code:35333,Error:Maximum number of session has been reached(per SourceElement)
```

Meaning

The number of copy sessions exceeds the maximum number of concurrent copy sessions.

Action

Terminate the unnecessary copy sessions.

■ Error:Method Not Supported

Message example

```
yyyy-mm-dd hh:mm:ss.621 10679 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-27ccceaa-97f4-4239-b2ed-9c5413641399 de5d75459b54454587bf0dbaead29be5
d66bb19054fb4c3a92c5ce8b8dbf6596 - - -] create_snapshot,snapshotname:snapshot-a6469255- 024e-
4257-a965-c6c8854b5263,source volume name:volume-b5136db2-ac55-4fd5-86e3-
6b6ee0d9e2f2,vol_instance.path:root/eternus:FUJITSU_StorageVolume.CreationClassName=
"FUJITSU_StorageVolume",SystemName="10.21.86.178",DeviceID="TPPC00152",SystemCreation
ClassName="FUJITSU_StorageComputerSystem",dest_volumename:FJosv_UwJYZYI9x,pool: root/
eternus:FUJITSU_ThinProvisioningPool.InstanceID="FUJITSU:TPP0001",Return code:1,Error: Method Not
Supported
```

Meaning

The "Create Snapshot" process failed.

Action

This message appears if the "Create Snapshot" function is executed when a Thin Provisioning Pool (TPP) is specified as the destination storage pool where snapshots are created in the ETERNUS DX S2 series. Check whether the process target volumes are normal volumes.

The status of the created snapshot is "Error". Delete this snapshot.

■ Error:Invalid Parameter

Message example

```
yyyy-mm-dd hh:mm:ss.661 32382 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-9b705585-0d04-4510-800c-e22b2c84b691 7ce46baba7d243dbbf20e09f1120c636
8b1e178060db4aad826feb196259fc2 - - -]
delete_volume,volumename:FJosv_MXj2PTWimVSXbYEkvzShw==,Return code:5,Error:Invalid Parameter
```

Meaning

The process could not be performed because another process is running on the ETERNUS AF/DX.

Action

- After a certain period of time, perform the process again.
- If the volume state is "error_deleting" or "error_extending", a recovery is possible by executing the following command.

```
cinder reset-state <volume>
```

For <volume>, specify the ID or name of the target volume.

■ no valid snap pool in configuration file

Message example

```
yyyy-mm-dd hh:mm:ss.533 22208 ERROR oslo_messaging.rpc.dispatcher
[req-9eabbbe4-5790-439c-9a49-59e830d5a2d4 dd78669e03b343a790b55f63d89f0314
6dbc0466b13041fa9a5f717ee139cfb5 - - -] Exception during message handling: Bad or
unexpected response from the storage volume backend API: create_snapshot:
3ddc3db9-44e9-458c-a307-7a336e1bbabd, no valid snap pool in configuration file.
```

Meaning

The snapshot pools defined in the driver configuration file cannot be used.

Action

Specify a valid snapshot pool.

■ Insufficient free space

Message example

```
yyyy-mm-dd hh:mm:ss.820 17310 WARNING cinder.scheduler.filters.capacity_filter
[req-dad8e42e-4347-43be-908b-758a39d3792a dd78669e03b343a790b55f63d89f0314
6dbc0466b13041fa9a5f717ee139cfb5 - - -] Insufficient free space for volume creation. Total
capacity is 0.00 on host XXX@YYY#ZZZ.
```

Meaning

The logical free space of each storage pool defined in the driver configuration file is smaller than the volume capacity to be created.

Action

Specify a storage pool with a capacity that is greater than or equal to the volume capacity to be created.

■ No free space in all storage pool for created volume

Message example

```
yyyy-mm-dd hh:mm:ss.980 23676 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_utils [-] No free space in all storage pool for
created volume, necessary space : 6000.
```

Meaning

No storage pools defined in the driver configuration file have a logical free space greater than or equal to the volume capacity to be created.

Action

Specify a storage pool with a capacity that is greater than or equal to the volume capacity to be created.

- all the EternusPools in driver configuration file are not exist. Please edit driver configuration file

Message example

```
yyyy-mm-dd hh:mm:ss.382 1588 WARNING
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-c7c5e041-6312-464b-ae1-136faade6516 - - - -] _find_pools, all the EternusPools in driver
configuration file are not exist. Please edit driver configuration file
```

Meaning

All the storage pools that are defined in the driver configuration file do not exist.

Action

Specify storage pools that exist.

- Error:Failed

Message example

```
yyyy-mm-dd hh:mm:ss.661 32382 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-95c5ba96-8ad4-4765-8e36-a9e45c6ec674 fe271d3cdb94a41b2fc555a94ae80f4 38e2dfc67
9bb43e591be2fbaca5f8dce - - -]
delete_volume,volumename:Fjosv_5NEwDFbCZIJTrEoMEkrFqg==,Return code:4,Error:Failed
```

Meaning

- The maximum number of volumes that are specified as the snapshot creation destination in the RAID group has been reached.
- The process could not be performed because an error was detected in an ETERNUS AF/DX process.

Action

- If this error occurs during a snapshot creation, reduce the number of volumes that are specified as the snapshot creation destination in the RAID group or specify a different RAID group.
- After a certain period of time, perform the process again.
- If the volume state is "error_deleting" or "error_extending", a recovery is possible by executing the following command.

```
cinder reset-state <volume>
```

For *<volume>*, specify the ID or name of the target volume.
After the recovery, execute the operation again as required.

■ The Logical Volume is in busy state

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_delete_volume, volumename: FJosv_8efhrYQGshp6smoiZQjGAA==, Return code: 32786,
Error: The Logical Volume is in busy state.
```

Meaning

A deletion instruction was executed for a volume or snapshot that is currently being used.

Action

● When a deletion instruction is executed for volumes

Information that was set directly with ETERNUS Web GUI or ETERNUS SF Storage Cruiser may have caused this error. Check the volume number from the meta data of the target volume and log in to the ETERNUS AF/DX to check the following items:

- Whether there are any copy sessions that use the target volume as the copy source or the copy destination
- Whether the target volume is registered to any affinity groups.

If either of the above is configured, delete or cancel the settings after making sure that they can be deleted.

If neither of the above is configured and the volume is not being used, the ETERNUS AF/DX may have returned a busy response. Restore the volume status in Cinder to Available and execute Delete Volume again.

● When a deletion instruction is executed for snapshots

The ETERNUS AF/DX may have returned a busy response. Restore the snapshot status in Cinder to Available and execute Delete Snapshot again.

■ volume not found

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_get_manage_volume_instance, source name or source id: {'u'source-name': 'u'xxx'}, volume not found.
```

Meaning

The volume that was specified with Manage Volume does not exist in the ETERNUS AF/DX.

Action

Check the specified ID or volume name and perform the operation again.

■ the pool of volume not in driver configuration file

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_find_pool_from_volume, vol_instance: root/eternus:FUJITSU_StorageVolume.
CreationClassName="FUJITSU_StorageVolume",SystemName="10.21.134.26",DeviceID="600000E00D2A0000
02A011502110000",SystemCreationClassName="FUJITSU_StorageComputerSystem",
the pool of volume not in driver configuration file.
```

Meaning

The volume that was specified with Manage Volume is in a pool that differs from the specified Eternus-Pool in the driver configuration file.

Action

After specifying the volume type that is specified as the volume creation destination for the storage pool in which the volume to import exists, perform the operation again.

■ volume type not support

Message example

```
yyyy-mm-dd hh:mm:ss ERROR [cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common]
_get_manage_volume_instance, volume type: 3, volume type not support.
```

Meaning

An invalid volume was specified for Manage Volume.

Action

Only Standard Volumes and TPVs can be specified for Manage Volume. Specify a valid volume and perform the operation again.

■ Error:Failed, Job:Specified user(xxxx) does not have Software role

Message example

```
yyyy-mm-dd hh:mm:ss.901 9715 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common [-]_check_user,Return code:4,
Error:Failed, Job:Specified user(xxxx) does not have Software role
```

Meaning

The specified user account does not have the "Software" role.

Action

Assign the "Software" role to the user account in `/etc/cinder/cinder.conf`. For details on the setting, refer to the ETERNUS Web GUI manuals.

■ cannot connect to ETERNUS.

Message example

```
yyyy-mm-dd hh:mm:ss.428 23510 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common [-]_exec_eternus_cli,cannot connect to
ETERNUS.SSH Result:255,retdata:None
```

Meaning

The ETERNUS AF/DX is inaccessible.

Action

- Confirm that the user account in `/etc/cinder/cinder.conf` is correct.
- Confirm that the public key of the execution server is registered in the ETERNUS AF/DX.
- Confirm that the communication path to the ETERNUS AF/DX is normal.

■ Error:Because OPC is being executed, the processing was discontinued.

Message example

```
yyyy-mm-dd hh:mm:ss.278 893 ERROR
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common
[req-d3ca000b-64b7-40d4-ae97-d77320b5fcb4 8574ce35383146088f43ef0d8a2f13340f81720937904
f4c9a7c7eb9c69ddffa - - -]
extend_volume,volumename:FJosv_XiEzzChsjD30As0fePVBmw==,Return code:4,
Error:Because OPC is being executed, the processing was discontinued.,PoolType:RAID_GROUP
```

Meaning

The volume expansion process cannot be executed because a physical copy (OPC) is being performed in the ETERNUS AF/DX.

Action

- Execute the process again after a while.
- If the volume status is "error_extending", execute the following command to recover the status.

```
cinder reset-state <volume>
```

Specify the volume ID or the volume name in `<volume>`.

■ Invalid value is input

Message example

```
yyyy-mm-dd hh:mm:ss.661 32382 ERROR  
cinder.volume.drivers.fujitsu.eternus_dx.eternus_dx_common  
[req-fa76ab5f-62f5-4c45-8d4d-e65ed62bcfe9 7ce46baba7d243dbbf20e09f1120c6368b1e178060db4a  
ad826febf196259fc2 - - -] _get_qos_category_by_value,Invalid value is input,key:maxBWS,value: value
```

Meaning

An invalid value has been entered for the QoS settings.

Action

Check the value entered.

Chapter 9

Obtaining the Required Information for Investigating Failures

When an error occurs, obtain the following information for failure investigations before contacting a Fujitsu engineer.

■ Method for obtaining the required information

Execute the following shell script to obtain the ".tar.bz2" formatted file that was created.

```
$ ./osvd_collect_log.sh <full | email> [ output path ]
```

Specify "full" or "email" for the argument.

- If "full" is specified, a single ".tar.bz2" formatted file is created.
- If "email" is specified, ".tar.bz2" formatted files that are divided into 5MB increments are created.

For "*output path*", specify a destination directory for the output. If this parameter is omitted, a log is output to the current directory. In addition, the destination directory must be created in advance.

If the specified directory does not exist, the following error message is displayed.

"Invalid output path was specified : *output path*"

Chapter 10

ETERNUS OpenStack VolumeDriver License

Apache License 2.0 is applied to the ETERNUS OpenStack VolumeDriver.

```
# Copyright (c) 2018 FUJITSU LIMITED
# Copyright (c) 2012 EMC Corporation.
# Copyright (c) 2012 OpenStack Foundation
# All Rights Reserved.
#
# Licensed under the Apache License, Version 2.0 (the "License"); you may
# not use this file except in compliance with the License. You may obtain
# a copy of the License at
#
#     http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS, WITHOUT
# WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the
# License for the specific language governing permissions and limitations
# under the License.
#
```

Appendix A

Driver Configuration File

This appendix provides information on the driver configuration file. In the driver configuration file, set the information about the ETERNUS AF/DX that is to be used. Set the following information.

- SMI-S connection IP address
Enter the IP address for the SMI-S connection of the ETERNUS AF/DX.
Enter the IP address between `<EternusIP>` and `</EternusIP>`.

Example for IPv4

```
<EternusIP>0.0.0.0</EternusIP>
```

Example for IPv6

```
<EternusIP>2001::2</EternusIP>
```

- SMI-S connection port number
Enter the number for the SMI-S connection port of the ETERNUS AF/DX.
Enter the port number between `<EternusPort>` and `</EternusPort>`.

Example:

```
<EternusPort>5988</EternusPort>
```

- SMI-S connection user
Enter the username for the SMI-S connection of the ETERNUS AF/DX.
Enter the connection user between `<EternusUser>` and `</EternusUser>`.

Example:

```
<EternusUser>smisuser</EternusUser>
```

Note

Configure the user that was created as described in "[Account used for the ETERNUS AF/DX](#)" (page 15) in "[2.4 Preliminary Settings](#)".

- SMI-S connection password
Enter the password for the SMI-S connection of the ETERNUS AF/DX.
Enter the password between <EternusPassword> and </EternusPassword>.

Example:

```
<EternusPassword>smispassword</EternusPassword>
```

Note

Configure the password that was created as described in ["Account used for the ETERNUS AF/DX" \(page 15\)](#) in ["2.4 Preliminary Settings"](#).

- The storage pool name that is to be used
Enter the RAID group name or the Thin Provisioning Pool (TPP) name that already exists in the ETERNUS AF/DX. Multiple names can be set.
Enter the RAID group name or the Thin Provisioning Pool (TPP) name between <EternusPool> and </EternusPool>.

Example:

```
<EternusPool>raid5_0001</EternusPool>
<EternusPool>raid5_0002</EternusPool>
```

- The storage pool name to use when creating a snapshot
Enter the RAID group name that already exists in the ETERNUS AF/DX.
Enter the RAID group name between <EternusSnapPool> and </EternusSnapPool>.

Example:

```
<EternusSnapPool>raid5_0001</EternusSnapPool>
```

Caution

If the contents (e.g. SMI-S connection IP address and the SMI-S connection user) of the driver configuration file contains errors and an operation is executed, the status continuously processes and is in a state where it does not complete. In this case, the error messages are not output to the Cinder log file.
If this occurs, confirm the contents of the driver configuration file and specify the correct parameters.

Example of the driver configuration file

```
<?xml version='1.0' encoding='UTF-8'?>
<FUJITSU>
<EternusIP>0.0.0.0</EternusIP>
<EternusPort>5988</EternusPort>
<EternusUser>smisuser</EternusUser>
<EternusPassword>smispassword</EternusPassword>
<EternusPool>raid5_0001</EternusPool>
<EternusSnapPool>raid5_0001</EternusSnapPool>
</FUJITSU>
```

ETERNUS OpenStack VolumeDriver 1.6
User's Guide

P2X0-1453-03ENZO

Date of issuance: May 2022
Issuance responsibility: FUJITSU LIMITED

- The content of this manual is subject to change without notice.
- This manual was prepared with the utmost attention to detail. However, Fujitsu shall assume no responsibility for any operational problems as the result of errors, omissions, or the use of information in this manual.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- The content of this manual may not be reproduced or distributed in part or in its entirety without prior permission from Fujitsu.


FUJITSU