

NEC Storage OpenStack Cinder Driver (Newton)

Installation Guide

This page is intentionally left blank.

Preface

This installation guide describes the installation and setup procedures, basic operations, log messages, and other details of the NEC Storage OpenStack Cinder driver (Newton).

The NEC Storage series disk array subsystem is referred to as the disk array in this manual unless clearly specified.

Terms

Term	Description
OpenStack	IaaS platform that uses open source software. The OpenStack Foundation operates the OpenStack community. IBM, Hewlett-Packard (HP), Dell, Red Hat, Cisco Systems, and many other vendors take part in its development. (NEC is a Gold Member.)
Newton	Release name of OpenStack 2016.2
Nova	Component that manages virtual machines and networks.
Cinder	Block Storage service for OpenStack. It manages block devices that can be connected to virtual machine instances. The interface to operate storage devices is defined and the dedicated driver is used for each used disk array to provide provisioning and other functions.
cinder-volume	Process running on the backend. The cinder driver is run from cinder-volume.
LD Set	Concept used in Access Control. This is configured as aggregation of logical disks.
Access Control	Name of the function that controls access to logical disks from application servers.
iSMCLI	Generic name for command components in the disk array. The cinder-volume-compatible driver uses the iSMCLI commands to reference or specify information.
LD	Stands for a logical disk. This is a unit that upper-level devices (host) use to access the disk array. Multiple physical disks are bound as a virtual media (RANK) for which RAID is built and logically divided into an easy-to-use size.
Volume	Disk (logical disk) managed by OpenStack.
BV (Base Volume)	This is the replication source volume for the snapshot function.
SV (Snapshot Volume)	This is the volume in which the snapshot function saves differences with the base volume.
SRA (Snapshot Reserve Area)	This is the area in which differential data is recorded in order to save a snapshot image.
MV (Master Volume)	This is a copy source volume for the data replication function.
RV (Replication Volume)	The data replication function replicates (copies) data from a master volume to a replication volume.
RSV (Replication reServed Volume)	This is a logical disk for retaining management information of the data replication function.
Snapshot method	This is a method whereby the NEC Storage Manager snapshot function creates a snapshot from OpenStack (this Cinder driver).
Data replication method	This is a method whereby the NEC Storage Manager data replication function creates a snapshot from OpenStack (this Cinder driver).
RV maintaining differences	RV whose differences are to be backed up when creating a Cinder snapshot by using

Term	Description
	the data replication method. _d is added at the end of the logical disk name of the relevant RV.
Full-access mode	For iSCSI, 16 or more servers (initiators) can access one LD Set in this mode.
ALUA	Stands for Asymmetric Logical Unit Access. This function identifies the differences in the performance of the access paths to the logical disks and selects the optimal path.
Optimal path	This is the path that provides the optimal performance for a logical disk.

Precaution

The `sudo` command must be issued by the user who can execute arbitrary commands with `sudo`.

Trademark Information

- * The OpenStack Word Mark and the OpenStack logos are trademarks or registered trademarks of the OpenStack Foundation in the United States and other countries.
- * Other product names and company names, etc. are trademarks or registered trademarks of the associated companies.

First Edition, October 2017

Contents

Chapter 1	Setup Procedure	8
1.1	For NEC Storage M Series (iSCSI Model).....	8
1.2	For NEC Storage M Series (FC Model)	10
Chapter 2	Restricting Access to Portals	11
2.1	Checking the Portals Assigned to the LD Set.....	11
2.2	Adding a Portal to the LD Set	11
2.3	Deleting a Portal to the LD Set.....	12
2.4	Examples of Multi-path Connection Configuration (iSCSI)	12
Chapter 3	About Relation of Server and its LD Set	14
3.1	Basic configuration.....	14
Chapter 4	Installing the Cinder Driver	16
4.1	Required Packages.....	16
4.2	Installing the Driver.....	16
4.2.1	Installing RPM Package	16
Chapter 5	Updating the Cinder Driver	17
Chapter 6	Setting Up the Cinder Driver	18
6.1	Parameters That Must Be Specified in cinder.conf	18
6.1.1	Parameters for the NEC Storage M Series.....	18
6.1.2	Other Parameters	20
6.2	Setting Example for /etc/cinder/cinder.conf	22
6.2.1	When using one config-group	22
6.2.2	When using multi config-group (Multi- Backend)	23
Chapter 7	Basic Operations Using the Cinder Commands.....	25
7.1	When Using One config-group (Backend)	25
7.2	When Using Multiple config-groups (Multi-Backend)	30
7.2.1	Settings Before Creating a Volume	30
7.2.2	Volume Creation and Other Operations	31
7.3	Other Functions	33
7.3.1	Using Images.....	33
7.3.2	Creating a Clone Volume	36
7.3.3	Expanding a Volume	38
7.3.4	Moving Volumes (Migration).....	39
7.3.5	Setting the QoS.....	42

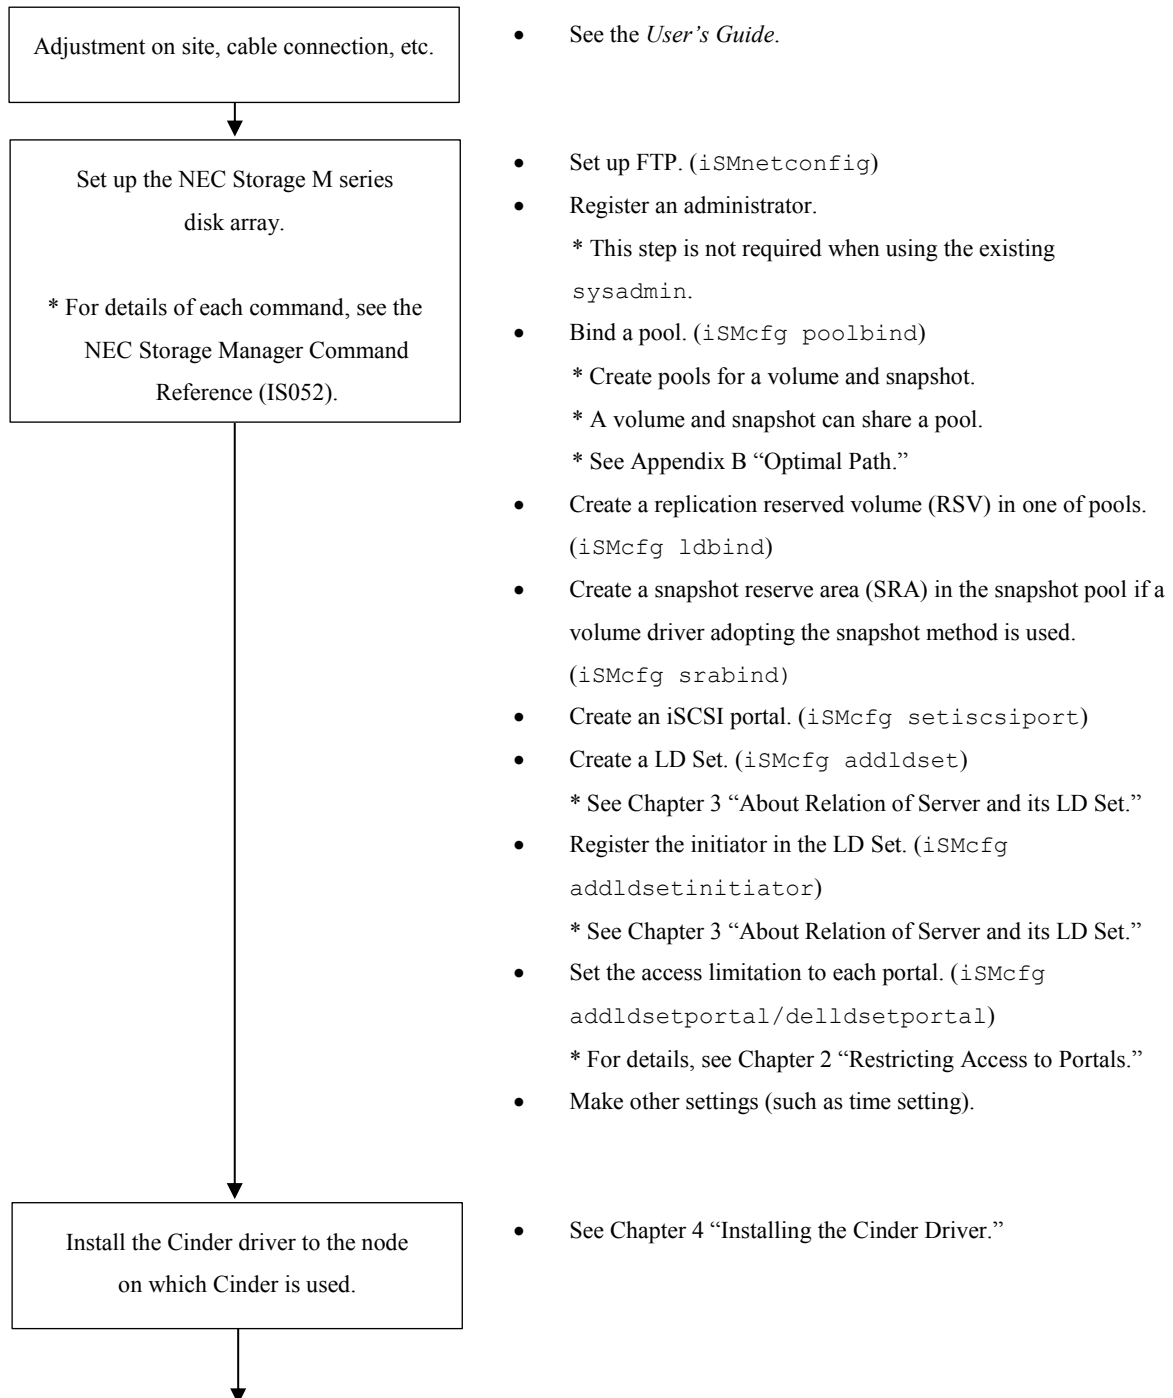
Chapter 8 Correspondence between Cinder Volume IDs and Logical Disk Names in the Disk Array	46
8.1 Command Specifications.....	46
8.2 Setting of Environmental Variables	46
8.3 Command Execution Example.....	47
8.4 Displaying the Usage	49
8.5 Displaying the Version.....	49
Chapter 9 Multi-path Connection	50
9.1 Setting Up DM-Multipath for RHEL	50
9.2 Setting Up multipath-tools for Ubuntu.....	51
9.3 Enabling the Nova Multi-path Function.....	52
9.4 Checking the Multi-path Configuration	53
Chapter 10 Snapshot Function Supporting Differential Backup.....	55
10.1 Creating a Snapshot.....	56
10.1.1 When no RV maintaining differences is paired with the target copy source logical disk	56
10.1.2 When an RV maintaining differences is paired with the copy source logical disk (MV).....	57
10.2 Deleting a Snapshot.....	58
10.2.1 When the target logical disk (RV) is retained as an RV maintaining differences	58
10.2.2 When deleting the target logical disk (RV).....	59
10.2.3 When the target logical disk is not paired with another logical disk	60
10.3 Creating a Volume from a Snapshot	61
10.3.1 When an MV is paired with the copy source logical disk (RV).....	61
10.3.2 When no MV is paired with the copy source logical disk.....	63
10.4 Deleting a Volume	64
10.4.1 When no RV is paired with the target logical disk.....	64
10.4.2 When an RV maintaining differences is paired with the target logical disk (The snapshot paired with the RV is already deleted.)	65
10.4.3 When a snapshot has been created from the target volume.....	66
Chapter 11 get_volume_stats Method.....	67
11.1 Overview	67
11.2 Timing When the get_volume_stats Method Is Called	67
11.3 Return Values.....	67
11.4 Selecting the Pool with the Largest Free Capacity.....	68
11.5 Return Values of the Pool Capacity and Free Pool Capacity.....	68
11.6 Considerations.....	68
Appendix A Notes	69
Appendix B Optimal Path.....	70


Appendix C Setting Full-access Mode.....	71
C.1 Restriction when using Full-access Mode	71
C.2 Creating an LD Set	71
C.3 Registering an Initiator	71
C.4 Deleting an Initiator.....	72
C.5 Checking the Registered Initiator	72
C.6 Checking the Full Access Mode Setting.....	73
Appendix D Log Messages.....	75
D.1 ERROR Level Messages	75
D.2 WARNING Level Messages	77
D.3 INFO Level Messages	80
Appendix E Collecting Failure Information	82
E.1 Log File and Configuration File of the Cinder Driver.....	82
E.2 NEC Storage Manager Logs.....	82
Appendix F Troubleshooting.....	83

Chapter 1 Setup Procedure

1.1 For NEC Storage M Series (iSCSI Model)

The procedure to install and set up the Cinder driver on the disk array is as follows.

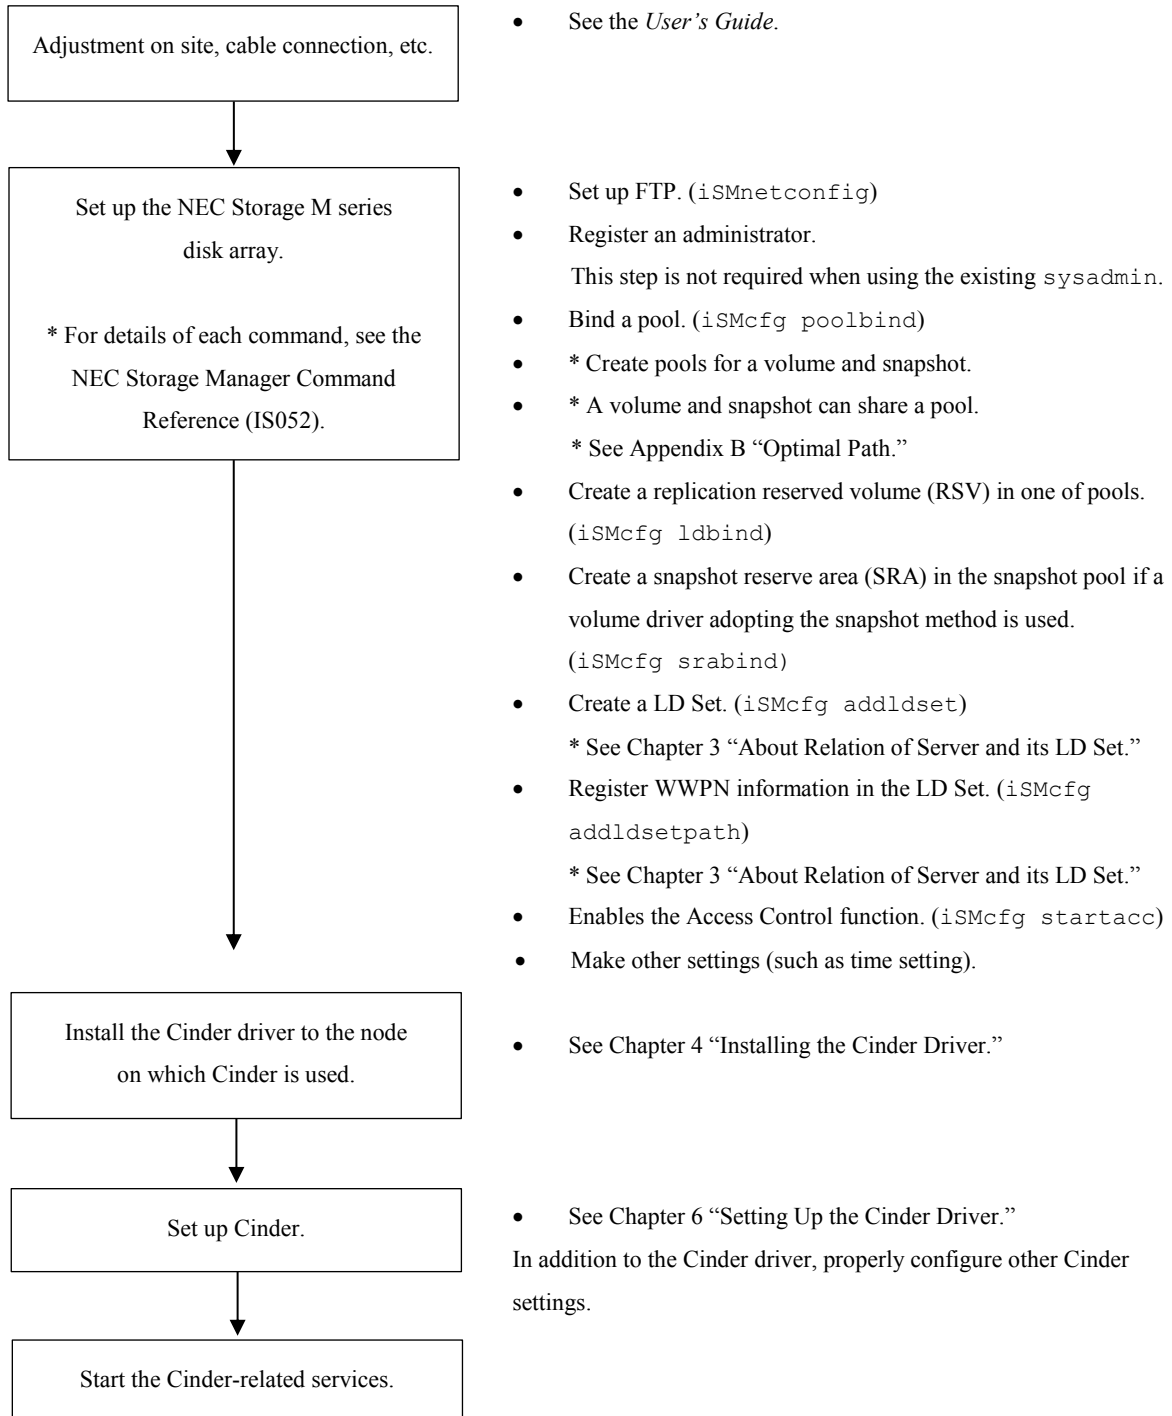




- See Chapter 6 “Setting Up the Cinder Driver.” In addition to the Cinder driver, properly configure other Cinder settings.

1.2 For NEC Storage M Series (FC Model)

The procedure to install and set up the Cinder driver on the disk array is as follows.



Chapter 2 Restricting Access to Portals

If it is necessary to restrict access to each portal, add or delete portals by using the following commands (only for iSCSI interfaces).

2.1 Checking the Portals Assigned to the LD Set

Check the portals assigned to the LD Set by using the `iSMview` command.

* Access from all portals is permitted by default.

Example:

```
iSMview -ac 200000255C3A21CC LX:OpenStack0

--- LD Set(iSCSI) Information ---
Platform           : LX
LD Set Name        : OpenStack0
Configuration Change :
Action mode        : Normal mode
Full Access Mode   : on
Explicit ALUA      : ---
Target Mode        : Multi-Target
Target Count       : 16
Portal             : 172.168.1.230:3260
                   : 172.168.1.231:3260
                   : 172.168.2.230:3260
                   : 172.168.2.231:3260
```

2.2 Adding a Portal to the LD Set

Add a portal to the LD Set by using the `iSMcfg addldsetportal` command.

Example:

```
iSMcfg addldsetportal -ldset LX:OpenStack1 -ipaddress 172.168.1.230
```

2.3 Deleting a Portal to the LD Set

Delete a portal from the LD Set by using the `iSMcfg delldsetportal` command.

Example:

```
iSMcfg delldsetportal -ldset LX:OpenStack0 -ipaddress 172.168.1.230
```

2.4 Examples of Multi-path Connection Configuration (iSCSI)

When using M series NEC Storage with multi-path connections, it is required to take account of the limitation values in iSCSI specification of M series NEC Storage to operate it. Examples of such case of configurations are as follows.

Example: Configuring M110 with two paths

LD Set	Number of LDs (Number of Targets)	Number of Portals (Number of Paths)	Number of Sessions
LD Set A	256	2 (*1)	512 (256 * 2 paths)
LD Set B	256	2 (*1)	512 (256 * 2 paths)
LD Set C	256	2 (*1)	512 (256 * 2 paths)
LD Set D	256	2 (*1)	512 (256 * 2 paths)
Total	1024	8	2048

Example: Configuring M310 with two paths

LD Set	Number of LDs (Number of Targets)	Number of Portals (Number of Paths)	Number of Sessions
LD Set A	1024	2 (*1)	2048 (1024 * 2 paths)
LD Set B	1024	2 (*1)	2048 (1024 * 2 paths)
LD Set C	1024	2 (*1)	2048 (1024 * 2 paths)
LD Set D	1024	2 (*1)	2048 (1024 * 2 paths)
Total	4096	8	8192

Example: Configuring M510 with four paths

LD Set	Number of LDs (Number of Targets)	Number of Portals (Number of Paths)	Number of Sessions
LD Set A	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set B	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set C	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set D	1024	4 (*1)	4096 (1024 * 4 paths)
Total	4096	16	16384

Example: Configuring M710 with four paths

LD Set	Number of LDs (Number of Targets)	Number of Portals (Number of Paths)	Number of Sessions
LD Set A	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set B	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set C	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set D	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set E	1024	4 (*1)	4096 (1024 * 4 paths)
LD Set F	1024	4 (*1)	4096 (1024 * 4 paths)
Total	6144	24	24576

(*1) Assign a unique portal in an LD Set and in all LD Sets.

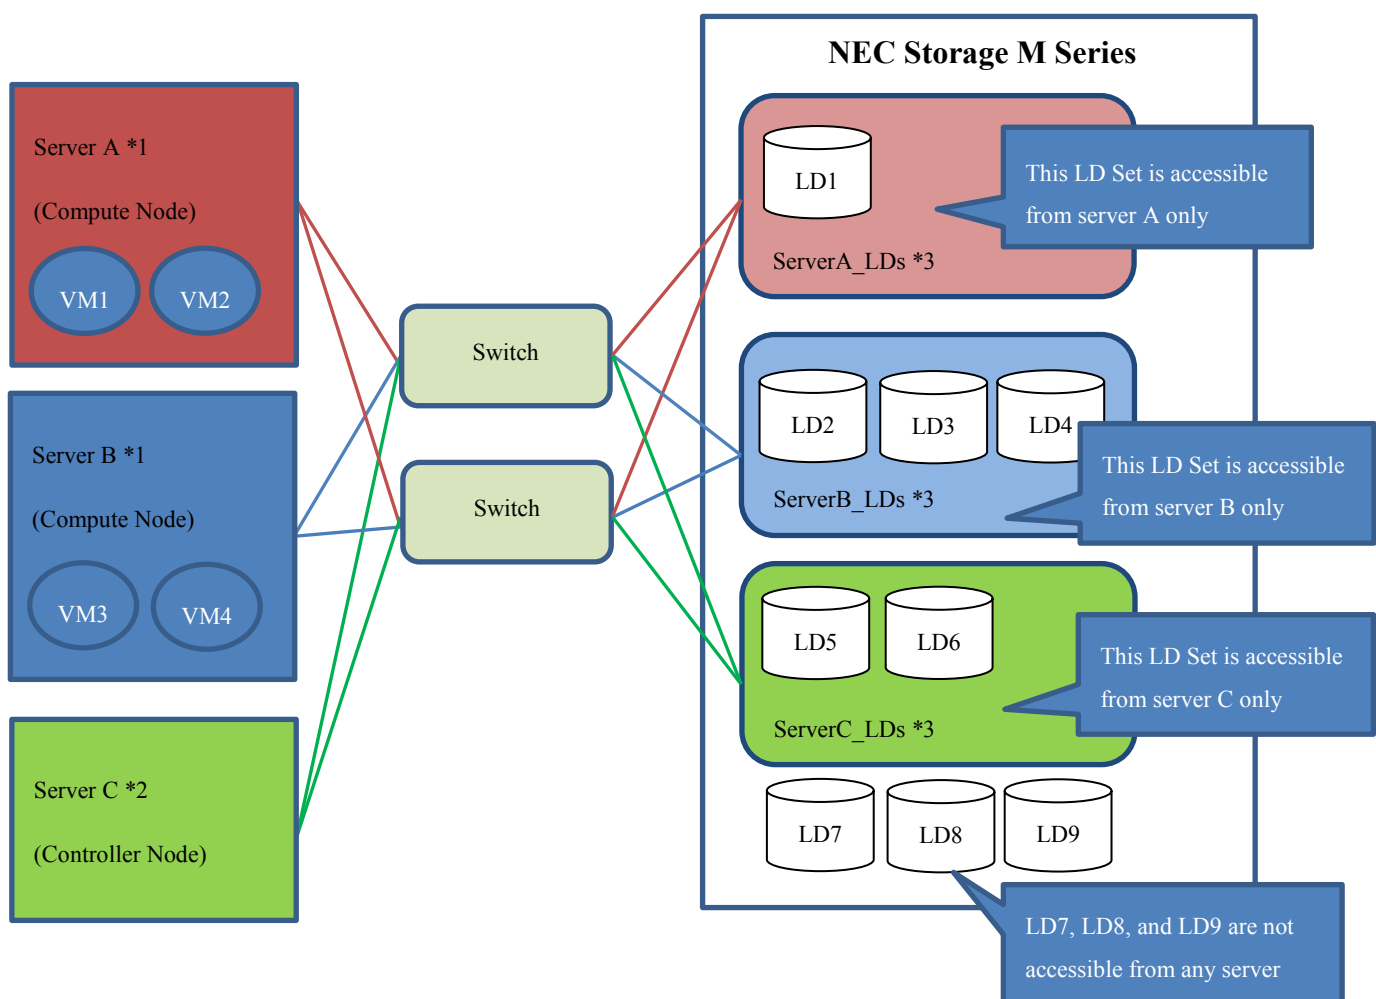
* The matrix below shows the specifications for NEC Storage M110, M310, M310F, M510, M710 and M710F

	M110	M310/M310F	M510	M710/M710F
Number of Sessions per Disk Array	2048	8192	16384	24576
Number of Sessions per Port	256	1024	1024	1024
Number of Portals per LD Set	8	8	16	24

Chapter 3 About Relation of Server and its LD Set

3.1 Basic configuration

In an OpenStack environment, create an LD Set on a server (compute or control node) basis so that each of the servers (compute and control nodes) uses a single LD Set. The following shows an example to create LD Sets.



Relation of server and its LD set

*1 It is when the Nova-Compute which runs on the server A or B attaches a volume to an VM instance that the target logical disk is assigned automatically to the LD Set. The target logical disk is deassigned from the LD set automatically when the volume is detached from the VM instance.

- *2 It is when the Cinder which runs on the server C does access to the logical disk (for image copying and so on) that the target logical disk is assigned automatically to the LD Set. When such process has finished, the target logical disk is released from the LD set.
- *3 The initiator (if iSCSI) or HBA WWPN (if FC) of the server which accesses the target LD Set is to be registered in the LD Set; In the case above, the initiator or HBA WWPN of server A is to be registered in ServerA_LDs, that of server B in ServerB_LDs, that of server C in ServerC_LDs, respectively.
- *4 You can have a configuration where an LD Set is shared by multiple servers (compute nodes). In this case, note that you can register up to 16 iSCSI initiators and 64 HBA WWPNs with an LD Set. For an iSCSI interface, set full-access mode to allow more than 16 servers (initiators) to access the same LD Set. (See “Appendix C Setting Full-access Mode”)
The configuration to set multiple servers (compute nodes) to a single LD Set is for backward compatibility, and not recommended for a 1.5.1 or later version driver.

Chapter 4 Installing the Cinder Driver

4.1 Required Packages

- Install the `sysfsutils` package (for FC connection).

When connecting a compute node to the M series disk array by Fibre Channel, prepare and install the `sysfsutils` package in the compute node.

4.2 Installing the Driver

4.2.1 Installing RPM Package

Install the `rpm` package in the installation medium on OpenStack controller nodes which run Cinder.

- * If the server does not have CD-ROM drive, transfer `rpm` package from another server and install it.

1. Mount the installation medium.

Create a mount directory. (Example: `/cdrom`)

Run the mount command to mount the installation medium.

```
$ /bin/mount -r /dev/cdrom /cdrom
```

2. Run the `rpm` command to install the package.

```
$ /bin/rpm -ivh /cdrom/LINUX-1.7.1/nec_cinder_volume-1.7.1-1.noarch.rpm
```

- * “1.7.1-1” in the above command description means the version of the Cinder driver. The above file name is a sample.

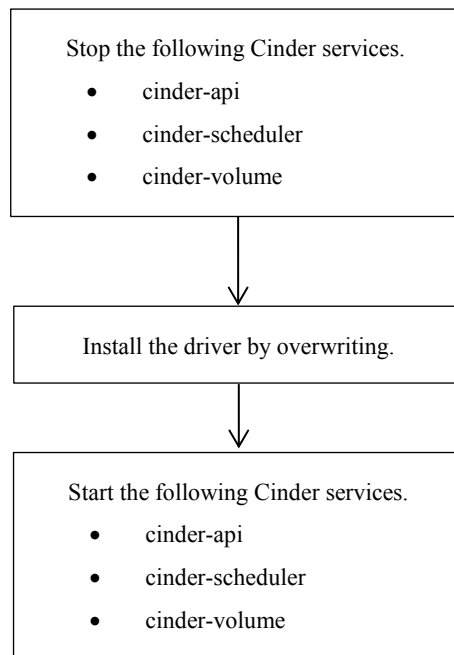
3. Unmount the installation medium.

```
/bin/umount /cdrom
```


Chapter 5 Updating the Cinder Driver

To update the Cinder driver, stop the Cinder services and install the new Cinder driver by overwriting the old one. After installation, start the Cinder services.

The procedure to update the Cinder driver is as follows.



- See Chapter 4 “Installing the Cinder Driver.”

Chapter 6 Setting Up the Cinder Driver

6.1 Parameters That Must Be Specified in `cinder.conf`

6.1.1 Parameters for the NEC Storage M Series

nec_ismcli_fip:

FIP for the disk array (Required)

Example: `nec_ismcli_fip = 192.168.1.13`

nec_ismcli_user:

iSMCLI login user (Required)

Example: `nec_ismcli_user = sysadmin`

nec_ismcli_password: ^{Note1}

iSMCLI login password

Example: `nec_ismcli_password = sys123`

nec_ismcli_privkey: ^{Note1}

RSA secret key file name for iSMCLI

Example: `nec_ismcli_privkey = /etc/cinder/id_rsa`

Note 1: Specify `nec_ismcli_password` or `nec_ismcli_privkey`.

When specifying the RSA secret key file, the specified file must exist in each cinder node and the `cinder-volume` process must be able to access the secret key specified in the file. In addition, the public key must have been registered in the target disk array.

nec_pools:

Specify the pool numbers to be used by separating them with commas. (Can be omitted)

If this parameter is omitted, all the pools on the disk array are used.

An LD is bound as the OpenStack volume from the one with the smallest number of LDs among used pools.

Example: `nec_pools = 0,1`

nec_ldset:

Specify the LD Set for the Compute Node. If this parameter is omitted, the driver automatically selects an appropriate LD Set. (NEC recommends omitting this parameter because the parameter is for backward compatibility.)

Example: `nec_ldset = LX:Server1`

nec_ldset_for_controller_node:

Specify the LD Set for the Controller Node. If this parameter is omitted, the driver automatically selects an appropriate LD Set. (NEC recommends omitting this parameter because the parameter is for backward compatibility.)

Example: `nec_ldset_for_controller_node = LX:OpenStack`

nec_ldname_format:

Format string to generate the logical disk name for the disk array from the ID of the volume recognized in cinder. (Can be omitted)

Specify `LX:%s` (fixed value). * **Note that an arbitrary format string cannot be specified.**

If this parameter is omitted, `LX:%s` is used.

Example: `nec_ldname_format = LX:%s`

nec_backup_pools:

Specify the numbers of the pools in which a snapshot (LD) has been created, delimited by commas. (Required)

An OpenStack snapshot (LD) is created in a pool whose number of LDs is the smallest among the used pools.

* When using the snapshot method, only one pool number can be specified.

* When using the data replication method, multiple pool numbers can be specified.

Example: `nec_backup_pools = 1`

nec_backup_ldname_format:

Format string to generate the logical disk name for the disk array from the ID of the snapshot recognized in cinder. (Can be omitted)

Specify `LX:%s` (fixed value). * **Note that an arbitrary format string cannot be specified.**

If this parameter is omitted, `LX:%s` is used.

Example: `nec_backup_ldname_format = LX:%s`

nec_iscsi_portals_per_cont:

Specify the number of paths per controller for the access paths from a compute node to storage through iSCSI (can be omitted). Either 1 or 2 can be specified. The default value is 1.

This value is enabled only when the Nova multi-path function is enabled (see “9.3 Enabling the Nova Multi-path Function”) and represents the number of iSCSI sessions after a volume is attached to an instance.

6.1.2 Other Parameters

volume_driver:

- Snapshot method
 - In the case of iSCSI interfaces

Specify `cinder.volume.drivers.nec.volume.MStorageISCSIDriver`.

Example: `volume_driver =`

```
cinder.volume.drivers.nec.volume.MStorageISCSIDriver
```
 - In the case of FC interfaces

Specify `cinder.volume.drivers.nec.volume.MStorageFCDriver`.

Example: `volume_driver =`

```
cinder.volume.drivers.nec.volume.MStorageFCDriver
```
- Data replication method
 - In the case of an iSCSI interface

Specify `cinder.volume.drivers.nec.volume.MStorageDDRISCSIDriver`.

Example: `volume_driver =`

```
cinder.volume.drivers.nec.volume.MStorageDDRISCSIDriver
```
 - In the case of an FC interface

Specify `cinder.volume.drivers.nec.volume.MStorageDDRFCDriver`.

Example: `volume_driver =`

```
cinder.volume.drivers.nec.volume.MStorageDDRFCDriver
```

enabled_backends:

Specify the config-groups of the backend to be used.

Example: When using four config-groups: `Storage1, Storage2, Storage3, and Storage4`.

```
enabled_backends = Storage1,Storage2,Storage3,Storage4
```

volume_backend_name:

Specify the backend name for `volume_backend_name` for each config-group.

Example: When specifying "Storage1" for `volume_backend_name` of a config-group.

```
[Storage1]
volume_backend_name = Storage1
```

- * If the same `volume_backend_name` is specified for multiple config-groups, group filtering by capacity filtering (free capacity) can be performed among the config-groups with the same `volume_backend_name`.

nec_backend_max_ld_count:

Specify the maximum number of volumes to be created.

Specify an appropriate value considering the maximum number of logical disks per LD Set (in the case of iSCSI, the maximum number of sessions per LD Set).

When the number of volumes reaches the value specified for this parameter, the `get_volume_stats` method returns 0 (GB) as the free capacity.

Example: When specifying `nec_backend_max_ld_count` for config-group [Storage1].

```
[Storage1]
nec_backend_max_ld_count = 256
```

- * If this parameter is omitted, the maximum number of logical disks per LD Set specified by the disk array specifications is used.
If the specified value is larger than the value of the disk array specifications, the maximum number of logical disks specified by the disk array specifications is also used.

nec_diskarray_name:

This parameter must be specified to configure multiple groups (multi-backend) by using the same storage device (storage device that has the same `nec_ismcli_fip`). Specify the disk array name targeted by the relevant config-group for this parameter.

Example: When specifying `nec_diskarray_name` for groups [Storage1] and [Storage2] for which the same `nec_ismcli_fip` is specified.

```
[Storage1]
nec_ismcli_fip=10.34.80.100
nec_diskarray_name = 200000255C3A21CC
[Storage2]
nec_ismcli_fip=10.34.80.100
nec_diskarray_name = 200000255C3A21CC
```

- * The character strings that can be specified for this parameter is the same as those that can be specified as the disk array name.
- * Exclusive access to the same storage device is controlled by using the disk array name specified

for this parameter.

- * If a different `nec_diskarray_name` is specified for each config-group for which the same storage device (storage device that has the same `nec_ismcli_fip`) is specified, exclusive access to the storage device may not properly function and result in an error as a command is being executed in another config-group (Cinder driver).

rpc_response_timeout:

Set the timeout value in seconds. If three or more volumes can be created at the same time, the reference value is 30 seconds multiplied by the number of volumes created at the same time.

Also, Specify nova parameters below in nova.conf file (/etc/nova/nova.conf).

`block_device_allocate_retries` : Specify 120

`block_device_allocate_retries_interval`: Specify the number of volumes created at the same time.

In addition, you need to edit the following value in the HAProxy configuration file (/etc/haproxy/haproxy.cfg) in an environment where HAProxy is used (*).

`timeout server` : Specify a value greater than `rpc_response_timeout`.

Run the following command after editing the value to reload the HAProxy settings.

```
# service haproxy reload
```

- * Note that the OpenStack environment set up using Red Hat OpenStack Platform Director may be set to use HAProxy.

6.2 Setting Example for /etc/cinder/cinder.conf

6.2.1 When using one config-group

- When using `nec_ismcli_password` to authenticate iSMCLI (Password authentication)

```
[DEFAULT]
enabled_backends = Storage1

[Storage1]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Storage1
nec_ismcli_fip = 192.168.1.10
nec_ismcli_user = sysadmin
nec_ismcli_password = sys123
nec_pools = 0
nec_backup_pools = 1
```

- When using `nec_ismcli_privkey` to authenticate iSMCLI
(Public key authentication)

```
[DEFAULT]
enabled_backends = Storage1

[Storage1]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Storage1
nec_ismcli_fip = 192.168.1.10
nec_ismcli_user = sysadmin
nec_ismcli_privkey = /etc/cinder/id_rsa
nec_pools = 0
nec_backup_pools = 1
```

6.2.2 When using multi config-group (Multi-Backend)

Configuration example:

- Four config-groups (backends)
 - Storage1, Storage2, Storage3, Storage4
- Two disk arrays
 - 200000255C3A21CC (192.168.1.10): Example for using config-group, Storage1 and Storage2
 - 2000000991000316 (192.168.1.20): Example for using config-group, Storage3 and Storage4

```
[DEFAULT]
enabled_backends = Storage1,Storage2,Storage3,Storage4

[Storage1]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Gold
nec_ismcli_fip = 192.168.1.10
nec_ismcli_user = sysadmin
nec_ismcli_password = sys123
nec_pools = 0
nec_backup_pools = 2
nec_diskarray_name = 200000255C3A21CC

[Storage2]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Silver
nec_ismcli_fip = 192.168.1.10
nec_ismcli_user = sysadmin
nec_ismcli_password = sys123
nec_pools = 1
```

```
nec_backup_pools = 3
nec_diskarray_name = 200000255C3A21CC

[Storage3]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Gold
nec_ismcli_fip = 192.168.1.20
nec_ismcli_user = sysadmin
nec_ismcli_password = sys123
nec_pools = 0
nec_backup_pools = 2
nec_diskarray_name = 2000000991000316

[Storage4]
volume_driver = cinder.volume.drivers.nec.volume.MStorageISCSIDriver
volume_backend_name = Silver
nec_ismcli_fip = 192.168.1.20
nec_ismcli_user = sysadmin
nec_ismcli_password = sys123
nec_pools = 1
nec_backup_pools = 3
nec_diskarray_name = 2000000991000316
```


Chapter 7 Basic Operations Using the Cinder Commands

It is necessary to set the environmental variables to run the Cinder command.

Example:

```
export OS_USERNAME=admin
export OS_PASSWORD=openstack
export OS_TENANT_NAME=admin
export OS_AUTH_URL=http://localhost:5000/v2.0/
```

7.1 When Using One config-group (Backend)

Operation example:

1. Display a volume list.

```
$ cinder list
+-----+-----+-----+-----+-----+
| ID | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

2. Create a 1-GB volume.

```
$ cinder create 1
+-----+-----+
| Property | Value |
+-----+-----+
| attachments | [] |
| availability_zone | nova |
| bootable | false |
| created_at | 2013-06-26T07:34:32.460094 |
| display_description | None |
| display_name | None |
| id | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 |
| metadata | {} |
| size | 1 |
| snapshot_id | None |
| source_volid | None |
| status | creating |
| volume_type | None |
+-----+-----+
```

Confirmation:

```
$ cinder list
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | available | None | 1 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

↑ The volume creation command (`cinder create`) runs asynchronously. Command execution is complete when Status becomes available.

3. Attach the created volume to the instance.

```
$ cinder list
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | available | None | 1 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

↑ Volume to attach

```
$ nova list
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| ID          | Name      | Status | Task State | Power State | Networks          |
+-----+-----+-----+-----+-----+-----+-----+-----+
| a91ce522-669d-4920-bfd1-fc87e65d4efe | instance1 | ACTIVE | None       | Running     | private=10.0.0.2 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

↑ Instance to which the volume is attached

Execute the volume attachment.

```
$ nova volume-attach instance1 24944c37-db40-45bc-ad11-fb3e9c54c9f0
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Property | Value          |
+-----+-----+-----+-----+-----+-----+-----+-----+
| device   | /dev/vdb       |
| serverId | a91ce522-669d-4920-bfd1-fc87e65d4efe |
| id       | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 |
| volumeId | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

↑ The volume has been attached as `/dev/vdb`.

Confirmation:

```
$ cinder list
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | in-use | None | 1 | None | false | a91ce522-669d-4920-bfd1-fc87e65d4efe |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

4. Detach the volume from the instance.

```
$ nova volume-detach instance1 24944c37-db40-45bc-ad11-fb3e9c54c9f0
```

Confirmation:

```
$ cinder list
```

```

+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | available | None | 1 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+

```

5. Create a snapshot. *1, *2, *3

```
$ cinder snapshot-create 24944c37-db40-45bc-ad11-fb3e9c54c9f0
```

```

+-----+-----+-----+-----+-----+-----+-----+
|      Property      |          Value          |
+-----+-----+-----+-----+-----+-----+-----+
|   created_at   | 2013-10-17T05:26:34.799868 |
| display_description | None |
| display_name | None |
| id | 720847df-b37b-4320-9b97-c2a23d4e9309 |
| metadata | {} |
| size | 1 |
| status | creating |
| volume_id | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 |
+-----+-----+-----+-----+-----+-----+-----+

```

Confirmation:

```
$ cinder snapshot-list
```

```

+-----+-----+-----+-----+-----+-----+-----+
|          ID          |          Volume ID          | Status | Display Name | Size |
+-----+-----+-----+-----+-----+-----+-----+
| 720847df-b37b-4320-9b97-c2a23d4e9309 | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | creating | None | 1 |
+-----+-----+-----+-----+-----+-----+-----+

```

↑ The snapshot creation command (`cinder snapshot-create`) runs asynchronously. Operation on the master volume can be resumed when Status becomes available.

* For details, see Chapter 10, “Snapshot Function Supporting Differential Backup.”

6. Create a volume from the created snapshot. *2, *3

```
$ cinder create --snapshot-id 720847df-b37b-4320-9b97-c2a23d4e9309 1
```

```
+-----+-----+
|      Property      |      Value      |
+-----+-----+
|      attachments   |      []         |
|  availability_zone |      nova       |
|      bootable      |      false      |
|      created_at    | 2013-10-17T05:34:39.029901 |
| display_description |      None       |
|      display_name  |      None       |
|      id            | 82f38460-919f-4722-a871-9e02e342b3a0 |
|      metadata      |      {}         |
|      size          |      1          |
|      snapshot_id   | 720847df-b37b-4320-9b97-c2a23d4e9309 |
|      source_volid  |      None       |
|      status        |      creating   |
|      volume_type   |      None       |
+-----+-----+
```

Confirmation:

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|      ID      | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | available | None | 1 | None | false | d |
| 82f38460-919f-4722-a871-9e02e342b3a0 | creating | None | 1 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ The volume creation command (`cinder create`) runs asynchronously. Operation on the restored volume can be resumed when `Status` becomes `available`.

7. Delete the snapshot. *3, *4

```
$ cinder snapshot-delete 720847df-b37b-4320-9b97-c2a23d4e9309
```

Confirmation:

```
$ cinder snapshot-list
```

```
+-----+-----+-----+-----+-----+
| ID | Volume ID | Status | Display Name | Size |
+-----+-----+-----+-----+-----+
| 720847df-b37b-4320-9b97-c2a23d4e9309 | 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | deleting | None | 1 |
+-----+-----+-----+-----+-----+
```

↑ deleting is displayed for Status while deleting the snapshot.

```
$ cinder snapshot-list
```

```
+-----+-----+-----+-----+-----+
| ID | Volume ID | Status | Display Name | Size |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

8. Delete the volume.

```
$ cinder delete 24944c37-db40-45bc-ad11-fb3e9c54c9f0
```

Confirmation:

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
| ID | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| 24944c37-db40-45bc-ad11-fb3e9c54c9f0 | deleting | None | 1 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ deleting is displayed for Status while deleting the volume.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
| ID | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
```

(*1) Number of snapshot generations

- For the snapshot method, up to 256 snapshot generations can be created.
- For the data replication method, up to three snapshot generations can be created.

(*2) When using the snapshot method, refer to “Snapshot reserve area” in Appendix “Notes”

(*3) When using the data replication method, refer to Chapter 10 “Snapshot Function Supporting Differential Backup.”

(*4) For the snapshot method, it may take some time to delete intermediate snapshot generations.

7.2 When Using Multiple config-groups (Multi-Backend)

7.2.1 Settings Before Creating a Volume

Operation example:

1. Generate volume types.

```
$ cinder type-create Gold
```

```
+-----+
|          ID          |      Name      |
+-----+
| 12b01baf-e89c-48fe-8c3b-12d93d778696 |      Gold      |
+-----+
```

```
$ cinder type-create Silver
```

```
+-----+
|          ID          |      Name      |
+-----+
| aa1c5c10-14d5-4b95-98ce-cfa78f164f0c |      Silver    |
+-----+
```

2. Display a list the volume types generated in step 1.

```
$ cinder type-list
```

```
+-----+
|          ID          |      Name      |
+-----+
| 12b01baf-e89c-48fe-8c3b-12d93d778696 |      Gold      |
| aa1c5c10-14d5-4b95-98ce-cfa78f164f0c |      Silver    |
+-----+
```

3. Associate the volume types generated in step 1 with `volume_backend_name` defined in `cinder.conf`.

```
$ cinder type-key Gold volume_backend_name=Gold
```

```
$ cinder type-key Silver set volume_backend_name=Silver
```

* In the above examples, two `volume_backend_name` parameters are defined.

7.2.2 Volume Creation and Other Operations

1. Create a 5-GB volume on volume type Gold.

```
$ cinder create --volume_type Gold 5
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2013-10-16T04:36:49.966918
display_description	None
display_name	None
id	e683842b-3cd2-4aaa-b679-cb1c8557bb32
metadata	{}
size	5
snapshot_id	None
source_volid	None
status	creating
volume_type	Gold

* Specify a volume type for --volume_type.

2. Create a 10-GB volume on volume type Silver.

```
$ cinder create --volume_type Silver 10
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2013-10-16T04:38:44.578144
display_description	None
display_name	None
id	e9c43e39-b7a2-4fcc-909a-482ae6b1eb2b
metadata	{}
size	10
snapshot_id	None
source_volid	None
status	creating
volume_type	Silver

* Specify a volume type for --volume_type.

- List the volumes created in steps 1 and 2.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| e683842b-3cd2-4aaa-b679-cb1c8557bb32 | available | None | 5 | Gold | false | |
| e9c43e39-b7a2-4fcc-909a-482ae6b1eb2b | available | None | 10 | Silver | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

* The backend where the volume is created can be identified by `Volume Type`.

* The following operations are the same as those in 7.1 “When Using One config-group (Backend).”

- Attach the created volume to the instance.
- Detach the volume from the instance.
- Create a snapshot.
- Create a volume from the created snapshot.
- Delete the snapshot.
- Delete the volume.

7.3 Other Functions

7.3.1 Using Images

Operation example:

1. Create a volume by using images.

- Check the registered images.

```
$ glance image-list
```

```
+-----+-----+-----+-----+
| ID              | Name              | Status | Server |
+-----+-----+-----+-----+
| d56a37c7-5fb4-42c9-ab45-6bf995079168 | cirros-0.3.1-x86_64-uec      | ACTIVE |      |
| 20f93e96-9813-4934-9935-7a50024f7592 | cirros-0.3.1-x86_64-uec-kernel | ACTIVE |      |
| c1d4e113-4946-427f-9980-95c32e344a09 | cirros-0.3.1-x86_64-uec-ramdisk | ACTIVE |      |
+-----+-----+-----+-----+
```

- Create a volume by selecting an image from the registered images.

```
$ cinder create --image-id d56a37c7-5fb4-42c9-ab45-6bf995079168 20
```

```
+-----+-----+
| Property      | Value              |
+-----+-----+
| attachments   | []                 |
| availability_zone | nova               |
| bootable      | false              |
| created_at    | 2014-04-15T07:14:26.193955 |
| display_description | None               |
| display_name  | None               |
| id            | fb2a8916-4a30-4002-b0cf-4f186c344169 |
| image_id      | d56a37c7-5fb4-42c9-ab45-6bf995079168 |
| metadata      | {}                 |
| size          | 20                 |
| snapshot_id   | None               |
| source_volid  | None               |
| status        | creating           |
| volume_type   | None               |
+-----+-----+
```

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
| ID              | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | creating | None | 20 | None | false |      |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ creating is displayed for Status while creating a volume.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | downloading | None | 20 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ downloading is displayed for Status while copying the selected image.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | available | None | 20 | None | true | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ Volume creation is complete when Status becomes available.

2. Upload the volume image.

Operation example:

- Display and check a list of volumes.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | available | None | 20 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

- Select a volume from the displayed volumes, and register it to the image service.

Example: The name of the image to be registered is `test-image`.

```
$ cinder upload-to-image fb2a8916-4a30-4002-b0cf-4f186c344169 test-image
```

```
+-----+-----+-----+-----+-----+-----+-----+
| Property | Value |
+-----+-----+-----+-----+-----+-----+
| container_format | bare |
| disk_format | raw |
| display_description | None |
| id | fb2a8916-4a30-4002-b0cf-4f186c344169 |
| image_id | cd9fceed-1409-48e6-a96a-2497f5ac34e1 |
| image_name | test-image |
| size | 20 |
| status | uploading |
| updated_at | 2014-04-15T07:14:59.000000 |
| volume_type | None |
+-----+-----+-----+-----+-----+-----+-----+
```

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | uploading | None | 20 | None | false | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ uploading is displayed for Status while registering the selected image.

```
$ cinder list
```

```
+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
+-----+-----+-----+-----+-----+-----+-----+
| fb2a8916-4a30-4002-b0cf-4f186c344169 | available | None | 20 | None | true | |
+-----+-----+-----+-----+-----+-----+-----+
```

↑ Uploading the volume image is complete when Status becomes available.

- Confirm that the image has been registered.

```
$ glance image-list
```

```
+-----+-----+-----+-----+-----+-----+
| ID          | Name          | Disk Format | Container Format | Size      | Status |
+-----+-----+-----+-----+-----+-----+
| d56a37c7-5fb4-42c9-ab45-6bf995079168 | cirros-0.3.1-x86_64-uec | ami        | ami              | 25165824 | active |
| df430cc2-3406-4061-b635-a51c16e488ac | cirros-0.3.1-x86_64-uec-kernel | aki        | aki              | 4955792  | active |
| 3cf852bd-2332-48f4-9ae4-7d926d50945e | cirros-0.3.1-x86_64-uec-ramdisk | ari        | ari              | 3714968  | active |
| cd9fceed-1409-48e6-a96a-2497f5ac34e1 | test-image          | raw        | bare             | 95165824 | active |
+-----+-----+-----+-----+-----+-----+
```

7.3.2 Creating a Clone Volume

- Create a clone volume.

Operation example:

1. Display and check a list of volumes.

```
$ cinder list
```

ID	Status	Display Name	Size	Volume Type	Bootable	Attached to
b7b3ff35-41de-482f-bf11-3e520e8b1c0d	available	None	2	None	false	

2. Create a clone volume by selecting a volume from the displayed volumes.

```
$ cinder create --source-uuid b7b3ff35-41de-482f-bf11-3e520e8b1c0d 2
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2014-04-25T04:25:29.613667
display_description	None
display_name	None
encrypted	False
id	7fd35865-4315-4a22-9d8f-7b18fde7a942
metadata	{}
size	2
snapshot_id	None
source_uuid	b7b3ff35-41de-482f-bf11-3e520e8b1c0d
status	creating
volume_type	None

```
$ cinder list
```

ID	Status	Display Name	Size	Volume Type	Bootable	Attached to
7fd35865-4315-4a22-9d8f-7b18fde7a942	creating	None	2	None	false	
b7b3ff35-41de-482f-bf11-3e520e8b1c0d	creating	None	2	None	false	

↑ creating is displayed for Status while creating a clone volume.

```
$ cinder list
```

ID	Status	Display Name	Size	Volume Type	Bootable	Attached to
7fd35865-4315-4a22-9d8f-7b18fde7a942	available	None	2	None	false	
b7b3ff35-41de-482f-bf11-3e520e8b1c0d	available	None	2	None	false	

↑ Volume creation is complete when Status becomes available.

Notes:

1. For the data replication method, when a maximum number (three generations) of replication volumes (RVs) has already been created for a clone source volume, it assumed as an error.
2. For the data replication method, when two snapshots have been created for a clone source volume, an error occurs if a new snapshot of the clone source volume is created before creation of a clone volume is complete. (This is because the number of pairs of the clone source volume temporarily reaches a maximum number (three generations.)

7.3.3 Expanding a Volume

- Expand a volume.

Operation example:

- Display and check a list of volumes.

```
$ cinder list
```

```
-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
-----+-----+-----+-----+-----+-----+-----+-----+
| 86130194-d424-4292-b956-bbaccf13030 | available | None | 10 | None | false | |
-----+-----+-----+-----+-----+-----+-----+-----+
```

- Select a volume from the displayed volumes and expand its size to 20 GB.

```
$ cinder extend 86130194-d424-4292-b956-bbaccf13030 20
```

```
$ cinder list
```

```
-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
-----+-----+-----+-----+-----+-----+-----+-----+
| 86130194-d424-4292-b956-bbaccf13030 | extending | None | 10 | None | false | |
-----+-----+-----+-----+-----+-----+-----+-----+
```

↑ extending is displayed for Status while expanding the volume size.

```
$ cinder list
```

```
-----+-----+-----+-----+-----+-----+-----+-----+
|          ID          | Status | Display Name | Size | Volume Type | Bootable | Attached to |
-----+-----+-----+-----+-----+-----+-----+-----+
| 86130194-d424-4292-b956-bbaccf13030 | extending | None | 20 | None | false | |
-----+-----+-----+-----+-----+-----+-----+-----+
```

↑ Expanding the volume is complete when Status becomes available.

Notes:

- For the data replication method, for the M110, M310, M310F, M510, M710 and M710F disk arrays, volumes for which a pair setting has been specified (a replication volume (RV) has been created) can be expanded. For other disk arrays, an error occurs if trying to expand such volumes.
- The snapshot method cannot be used to expand a volume (BV) that has snapshots.
- The volume status must be available.

7.3.4 Moving Volumes (Migration)

- Move a volume to another backend.

Operation example:

- Check the current backend information (`os-vol-host-attr:host`) of the target volume.

```
$ cinder show 8ac21714-6c93-49a1-821b-22c3a7d545d2
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2014-05-22T08:10:06.000000
display_description	None
display_name	None
encrypted	False
id	8ac21714-6c93-49a1-821b-22c3a7d545d2
metadata	{}
os-vol-host-attr:host	server1@Storage1#Gold
os-vol-mig-status-attr:migstat	None
os-vol-mig-status-attr:name_id	None
os-vol-tenant-attr:tenant_id	d7afb0e4c33c460dbfaa7134e9232eb8
size	10
snapshot_id	None
source_volid	None
status	available
volume_type	None

* The backend is displayed in the format `<node-name>@<config-group-name>#<backend-name>`.

- Check `enabled_backends` of `cinder.conf` and the backend name of each config-group.

```
-----
enabled_backends = Storage1, Storage2, Storage3, Storage4

[Storage1]
volume_backend_name=Gold

[Storage2]
volume_backend_name=Silver

[Storage3]
volume_backend_name=Gold

[Storage4]
volume_backend_name=Silver
-----
```

* The character string in `[]` indicates a config-group.

The name specified for `volume_backend_name` is the backend name.

- Check the node name of Cinder.

```
$ cinder-manage host list
server1@Storage1 nova
server1@Storage2 nova
server1@Storage3 nova
server1@Storage4 nova
```

* The node name is displayed in the format <node-name>@<config-group-name>.

Storage1, Storage2, Storage3, and Storage4 are the config-group names of each backend.

- Move (migrate) the target volume to **server1@Storage3#Gold**.

```
$ cinder migrate 8ac21714-6c93-49a1-821b-22c3a7d545d2
server1@Storage3#Gold
```

* The destination backend is displayed in the format <node-name>@<config-group-name>#<backend-name>.

```
$ cinder show 8ac21714-6c93-49a1-821b-22c3a7d545d2
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2014-05-22T08:10:06.000000
display_description	None
display_name	None
encrypted	False
id	8ac21714-6c93-49a1-821b-22c3a7d545d2
metadata	{}
os-vol-host-attr:host	server1@Storage1#Gold
os-vol-mig-status-attr:migstat	migrating
os-vol-mig-status-attr:name_id	None
os-vol-tenant-attr:tenant_id	d7afb0e4c33c460dbfaa7134e9232eb8
size	10
snapshot_id	None
source_volid	None
status	available
volume_type	None

↑ os-vol-mig-status-attr:migstat is migrating is displayed for os-vol-mig-status-attr:migstat while moving (migrating the volume).


```
$ cinder show 8ac21714-6c93-49a1-821b-22c3a7d545d2
```

Property	Value
attachments	[]
availability_zone	nova
bootable	false
created_at	2014-05-22T08:10:06.000000
display_description	None
display_name	None
encrypted	False
id	8ac21714-6c93-49a1-821b-22c3a7d545d2
metadata	{}
os-vol-host-attr:host	server1@Storage3#Gold
os-vol-mig-status-attr:migstat	None
os-vol-mig-status-attr:name_id	3bb0cf66-cd67-4d99-b184-2203594dd132
os-vol-tenant-attr:tenant_id	d7afb0e4c33c460dbfaa7134e9232eb8
size	10
snapshot_id	None
source_volid	None
status	available
volume_type	None

↑ Moving (migrating) the volume is complete when `os-vol-host-attr:host` becomes the destination host (`server1@Storage3#Gold`) and `os-vol-mig-status-attr:migstat` becomes `None`.

Note:

1. A volume that contains snapshots cannot be moved.
2. For the data replication method, a volume (MV) that is paired with an RV in the maintaining differences status cannot be moved even if the MV does not contain any snapshots.
3. The volume status must be `available`.

7.3.5 Setting the QoS

- Set the QoS.

* When a volume is created, the QoS is set by specifying the volume type that is associated with the QoS specification.

- Keys that can be set in the QoS specification

No.	Key	Description	Remark
1	upperlimit	Set the I/O upper limit.	Can be omitted (*1)
2	lowerlimit	Set the I/O lower limit.	Can be omitted (*1)
3	upperreport	Specify whether to output data in the operation log of NEC Storage Manager when I/O is suppressed because it has reached the upper limit.	Can be omitted

* Keys other than the above are invalid.

- Values that can be specified for the QoS specification keys

No.	Key	Specifiable Values	Remark
1	upperlimit	0, 10 to 1000000	If 0 is specified, the I/O upper limit is not set. (*2)
2	lowerlimit	0, 10 to 1000000	If 0 is specified, the I/O lower limit is not set. (*2)
3	upperreport	on, off	on: Output to the operation log of NEC Storage Manager. off: Do not output to the operation log of NEC Storage Manager. (Default)

* If a value other than the above is set as upperlimit or lowerlimit, creating a volume fails.

If a value other than the above is set as upperreport, the system behavior is same as when off is specified.

(*1) If omitted, the I/O upper and lower limits are not set.

However, if the initial value of I/O control has been set by the `iSMioc setdfllimit` command, the set initial value is used.

* The `iSMioc setdfllimit` command is used to set the initial value of I/O control.

For more information, see “Storage Manager Command Reference.”

(*2) If the initial value of I/O control has been set by the `iSMioc setdfllimit` command and 0 is set as both upperlimit and lowerlimit, the I/O upper and lower limits are not set.

Operation example:

1. Create a volume type.

```
$ cinder type-create Gold
+-----+-----+
|                ID                | Name |
+-----+-----+
| d2cf1173-4a9b-4782-8b92-b9c33a52a3b3 | Gold |
+-----+-----+
```

2. Create the QoS specifications.

```
$ cinder qos-create QoSGold upperlimit=3000 lowerlimit=200
+-----+-----+
| Property | Value |
+-----+-----+
| consumer | back-end |
| id       | 67474e83-6fb2-4f1e-99bf-61b534b7cb6d |
| name     | QoSGold |
| specs    | {u'upperlimit': u'3000', u'lowerlimit': u'200'} |
+-----+-----+
```

3. Associate the QoS specifications and the volume type.

```
$ cinder qos-associate 67474e83-6fb2-4f1e-99bf-61b534b7cb6d
d2cf1173-4a9b-4782-8b92-b9c33a52a3b30
```

4. Specify the volume type to create a 100 GB volume.

```
$ cinder create 100 --volume-type Gold
```

* 3000 and 200 are set as the I/O upper and lower limits, respectively, of the created volume.

5. Change the I/O upper limit (upperlimit). (3000 → 5000) (*3)

```
$ cinder qos-key 67474e83-6fb2-4f1e-99bf-61b534b7cb6d set
upperlimit=5000
```

6. Check the QoS specifications after changing the I/O upper limit (upperlimit).

```
$ cinder qos-show 67474e83-6fb2-4f1e-99bf-61b534b7cb6d
+-----+-----+
| Property | Value |
+-----+-----+
| consumer | Value |
| id       | 67474e83-6fb2-4f1e-99bf-61b534b7cb6d |
| name     | QoSGold |
| specs    | {u'upperlimit': u'5000', u'lowerlimit': u'200'} |
+-----+-----+
```

7. Specify the volume type to create a 50 GB volume.

```
$ cinder create 50 --volume-type Gold
```

* 5000 and 200 are set as the I/O upper and lower limits, respectively, of the created volume.

8. Delete the I/O lower limit (lowerlimit). (*3)

```
$ cinder qos-key 67474e83-6fb2-4f1e-99bf-61b534b7cb6d unset
lowerlimit
```

9. Check the QoS specifications after deleting the I/O lower limit (lowerlimit).

```
$ cinder qos-show 67474e83-6fb2-4f1e-99bf-61b534b7cb6d
+-----+-----+-----+
| Property |          Value          |
+-----+-----+-----+
| consumer |          Value          |
|   id    | 67474e83-6fb2-4f1e-99bf-61b534b7cb6d |
|  name   |           QoSGold      |
|  specs  | {u'upperlimit': u'5000'|
+-----+-----+-----+
```

10. Specify the volume type to create a 10 GB volume.

```
$ cinder create 10 --volume-type Gold
```

* 5000 is set as the I/O upper limit of the created volume.

(*3) The I/O upper and lower limits of a volume that has already been created will not change.

For how to change the I/O upper and lower limits by using NEC Storage Manager, refer to the following:

- Example of how to change the I/O upper and lower limits by using NEC Storage Manager

1. Check the QoS specifications list.

```
$ cinder qos-list
+-----+-----+-----+-----+
|          ID          | Name | Consumer |          specs          |
+-----+-----+-----+-----+
| 67474e83-6fb2-4f1e-99bf-61b534b7cb6d | QoSGold | back-end | {u'upperlimit': u'3000'} |
+-----+-----+-----+-----+
```

2. Check the volume type that is associated with the QoS specifications (QoSGold).

```
$ cinder qos-get-association 67474e83-6fb2-4f1e-99bf-61b534b7cb6d
+-----+-----+-----+-----+
| Association_Type | Name |          ID          |
+-----+-----+-----+-----+
| volume_type     | Gold | f02a7d95-c45e-4f6f-a63f-8ae4a5d9477a |
+-----+-----+-----+-----+
```

3. Acquire the ID of the volume that is associated with the volume type (Gold).

```
$ cinder list | grep Gold
| 00c500b7-fc55-4316-918e-daf9e3e4d87f | available | vol_0147 | 1 | Gold |
| 030c795e-dc01-4dc4-995f-feebd8a70aff | available | vol_0275 | 1 | Gold |
```

4. Check the logical disk name of the disk array that is associated with the target volume ID.

- RHEL

```
$ cd /usr/lib/python2.7/site-packages/cinder/volume/drivers/nec
```

- Ubuntu

```
$ cd /usr/local/lib/python2.7/dist-packages/cinder/volume/drivers/nec
```

- Change the name to the logical disk name (LD Name) of the disk array.

```
$ python volume_list.py 00c500b7-fc55-4316-918e-daf9e3e4d87f
+-----+-----+
|          ID          |      LD Name      |
+-----+-----+
| 00c500b7-fc55-4316-918e-daf9e3e4d87f | 1S1tNffqZAOqDPRkbzIcR |
+-----+-----+
```

```
$ python volume_list.py 030c795e-dc01-4dc4-995f-feebd8a70aff
+-----+-----+
|          ID          |      LD Name      |
+-----+-----+
| 030c795e-dc01-4dc4-995f-feebd8a70aff | 5kg0vtWxoRfkSKY8ne5J1 |
+-----+-----+
```

5. Log in to the target disk array and change the I/O control values by using the `iSMioc setlimit` command.

Example:

```
# iSMioc setlimit -ldname LX:1S1tNffqZAOqDPRkbzIcR -upperlimit 5000
-lowerlimit 200
```

* Add the OS type (LX:) to the logical disk name.

* For more information about the `iSMioc setlimit` command, see “Storage Manager Command Reference.”

* For more information about the I/O control function, refer to the “I/O Control User’s Manual.”

Chapter 8 Correspondence between Cinder Volume IDs and Logical Disk Names in the Disk Array

The NEC Storage M series Cinder driver for OpenStack (Kilo) creates a logical disk name by converting the Cinder volume ID (UUID) into a base-62 number. Therefore, their correspondence is difficult to understand.

This chapter describes the command to output correspondence between the Cinder volume ID and logical disk name.

8.1 Command Specifications

Name	volume_list.pyc	
Syntax	<pre>python volume_list.pyc [list snapshot-list] [--all-tenants] python volume_list.pyc [volume-ID] python volume_list.pyc [--help -h] python volume_list.pyc [--version -v]</pre>	
Description	This command lists the volume IDs managed in Cinder and LD names managed in the storage device.	
Options	list	Lists the volume IDs managed in Cinder and LD names managed in the storage device.
	snapshot-list	Lists the snapshot IDs managed in Cinder and LD names managed in the storage device.
	--all-tenants	Specify this option to display the volumes or snapshots for all the tenants.
	volume-ID	Lists the specified volume ID and LD names managed in the storage device.
	--help -h	Displays the usage of command.
	--version -v	Displays the version of command.
Return values	This command does not return any value.	
Messages	Invalid UUID format	Argument is not uuid.
	Invalid argument	Argument count is illegal.

8.2 Setting of Environmental Variables

It is necessary to set the environmental variables to run the Cinder command.

Example:

```
export OS_USERNAME=admin
export OS_PASSWORD=openstack
export OS_TENANT_NAME=admin
export OS_AUTH_URL=http://localhost:5000/v2.0/
```

8.3 Command Execution Example

1. Move to the installation destination.

* For RHEL

```
$ cd /usr/lib/python2.7/site-packages/cinder/volume/drivers/nec
```

* For Ubuntu

```
$ cd /usr/local/lib/python2.7/dist-packages/cinder/volume/drivers/nec
```

2. Run the command (displaying all volumes).

```
$ python volume_list.pyc
```

```
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
| 1e5e4886-f394-41dc-b55d-d3f6a40612cc | vIr9qyun7XKLVoCOWAgYu |
| 5b85dfce-76bc-4e95-a943-6c22ac8e35dc | 2mhTuYrCGy8YnJcgRQBCPs |
| d1935c4b-f2c3-4f4a-8cfe-a46e95550e1e | 6NSjtgPWTieAMhvsJspHA6 |
| e2abcece-2863-485d-afa3-378cd43a8557 | 6timV9ZImmJlGxJzyoGgid |
| fd761e59-4940-47f4-baf3-1f49b2754de4 | 7iH2Q85fN8DTwFOIpKwoeC |
+-----+-----+
↑ Cinder volume ID          ↑ Logical disk name
```

3. Run the command (when no volume has been created).

```
$ python volume_list.pyc
```

```
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
+-----+-----+
```

4. Run the command (listing all the snapshots).

```
$ python volume_list.pyc snapshot-list
```

```
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
| 3d9d6a46-854b-490a-8a1c-e785f9d31d12 | 1sGTC0Rvk1CwAsE1YAYz4M |
| 9f0c0fb0-6f9e-4dd5-bel0-aaac80f39733 | 4q7I6V6T8c8BVFuE8UfWVP |
| c8ba6f90-698b-417e-afbc-755ad0ae5922 | 661fL9sXV64aGPxa4k2Yuw |
+-----+-----+
↑ Cinder snapshot ID          ↑ Logical disk name
```

5. Run the command (displaying a specific volume ID only).

```
$ python volume_list.pyc 6NSjtgPWTieAMhvsJspHA6
+-----+-----+
|          LD Name          |          ID          |
+-----+-----+
| 6NSjtgPWTieAMhvsJspHA6 | d1935c4b-f2c3-4f4a-8cfe-a46e95550e1e |
+-----+-----+
↑ Logical disk name      ↑ Cinder volume ID
```

6. Run the command (displaying a specific volume ID only).

```
$ python volume_list.pyc d1935c4b-f2c3-4f4a-8cfe-a46e95550e1e
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
| d1935c4b-f2c3-4f4a-8cfe-a46e95550e1e | 6NSjtgPWTieAMhvsJspHA6 |
+-----+-----+
↑ Cinder volume ID      ↑ Logical disk name
```

7. Run the command (listing all the volumes of all the tenants).

```
$ python volume_list.pyc --all-tenants
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
| 1e5e4886-f394-41dc-b55d-d3f6a40612cc | vIr9qyun7XKLVoCOWAgYu |
| 5b85dfce-76bc-4e95-a943-6c22ac8e35dc | 2mhTuYrCGy8YnJcgRQBCPs |
| d1935c4b-f2c3-4f4a-8cfe-a46e95550e1e | 6NSjtgPWTieAMhvsJspHA6 |
| e2abcece-2863-485d-afa3-378cd43a8557 | 6timV9ZImmJlqxJzyoGgid |
| fd761e59-4940-47f4-baf3-1f49b2754de4 | 7ih2Q85fN8DTwFOIpKWoeC |
| 72d4c228-c4f1-4117-87f0-4fd3147a7d22 | 3UgM53C4R8MrPgeCFcqvQA |
| e5d6f9f1-2e1f-4a96-a767-ba0aa7f8eb89 | 6zhUDKpLnOadJRZFTziLwP |
+-----+-----+
↑ Cinder volume ID      ↑ Logical disk name
```

8. Run the command (listing all the snapshots of all the tenants).

```
$ python volume_list.pyc snapshot-list --all-tenants
+-----+-----+
|          ID          |          LD Name          |
+-----+-----+
| 3d9d6a46-854b-490a-8a1c-e785f9d31d12 | 1sGTCORvk1CwAsElYAYz4M |
| 9f0c0fb0-6f9e-4dd5-be10-aaac80f39733 | 4q7I6V6T8c8BVFuE8UfWVP |
| c8ba6f90-698b-417e-afbc-755ad0ae5922 | 661fL9sXV64aGPxa4k2Yuw |
| 88333a78-ea45-4367-a4ca-c7306e7ef86e | 490MoH7rwmboMA49B9idiQ |
| 959135b5-90a5-4f7b-b20c-dbabf41da029 | 4YEDVATDPVGFY622WIKngv |
| f4331af0-7aae-4980-9cc7-b36c39898c37 | 7QnTwQMf4PSPg2EmkT9gW7 |
+-----+-----+
↑ Cinder snapshot ID    ↑ Logical disk name
```


8.4 Displaying the Usage

```
$ python volume_list.py --help
```

```
Version : 1.0
```

```
usage: python volume_list.py [--version] [-v] [--help] [-h]
        [UUID] [ldname] [list] [snapshot-list]
```

Show mappings of Cinder volume/snapshot ID and logical disk name in storage.

Optional arguments:

<code>--version, -v</code>	show program's version number and exit
<code>--help, -h</code>	show this help and exit
<code>UUID</code>	filter by uuid of Cinder volume/snapshot
<code>ldname</code>	filter by logical disk name in storage
<code>list</code>	show all volumes (default)
<code>snapshot-list</code>	show all snapshots

Notes:

The environment variables below must be set in advance.

`OS_USERNAME`, `OS_PASSWORD`, `OS_TENANT_NAME`, `OS_AUTH_URL`

8.5 Displaying the Version

```
$ python volume_list.py --version
```

```
Version : 1.0
```

```
$ python volume_list.py -v
```

```
Version : 1.0
```

Chapter 9 Multi-path Connection

9.1 Setting Up DM-Multipath for RHEL

- Installation of DM-Multipath is required if not installed yet.

```
# yum install device-mapper-multipath
```
- Create a multi-path configuration file (`/etc/multipath.conf`) by using the `mpathconf` utility.

```
# mpathconf -enable
```

* Set `user_friendly_names` of `/etc/multipath.conf` to `no`.

```
defaults {
    user_friendly_names no
    find_multipaths yes
}
```

* Add the following setting in the multi-path configuration file (`/etc/multipath.conf`). *1

```
devices {
    device {
        vendor "NEC"
        product "DISK ARRAY"
        path_grouping_policy group_by_prio
        prio alua
        hardware_handler "1 alua"
    }
}
```
- Enable auto startup of the `multipathd` service.

```
# systemctl enable multipathd
```
- Restart the `multipathd` service.

```
# systemctl restart multipathd
```

(*1) It is recommended to set a path grooming policy according to the ALUA-based path priority.

9.2 Setting Up multipath-tools for Ubuntu

- Installation of multipath-tools is required if not installed yet.

```
# apt-get install multipath-tools sysv-rc-conf
```

- Create a multi-path configuration file (/etc/multipath.conf).

```
# vi /etc/multipath.conf
```

* Add the following setting in the multi-path configuration file (/etc/multipath.conf). *1

```
defaults{
    user_friendly_names no
}

devices {
    device {
        vendor "NEC"
        product "DISK ARRAY"
        path_grouping_policy group_by_prio
        prio alua
        hardware_handler "1 alua"
    }
}
```

- Enable auto startup of the multipath-tools service.

```
# sysv-rc-conf multipath-tools on
```

- Restart the multipath-tools service.

```
# service multipath-tools restart
```

(*1) It is recommended to set a path grooming policy according to the ALUA-based path priority.

9.3 Enabling the Nova Multi-path Function

The Nova multi-path function is enabled according to the following procedure.

(Only for iSCSI interfaces)

- Add the following description to the `[libvirt]` section in `nova.conf`.

```
iscsi_use_multipath = True
```

- Restart `nova-compute`.

```
service openstack-nova-compute restart
```

9.4 Checking the Multi-path Configuration

After attaching a volume to an instance, check the multi-path configuration by using the `multipath` command.

- Example 1: When `iscsi_use_multipath` in `nova.conf` is `False` (multi-path disabled)

```
# multipath -ll
200255c3a21cc0000 dm-3 NEC,DISK ARRAY
size=1.0G features='0' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=130 status=active
   `- 10:0:0:0 sdb 8:16 active ready running
```

⇒ There is a single path.

- Example 2: When `iscsi_use_multipath` in `nova.conf` is `True` (multi-path enabled)

```
# multipath -ll
200255c3a21cc0000 dm-3 NEC,DISK ARRAY
size=1.0G features='0' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=130 status=active
| `- 11:0:0:0 sdb 8:16 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
   `- 12:0:0:0 sdc 8:32 active ready running
```

⇒ There are multiple paths.

- Example 3: If a fault occurs on a path

```
# multipath -ll
mpathb (200255c3a21cc0000) dm-3 NEC,DISK ARRAY
size=1.0G features='0' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=0 status=enabled
| `- 17:0:0:0 sdb 8:16 failed faulty running
`-+- policy='round-robin 0' prio=50 status=active
   `- 18:0:0:0 sdc 8:32 active ready running
```

* Check the system log (`/var/log/messages`).

```
Oct 15 17:12:55 host01 multipathd: mpathb: sdb - tur checker reports
path is down
```

⇒ `multipathd` detects that the path went down.

```
Oct 15 17:12:55 host01 multipathd: mpathb: remaining active paths: 1
```

⇒ `multipathd` detects that there is only one active path.

- Example 4: If faults occur on all the paths

```
# multipath -ll
mpathb (200255c3a21cc0000) dm-3 NEC,DISK ARRAY
size=1.0G features='0' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=0 status=enabled
| `-- 17:0:0:0 sdb 8:16 failed faulty running
`+- policy='round-robin 0' prio=0 status=enabled
  `-- 18:0:0:0 sdc 8:32 failed faulty running
```

* Check the system log (/var/log/messages).

```
Oct 15 17:23:52 host01 multipathd: mpathb: sdc - tur checker reports
path is down
```

⇒ multipathd detects that the path went down.

```
Oct 15 17:23:52 host01 multipathd: mpathb: remaining active paths: 0
```

⇒ multipathd detects that there is only one active path.

- * **If faults occur on all the paths, a writing error may occur when accessing the relevant volume after the paths are restored.**

In this case, unmount the relevant volume after restoring the path and mount it again before accessing it.

- Example 5: If either of the paths is restored

```
# multipath -ll
mpathb (200255c3a21cc0000) dm-3 NEC,DISK ARRAY
size=1.0G features='0' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=130 status=active
| `-- 17:0:0:0 sdb 8:16 active ready running
`+- policy='round-robin 0' prio=0 status=enabled
  `-- 18:0:0:0 sdc 8:32 failed faulty running
```

* Check the system log (/var/log/messages).

```
Oct 15 17:32:11 host01 multipathd: mpathb: sdb - tur checker reports
path is up
```

⇒ multipathd detects that a path has been restored.

```
Oct 15 17:32:11 host01 multipathd: mpathb: remaining active paths: 1
```

⇒ multipathd detects that there is only one active path.

Chapter 10 Snapshot Function Supporting Differential Backup

- * Please read this chapter when using the data replication method.
- * An operational overview of this function is provided on the following pages.

10.1 Creating a Snapshot

10.1.1 When no RV maintaining differences is paired with the target copy source logical disk

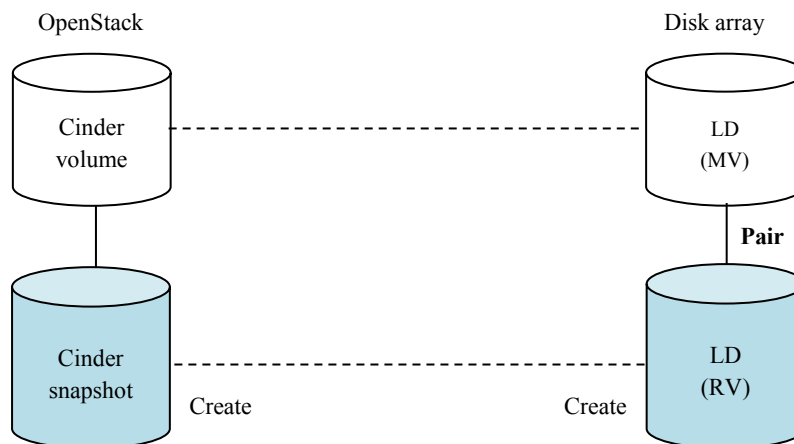
Before running the command:



- Create new RVs, specify the pair setting, and then execute replicate (fully copy) and separate operations.

The pair setting is maintained after the separate operation is complete.

After running the command:



Note: Up to three RVs (three generations) can be created for a single MV.

If an attempt is made to create a fourth (generation of) snapshot, an error occurs.

* Take measures on the upper level management machine not to create more than three generations of snapshots.

* The snapshot resulting in an error can be deleted by using the `cinder snapshot-delete` command.


* RV maintaining differences

RV that has `_d`, which means a backup target, at the end of its logical disk name.

Although the snapshot was deleted from OpenStack, the logical disk (RV) is not deleted from the disk array and maintains the pair setting with an MV.

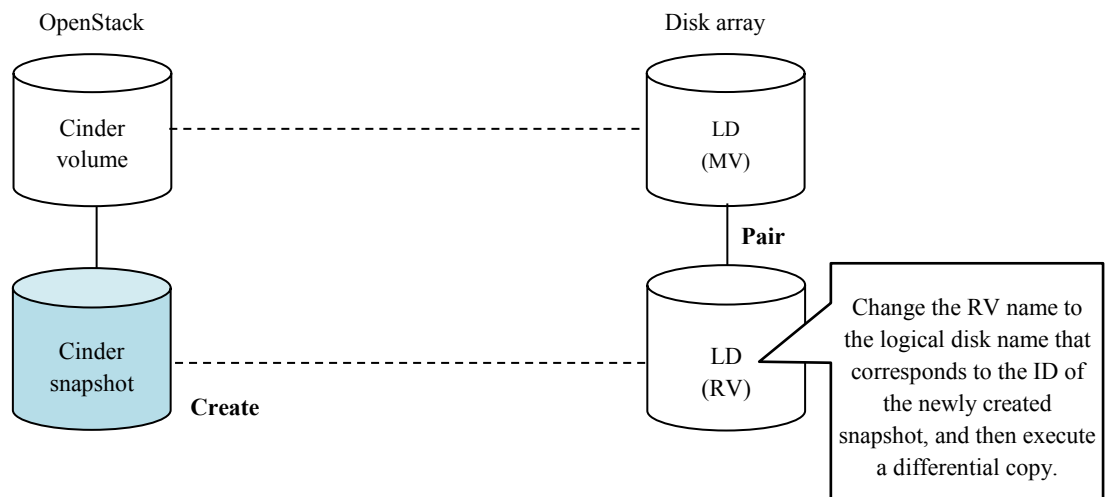
10.1.2 When an RV maintaining differences is paired with the copy source logical disk (MV)

Before running the command:



- Change the RV name to the logical disk name that corresponds to the ID of the newly created snapshot, and then execute replicate (differential copy) and separate operations. The pair setting is maintained after the separate operation is complete.

After running the command:




* If there are multiple RVs maintaining differences, carry out this process on one of them (selected by the Cinder driver).

10.2 Deleting a Snapshot

10.2.1 When the target logical disk (RV) is retained as an RV maintaining differences

Before running the command:

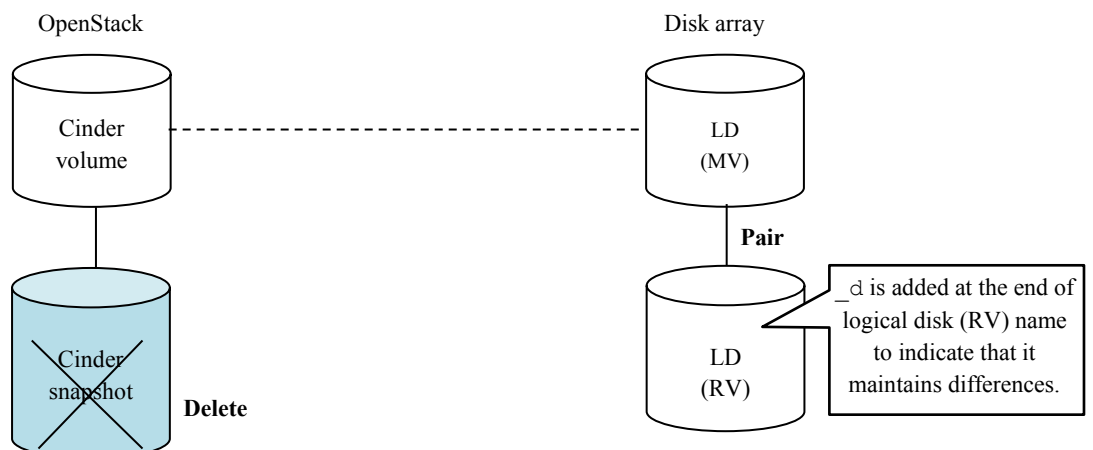


- Delete only the snapshot on OpenStack.

The logical disk (RV) on the disk array is not deleted and remains paired with an MV.


`_d` is added at the end of the logical disk (RV) name to indicate that it maintains differences.

After running the command:



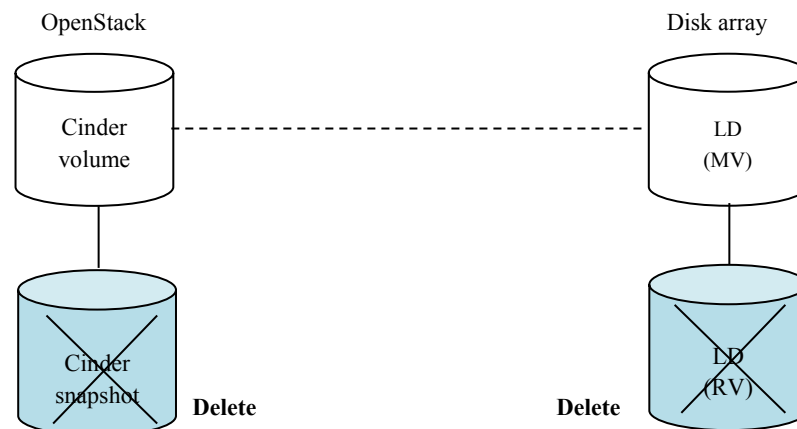
10.2.2 When deleting the target logical disk (RV)

Before running the command:



- Delete the snapshot by setting metadata (`force-delete=True`) for the snapshot by using the `snapshot-metadata` command. This deletes the logical disk (RV) from the disk array at the same time.*1

After running the command:



*1: Command example:

```
$ cinder snapshot-metadata 720847df-b37b-4320-9b97-c2a23d4e9309 set force-
delete=True
```

```
+-----+
| Property | Value |
+-----+
| force-delete | True |
+-----+
```

```
$ cinder snapshot-delete 720847df-b37b-4320-9b97-c2a23d4e9309
```


10.2.3 When the target logical disk is not paired with another logical disk

Before running the command:



- Delete the target logical disk regardless of whether metadata (`force-delete=True`) is set.

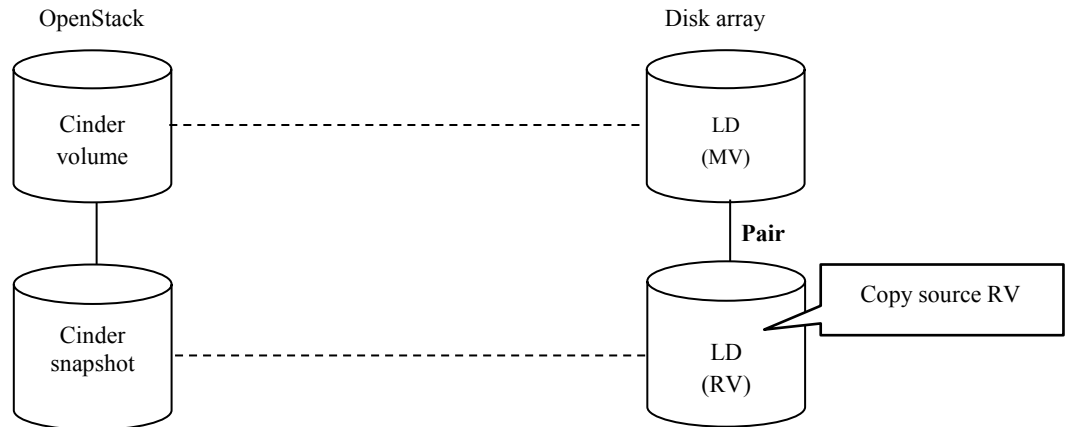
After running the command:



10.3 Creating a Volume from a Snapshot

10.3.1 When an MV is paired with the copy source logical disk (RV)

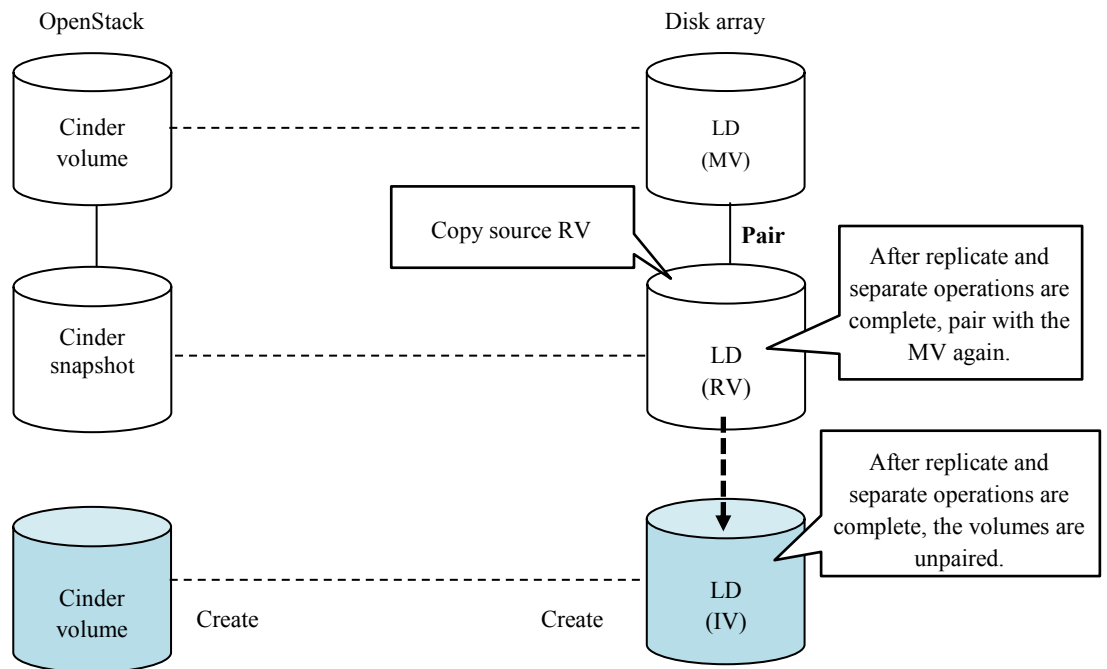
Before running the command:



- Unpair the RV and MV, pair the newly built logical disk and the copy source logical disk (RV), and then execute replicate (fully copy) and separate operations.

After the separate operation is complete, unpair the volumes, and then pair the RV and MV again.^{*1,*2}

Before running the command:



- *1: If snapshots of the MV are created and the number of RVs reaches the upper limit while running this command, the MV and RV will not be re-paired.
The message indicating that re-pairing the volumes failed is output to the log.
Snapshots without the pair setting can be deleted by using the `snapshot-delete` command.

- *2: Replicate operation is not performed if the re-pair setting is specified. Therefore, if the re-pair setting is specified (difference information is lost), a full copy will be executed rather than a differential copy at the next replication operation (snapshot creation) for the relevant RV.

10.3.2 When no MV is paired with the copy source logical disk

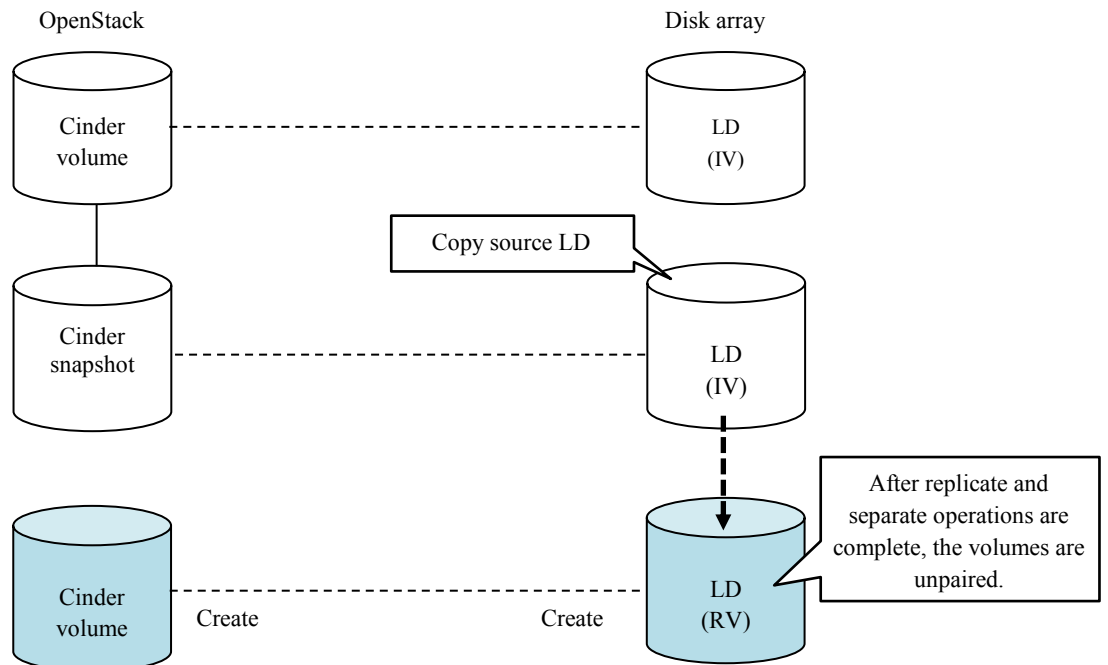
Before running the command:



- Bind a new logical disk, pair it with the copy source logical disk, and then execute replicate (fully copy) and separate operations.

After the separate operation is complete, the volumes are unpaired.

After running the command:



10.4 Deleting a Volume

10.4.1 When no RV is paired with the target logical disk

Before running the command:



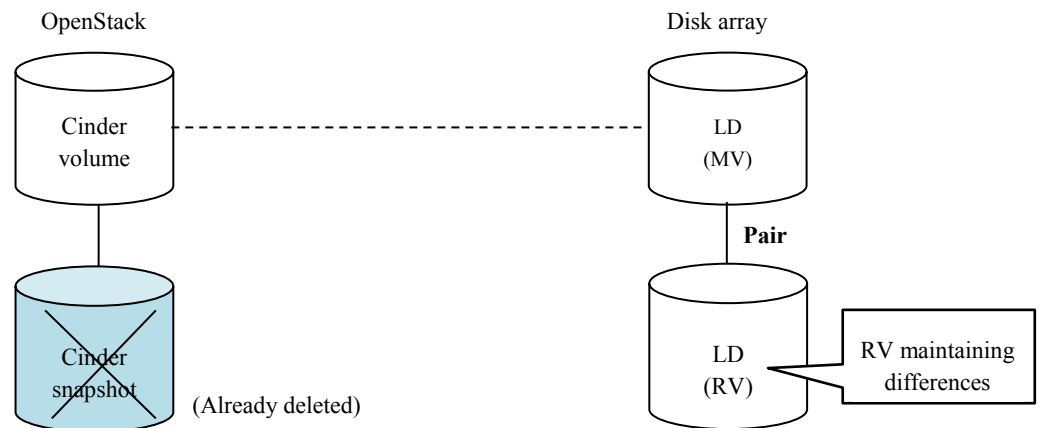
- Delete the target volume and logical disk.

After running the command:



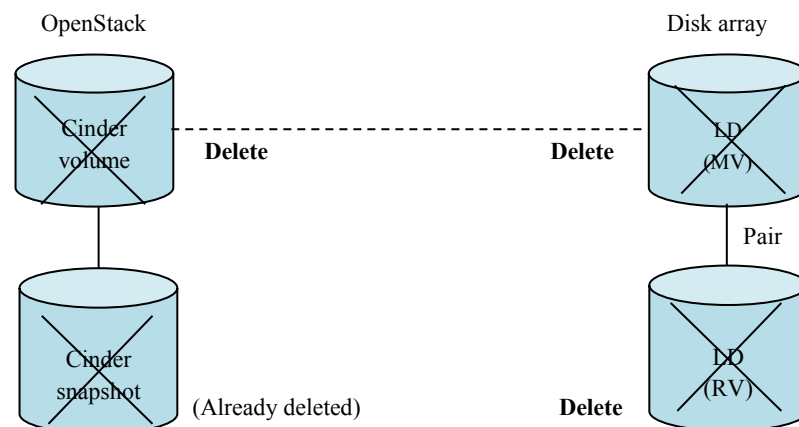
10.4.2 When an RV maintaining differences is paired with the target logical disk (The snapshot paired with the RV is already deleted.)

Before running the command:



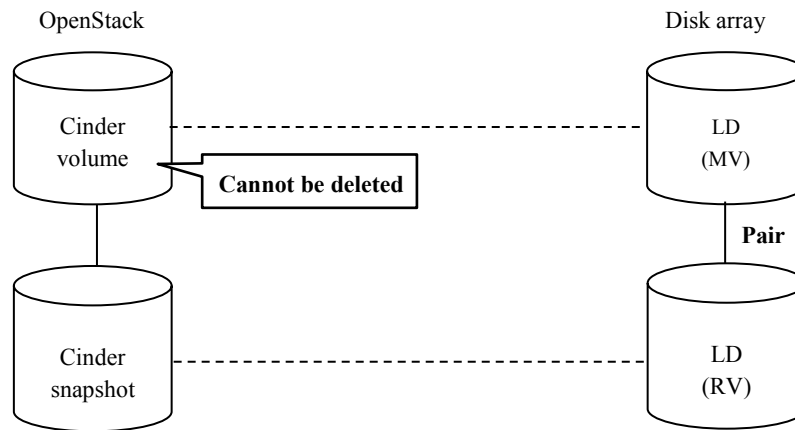
- Delete the target volume and logical disk, and RV maintaining differences that is paired with the target logical disk.

After running the command:



10.4.3 When a snapshot has been created from the target volume

Before running the command:



- The volume cannot be deleted because there is a dependent snapshot. If an attempt is made to delete the volume, an error occurs on the Cinder manager. * This is the OpenStack specifications.

Chapter 11 `get_volume_stats` Method

11.1 Overview

- Grouping by the capacity filter between the config-groups with the same `volume_backend_name` is allowed by returning the free capacity (`free_capacity_gb`) to the Cinder manager by using the `get_volume_stats` method implemented in this Cinder driver.

11.2 Timing When the `get_volume_stats` Method Is Called

- This method is called at one-minute intervals from the Cinder manager. This method runs asynchronously with other methods (such as volume creation). * This is the OpenStack specifications.

11.3 Return Values

- `total_capacity_gb`: Capacity of the pool that has the largest free capacity (GB)
- `free_capacity_gb`: Free capacity of the pool that has the largest free capacity (GB)
- `reserved_percentage`: 0
- `volume_backend_name`: `backend_name` or
`'MStorageISCSIDriver'` or
`'MStorageFCDriver'` or
`'MStorageDDRISCSIDriver'` or
`'MStorageDDRFCDriver'`
- `vendor_name`: `'NEC'`
- `driver_version`: `'1.7.1' *1`
- `storage_protocol`: `'iSCSI'` or `'FC'`
- `QoS_support`: `True`

*1 “1.5.1” indicates the Cinder driver version. The above file name is just an example.

11.4 Selecting the Pool with the Largest Free Capacity

- The pool with the largest free capacity is selected from the pools specified for `nec_pools` in `cinder.conf`.
- If `nec_pools` is omitted in `cinder.conf`, select the pool with the largest free capacity from all the pools.

11.5 Return Values of the Pool Capacity and Free Pool Capacity

- The values of the following items are returned.
`total_capacity_gb:` Pool capacity
`free_capacity_gb:` Free pool capacity

11.6 Considerations

- Because the `get_volume_stats` method runs asynchronously with other methods (such as volume creation), the Cinder driver volume creation process may be called before the Cinder manager detects insufficient free capacity, causing an error in the Cinder driver depending on the timing. Therefore, keep sufficient free capacity during operation.
* The reason for error can be checked in the log message.

Appendix A Notes

- Start the NEC Storage M series disk array before starting the cinder node.
- Make sure that the services necessary for volume operations are running, especially when the services are restarted after the settings have been changed.

- Snapshot reserve area

If the snapshot reserve area is insufficient when using the snapshot method, the snapshot image cannot be retained and the backup data allocated in the SV is lost. Therefore, it is necessary to allocate enough space to the snapshot reserve area, as well as to monitor the snapshot reserve area by using NEC Storage Manager so that the area will not be insufficient, and expand the area if necessary.

When the snapshot reserve area runs out of free space, delete failed volumes manually by using CLI.

* For more information, refer to the following manuals:

- Snapshot User's Manual (Function Guide)
 - Snapshot reserve area management
- Storage Manager Command Reference
 - iSMcfg sraexpand (Expanding the snapshot reserve area (SRA))

Appendix B Optimal Path

The initial optimal path is specified for each pool by a disk array.

Create a pool so that optimal paths are distributed to the controllers properly.

* The optimal paths are determined as follows according to the pool number.

- If the pool number is even, the optimal path is CONT0.
- If the pool number is odd, the optimal path is CONT1.

The optimal path can be changed for each logical disk.

For more information, refer to the description of the `isMcfg ldchgattr` command (`-owner` option) in the “Storage Manager Command Reference.”

Appendix C Setting Full-access Mode

Full-access mode allows 16 or more servers (initiators) to access one LD Set. To use full-access mode, create an LD Set, and register/delete an initiator by using the procedures described in this chapter.

C.1 Restriction when using Full-access Mode

Restrictions on using Full-access Mode.

- You need to create a Full-access Mode LD set and specify the name of the LD set for the `nec_ldset` parameter in `cinder.conf` to use Full-access mode.
- Complete iSCSI RemoteDataReplication hosts configuration before creating Full-access Mode LD set. The iSCSI RemoteDataReplication hosts configuration menu is inaccessible if one or more Full-access Mode LD sets exist.
- Use command-line interface described below in this chapter to configure LD sets. GUI does not support Full-access Mode LD sets.
- Create a Full-access Mode LD set before registering iSCSI initiators which access Full-access Mode LD set. The initiators may be unregistered unexpectedly if no Full-access Mode LD set exists.
- Follow the procedure "Nondisruptive update by commands" in *Update Guide* to nondisruptively update Storage Control Software when Full-access mode is used. While updating, the command "`iSMcfg arrayupdate -mode prepare -file ooooo.inf`" may exit abnormally by timeout error (iSM31222). In this case, ignore the error and continue the procedure. While some path failures may be detected by Device-Mapper Multipath, these errors do not impact application server I/O.

C.2 Creating an LD Set

- Set on (multi-target mode) to the `-multitarget` option.
- Specify the `-fullaccess` option (full-access mode).

Example:

```
iSMcfg addldset -ldset LX:OpenStack0 -type iscsi -multitarget on -fullaccess
```

- * Be sure to create an LD Set before registering an initiator.
- * The `-fullaccess` option of the `iSMcfg addldset` command is not described in the "Storage Manager Command Reference" corresponding to NEC Storage Manager Version 9.3 or earlier.

C.3 Registering an Initiator

- To enable full-access mode when creating an LD Set, register an initiator by using the `iSMcfg addinitiator` command.

Example:

```
iSMcfg addinitiator -initiatorname iqn.1994-05.com.redhat:b3349cf7eb7
```

C.4 Deleting an Initiator

- To enable full-access mode when creating an LD Set, delete an initiator by using the `iSMcfg delinitiator` command.

Example:

```
iSMcfg delinitiator -initiatorname iqn.1994-05.com.redhat:b3349cf7eb7
```

C.5 Checking the Registered Initiator

- To check the initiators that have been registered to the disk array, check “Initiator List” by using the `iSMview -a` command.

Example:

```
iSMview -a
```

```
--- Access Control Information ---
```

```
Access Control Mode : off
Licensed Path Count : nolim
Used Path Count(FC)   : 0
Used Path Count(iSCSI) : 0
Used Path Count(Total) : 0
Registered Initiator Count : 3
```

```
--- LD Set(iSCSI) List ---
```

Platform	LD Set Name	Initiator Count	LD Count
LX	OpenStack0	0	0
LX	OpenStack1	0	0

```
--- Initiator List ---
```

```
Initiator Name
iqn.1994-05.com.redhat:efa044b15e8
iqn.1994-05.com.redhat:2fbd4c154b7
iqn.1994-05.com.redhat:289ee7a72fd
```


C.6 Checking the Full Access Mode Setting

- To check whether the full access mode is enabled, use the `iSMview -ac` or `iSMview -a` command. If the `iSMview -ac` command is executed for the full access mode LD Set, `on` is displayed for Full Access Mode of LD Set (iSCSI) Information.

Example:

```
iSMview -ac LX: iSCSI

--- LD Set(iSCSI) Information ---
Platform           : LX
LD Set Name        : iSCSI
Configuration Change :
Action mode        : Normal mode
Full Access Mode   : on
```

- If the `iSMview -a` command is executed for the full access mode LD Set, `0` is still displayed for Initiator Count of LD Set(iSCSI) List, and the number of initiators is displayed for Registered Initiator Count.

Example:

```
iSMview -a

--- Access Control Information ---
Access Control Mode : off
Licensed Path Count : nolim
Used Path Count(FC) : 0
Used Path Count(iSCSI) : 0
Used Path Count(Total) : 0
Registered Initiator Count : 3

--- LD Set(iSCSI) List ---
Platform  LD Set Name      Initiator Count  LD Count
LX        OpenStack0      0                0
LX        OpenStack1      0                0

--- Initiator List ---
```

Initiator Name

iqn.1994-05.com.redhat:efa044b15e8

iqn.1994-05.com.redhat:2fbd4c154b7

iqn.1994-05.com.redhat:289ee7a72fd

Appendix D Log Messages

D.1 ERROR Level Messages

Message	Explanation of message
'<iSMCLI command line>' failed. status=<exit_status>,out="<stdout>", err="<stderr>". <exit_status>: Exit status <stdout>: Contents of standard output <stderr>: Standard error output	The iSMCLI command resulted in an error. * For more information about errors, refer to the <i>NEC Storage Manager Messages Handbook</i> (IS010).
All Logical Disk numbers are used.	The number of volumes exceeds the maximum number of volumes that can be created.
Specified Logical Disk <LD name> has been copied. <LD name> : Logical Disk name (nickname)	The target volume cannot be operated because it is being replicated.
Logical Disk Set could not be found.	The target LD Set could not be detected.
<tag>[@name="<attr>"] not found. line=<line> out="<xml>" <tag> : XML tag <attr> : XML attribute <line> : Line number in the XML file <xml> : XML file contains the storage configuration information	A <tag> tag whose name attribute is <attr> could not be found in the storage configuration information (XML).
Not enough pool capacity. pool_number=<pool>, size_increase=<size> <pool> : Target pool number <size> : The required free size of the pool	There is not enough free space in the target pool.
Logical Disk has unbound already (name=<name>, id=<id>). <name> : Logical Disk name (nickname) <id> : Target volume ID	The logical disk corresponding to the specified volume ID could not be found. The target logical disk has already been deleted. (This message is output to the log but volume deletion does not result in an error.)
Logical Disk `<name>` has unbound already. <name> : Logical Disk name (nickname)	The LD specified with <name> was not found when the volume is deleted. (This message is output to the log but volume deletion

	does not result in an error.)
RPL Attribute Error. RPL Attribute = <RPL Attribute>. <RPL Attribute> : RPL type	The attribute of the volume to be backed up or restored is invalid. (Check the status with the iSM client, change [RPL Type] to [IV], and then run the command.)
Cannot create clone volume. number of pairs reached 3. ldbname=<name> <name> : Logical Disk name (nickname)	A clone cannot be created because the maximum number of RVs (three RVs) have already been created for the target MV.
Cannot create snapshot. number of pairs reached 3. ldbname=<name> <name> : Logical Disk name (nickname)	A snapshot cannot be created because the maximum number of RVs (three RVs) have already been created for the target MV.
Specified Logical Disk <name> is not available. <name> : Logical Disk name (nickname)	The specified volume is in unusable state.
Logical Disk `<name>` does not exist. <name> : Logical Disk name (nickname)	The target logical disk could not be found. (The volume could have been deleted by other than OpenStack.)
backup_pools is not set.	nec_backup_pools is not set. (This must be always set when using the backup function.)
Cannot create volume from snapshot, because the snapshot data does not exist. bvname=<bvname>, svname=<svname> <bvname> : BV name (nickname) <svname> : SV name (nickname)	Since the snapshot data has been deleted, a volume cannot be created by using this snapshot.
Cannot delete snapshot. Logical Disk=<name> <name> : Logical Disk name (nickname)	The snapshot of the target volume cannot be deleted because it could not be found.

D.2 WARNING Level Messages

Message	Explanation of message
<p>Failed to Create Volume (Volume ID = <id>, Size = <size>GB) (<message>)</p> <p><id> : Target volume ID <size> : Target volume capacity <message> : Detail information message</p>	Creating a volume failed.
<p>Failed to Extend Volume (Volume ID = <id>, New Size = <size>GB, Old Size = <oldsize>GB) (<message>)</p> <p><id> : Target volume ID <size> : After extended target volume capacity <oldsize> : Before extended target volume capacity <message> : Detail information message</p>	Expanding the volume failed.
<p>Failed to Create Cloned Volume (Volume ID = <id>, Source Volume ID = <source id>) (<message>)</p> <p><id> : Target volume ID <source id> : Source volume ID <message> : Detail information message</p>	Creating a clone volume failed.
<p>Failed to Migrate Volume (Volume ID = <id>, Destination Host = <host>) (<message>)</p> <p><id> : Target volume ID <host> : Destination host information <message> : Detail information message</p>	Migrating the volume failed.
<p>Failed to Create iSCSI Export (Volume ID = <id>, Initiator Name = <name>) (<message>)</p> <p><id> : Target volume ID <name> : Initiator name of exporting target <message> : Detail information message</p>	Exporting the volume (via iSCSI) failed.

<p>Failed to Create FC Export (Volume ID = <id>, Initiator WWPNs = <wwpn>) (<message>)</p> <p><id> : Target volume ID <wwpn> : World Wide Port Name of exporting target <message> : Detail information message</p>	Exporting the volume (via FC) failed.
<p>Failed to Remove Export (Volume ID = <id>) (<message>)</p> <p><id> : Target volume ID <message> : Detail information message</p>	Canceling exporting the volume failed.
<p>Failed to Initialize iSCSI Connection (Volume ID = <id>, Connector = <info>) (<message>)</p> <p><id> : Target volume ID <info> : Connection host information <message> : Detail information message</p>	Connecting to the volume (via iSCSI) failed.
<p>Failed to Terminate iSCSI Connection (Volume ID = <id>, Connector = <info>) (<message>)</p> <p><id> : Target volume ID <info> : Connection host information <message> : Detail information message</p>	Disconnecting the volume (from iSCSI) failed.
<p>Failed to Initialize FC Connection (Volume ID = <id>, Connector = <info>) (<message>)</p> <p><id> : Target volume ID <info> : Connection host information <message> : Detail information message</p>	Connecting to the volume (via FC) failed.
<p>Failed to Terminate FC Connection (Volume ID = <id>, Connector = <info>) (<message>)</p> <p><id> : Target volume ID <info> : Connection host information <message> : Detail information message</p>	Disconnecting the volume (from FC) failed.
<p>Failed to Delete Volume (Volume ID = <id>) (<message>)</p> <p><id> : Target volume ID <message> : Detail information message</p>	Deleting the volume failed.

<p>Failed to Create Snapshot (Snapshot ID = <id>s, Snapshot Volume ID = <snap vol id>) (<message>)</p> <p><id> : Target snapshot ID <snap vol id> : Target snapshot volume ID <message> : Detail information message</p>	<p>Creating a snapshot failed.</p>
<p>Failed to Create Volume from Snapshot (Volume ID = <id>, Snapshot ID = <snap id>, Snapshot Volume ID = <snap vol id>) (<message>)</p> <p><id>: Target volume ID <snap id> : Source snapshot ID <snap vol id> : Source snapshot volume ID <message> : Detail information message</p>	<p>Creating a volume from the snapshot failed.</p>
<p>Failed to Delete Snapshot (Snapshot ID = <id>, Snapshot Volume ID = <snap vol id>) (<message>)</p> <p><id> : Target snapshot ID <snap vol id> : Target snapshot volume ID <message> : Detail information message</p>	<p>Deleting the snapshot failed.</p>

D.3 INFO Level Messages

Message	Explanation of message
<p>Created Volume (Volume ID = <id>, Size = <size>GB)</p> <p><id> : Target volume ID <size> : Target volume capacity</p>	A volume has been created.
<p>Extended Volume (Volume ID = <id>, New Size = <size>GB, Old Size = <oldsize>GB)</p> <p><id> : Target volume ID <size> : After extended target volume capacity <oldsize> : Before extended target volume capacity</p>	The volume has been expanded.
<p>Created Cloned Volume (Volume ID = <id>, Source Volume ID = <source id>)</p> <p><id> : Target volume ID <source id> : Source volume ID</p>	A clone volume has been created.
<p>Migrated Volume (Volume ID = <id>, Destination Host = <host>)</p> <p><id> : Target volume ID <host> : Destination host information</p>	The volume has been migrated.
<p>Created iSCSI Export (Volume ID = <id>, Initiator Name = <name>)</p> <p><id> : Target volume ID <name> : Initiator name of exporting target</p>	The volume has been exported (via iSCSI).
<p>Created FC Export (Volume ID = <id>, Initiator WWPNs = <wwpn>)</p> <p><id> : Target volume ID <wwpn> : World Wide Port Name of exporting target</p>	The volume has been exported (via FC).
<p>Removed Export (Volume ID = <id>)</p> <p><id> : Target volume ID</p>	Exporting the volume has been canceled.

<p>Initialized iSCSI Connection (Volume ID = <id>, Connector = <info>)</p> <p><id> : Target volume ID <info> : Connection host information</p>	The volume has been connected (via iSCSI).
<p>Terminated iSCSI Connection (Volume ID = <id>, Connector = <info>)</p> <p><id> : Target volume ID <info> : Connection host information</p>	The volume has been disconnected (from iSCSI).
<p>Initialized FC Connection (Volume ID = <id>, Connector = <info>)</p> <p><id> : Target volume ID <info> : Connection host information</p>	The volume has been connected (via FC).
<p>Terminated FC Connection (Volume ID = <id>, Connector = <info>)</p> <p><id> : Target volume ID <info> : Connection host information</p>	The volume has been disconnected (from FC).
<p>Deleted Volume (Volume ID = <id>)</p> <p><id> : Target volume ID</p>	The volume has been deleted.
<p>Created Snapshot (Snapshot ID = <id>s, Snapshot Volume ID = <snap vol id>)</p> <p><id> : Snapshot ID <snap vol id> : Snapsthot volume ID</p>	A snapshot has been created.
<p>Created Volume from Snapshot (Volume ID = <id>, Snapshot ID = <snap id>, Snapshot Volume ID = <snap vol id>)</p> <p><id>: Target volume ID <snap id> : Source snapshot ID <snap vol id> : Source snapshot volume ID</p>	A snapshot has been created from the volume.
<p>Deleted Snapshot (Snapshot ID = <id>, Snapshot Volume ID = <snap vol id>)</p> <p><id> : Target snapshot ID <snap vol id> : Target snapshot volume ID</p>	The snapshot has been deleted.

Appendix E Collecting Failure Information

E.1 Log File and Configuration File of the Cinder Driver

The Cinder driver log is output to the directory specified in `cinder.conf`. When a fault occurs, collect the relevant log file and configuration file (`cinder.conf`) in the specified directory.

To output debug information, specify `True` for the `debug` parameter.

`cinder.conf` setting example

```
[DEFAULT]
log_dir=/var/log/cinder/
debug=True
```

E.2 NEC Storage Manager Logs

Collect fault information by using the `ismgather` command.

* For details, see the *NEC Storage Manager Command Reference* (IS052).

Appendix F Troubleshooting

Q. cinder-volume does not start.

A.1. Check whether the settings for Cinder are correct.

- If the following message is output to the Cinder log, set `nec_ismcli_fip` in `cinder.conf`.

```
MStorageISCSIDriver: nec_ismcli_fip is not set.
```
- If the following message is output to the Cinder log, set `nec_ismcli_user` in `cinder.conf`.

```
MStorageISCSIDriver: nec_ismcli_user is not set.
```
- If the following message is output to the Cinder log, set `nec_ismcli_password` or `nec_ismcli_privkey` in `cinder.conf`.

```
MStorageISCSIDriver: nec_ismcli_password or nec_ismcli_privkey
must be set.
```
- If the following message is output to the Cinder log, an incorrect value (such as characters) may have been set to a parameter for which a numeric value must be specified. Specify a correct value.

```
invalid literal for int() with base 10: 'xxx'
```

A.2. Check whether the disk array has been started and iSMCLI is usable.

- If the following message is output to the Cinder log, connection with a storage device by using iSMCLI may have failed.

```
[Errno 113] EHOSTUNREACH
```

A.3. Check whether the LD that cinder-volume is using has been deleted.

- If the following message is output to the Cinder log, startup failed because the presence of the LD used by cinder-volume could not found.
 If the volume is unnecessary, delete it (by using `cinder delete` or other methods), restart cinder-volume after the volume status becomes `deleting`, and then delete the volume again by using the cinder management command (`cinder-manage`).

```
MStorageISCSIDriver: LD `<LD name>` not found.
```

Q. Volumes cannot be created or deleted.

A.1. Check whether the settings for Cinder are correct.

- If the following exception stack is output to the Cinder log, `nec_ismcli_fip` may have an invalid IP address format.

```

-----
Error: ['Traceback (most recent call last):\n',
      File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
      File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
      File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
      File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n
with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
      File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n
self._connect()\n',
      File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 109, in _connect\n
self.ssh.connect(fip, username=user, password=passwd)\n',
      File "/usr/local/lib/python2.7/dist-packages/paramiko/client.py", line 291, in
connect\n
for (family, socktype, proto, canonname, sockaddr) in socket.getaddrinfo(hostname, port,
socket.AF_UNSPEC, socket.SOCK_STREAM):\n',
'gaierror: [Errno -2] Name or service not known\n']
-----

```

- If the following exception stack is output to the Cinder log, the IP address specified for `nec_ismcli_fip` may be incorrect, and therefore, connection may have failed.

```

-----
Error: ['Traceback (most recent call last):\n',
      File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
      File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
      File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
      File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n
with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
      File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n
self._connect()\n',

```

```

File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 109, in _connect\n
self.ssh.connect(fip, username=user, password=passwd)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/client.py", line 305, in
connect\n    retry_on_signal(lambda: sock.connect(addr))\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/util.py", line 278, in
retry_on_signal\n    return function()\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/client.py", line 305, in
<lambda>\n    retry_on_signal(lambda: sock.connect(addr))\n',
File "/usr/local/lib/python2.7/dist-packages/eventlet/greenio.py", line 194, in
connect\n    socket_checkerr(fd)\n',
File "/usr/local/lib/python2.7/dist-packages/eventlet/greenio.py", line 46, in
socket_checkerr\n    raise socket.error(err, errno.errorcode[err])\n',
'error: [Errno 110] ETIMEDOUT\n']
-----

```

- If the following exception stack is output to the Cinder log, the value specified for `nec_ismcli_user` or `nec_ismcli_password` may be incorrect or the corresponding public key specified for `nec_ismcli_privkey` may not have been registered in the disk array, causing an authentication error.

```

-----
Error: ['Traceback (most recent call last):\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n    with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n
self._connect()\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 109, in _connect\n
self.ssh.connect(fip, username=user, password=passwd)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/client.py", line 342, in
connect\n    self._auth(username, password, pkey, key_filenames, allow_agent,
look_for_keys)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/client.py", line 533, in _auth\n
raise saved_exception\n', 'AuthenticationException: Authentication failed.\n']
-----

```

- If the following exception stack is output to the Cinder log, the secret key file path or file name specified for `nec_ismcli_privkey` may be incorrect.

```

-----
Error: ['Traceback (most recent call last):\n', File
"/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n
self._connect()\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 102, in _connect\n
privkey, password=password)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 198, in
from_private_key_file\n key = cls(filename=filename, password=password)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 51, in
__init__\n self._from_private_key_file(filename, password)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 163, in
_from_private_key_file\n data = self._read_private_key_file('\RSA', filename,
password)\n',
File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 279, in
_read_private_key_file\n f = open(filename, 'r')\n',
"IOError: [Errno 2] No such file or directory: '/etc/cinder/id_rsa'\n"]
-----

```

- If the following exception stack is output to the Cinder log, the user may not have the access right for the secret key file specified for `nec_ismcli_privkey`.

```

-----
Error: ['Traceback (most recent call last):\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n

```

```

self._connect()\n',
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 102, in _connect\n
privkey, password=passphrase)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 198, in
from_private_key_file\n    key = cls(filename=filename, password=password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 51, in
__init__\n    self._from_private_key_file(filename, password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 163, in
_from_private_key_file\n    data = self._read_private_key_file('\RSA', filename,
password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 279, in
_read_private_key_file\n    f = open(filename, 'r')\n',
"IOError: [Errno 13] Permission denied: '/etc/cinder/id_rsa'\n"]
-----

```

- If the following exception stack is output to the Cinder log, the secret key file specified for `nec_ismcli_privkey` is encrypted.

```

-----
Error: ['Traceback (most recent call last):\n',
  File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
  File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
  File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 880, in
create_volume\n    with NECiSMCLI(self._fmt, self._local_conf) as cli:\n',
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 144, in __enter__\n
self._connect()\n',
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 102, in _connect\n
privkey, password=passphrase)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 198, in
from_private_key_file\n    key = cls(filename=filename, password=password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 51, in
__init__\n    self._from_private_key_file(filename, password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/rsakey.py", line 163, in
_from_private_key_file\n    data = self._read_private_key_file('\RSA', filename,
password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 280, in
_read_private_key_file\n    data = self._read_private_key(tag, f, password)\n',
  File "/usr/local/lib/python2.7/dist-packages/paramiko/pkey.py", line 323, in

```

```

_read_private_key\n
raise PasswordRequiredException('\Private key file is encrypted')\n',
'PasswordRequiredException: Private key file is encrypted\n']
-----

```

- If the following exception stack is output to the Cinder log, the LD Set could not be found in the disk array.

```

-----
Traceback (most recent call last):
  File "/opt/stack/cinder/cinder/openstack/common/rpc/amqp.py", line 430, in
_process_data
    rval = self.proxy.dispatch(ctxt, version, method, **args)
  File "/opt/stack/cinder/cinder/openstack/common/rpc/dispatcher.py", line 133, in
dispatch
    return getattr(proxyobj, method)(ctxt, **kwargs)
  File "/opt/stack/cinder/cinder/volume/manager.py", line 288, in create_volume
    LOG.error(_("volume %s: create failed"), volume_ref['name'])
  File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
  File "/opt/stack/cinder/cinder/volume/manager.py", line 281, in create_volume
    model_update = self.driver.create_export(context, volume_ref)
  File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner
    retval = f(*args, **kwargs)
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1116, in
create_export
    return self._do_export(_ctx, volume)
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1069, in _do_export
    % (self._conf_ldset, ))
Error: MStorageISCSIDriver: LD Set `XXX` not found.
-----

```

- If the following exception stack is output to the Cinder log, the disk array does not have usable pools. (The following cases are possible:
 - There are no pools.
 - The pools specified for `nec_pools` do not exist.
 - There are no pools with a free capacity larger than the volume to be created.)

```

-----
Error: ['Traceback (most recent call last):\n',
  File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
  File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n

```



```

model_update = self.driver.create_volume(volume_ref)\n',
    File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
    File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 890, in
create_volume\n
raise exception.Error(self._msg('\nNo available pools found.\n'))\n',
'Error: MStorageISCSIDriver: No available pools found.\n']
-----

```

A.2. Check whether the environment has been properly built.

- If the following exception stack is output to the Cinder log, the disk array does not have LD Sets.

```

-----
Traceback (most recent call last):
  File "/opt/stack/cinder/cinder/openstack/common/rpc/amqp.py", line 430, in
_process_data
    rval = self.proxy.dispatch(ctxt, version, method, **kwargs)
  File "/opt/stack/cinder/cinder/openstack/common/rpc/dispatcher.py", line 133, in
dispatch
    return getattr(proxyobj, method)(ctxt, **kwargs)
  File "/opt/stack/cinder/cinder/volume/manager.py", line 288, in create_volume
    LOG.error(_("volume %s: create failed"), volume_ref['name'])
  File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
  File "/opt/stack/cinder/cinder/volume/manager.py", line 281, in create_volume
    model_update = self.driver.create_export(context, volume_ref)
  File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner
    retval = f(*args, **kwargs)
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1116, in
create_export
    return self._do_export(_ctx, volume)
  File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1058, in _do_export
    raise exception.Error(self._msg('No LD Sets found.'))
Error: MStorageISCSIDriver: No LD Sets found.
-----

```

- If the following exception stack is output to the Cinder log, there are no iSCSI portals that the LD Set can use.

```

-----
Traceback (most recent call last):
  File "/opt/stack/cinder/cinder/openstack/common/rpc/amqp.py", line 430, in
_process_data

```

```

    rval = self.proxy.dispatch(ctxt, version, method, **args)
File "/opt/stack/cinder/cinder/openstack/common/rpc/dispatcher.py", line 133, in
dispatch
    return getattr(proxyobj, method)(ctxt, **kwargs)
File "/opt/stack/cinder/cinder/volume/manager.py", line 288, in create_volume
    LOG.error(_("volume %s: create failed"), volume_ref['name'])
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/cinder/cinder/volume/manager.py", line 281, in create_volume
    model_update = self.driver.create_export(context, volume_ref)
File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner
    retval = f(*args, **kwargs)
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1116, in
create_export
    return self._do_export(_ctx, volume)
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 1074, in _do_export
    % (ldset['ldsetname'], )))
Error: MStorageISCSIDriver: LD Set `XXX` has no portals.
-----

```

- If the following exception stack is output to the Cinder log, all the LD numbers are in use. Delete unnecessary logical disks.

```

-----
Error: ['Traceback (most recent call last):\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 250, in create_volume\n
image_location)\n',
File "/opt/stack/cinder/cinder/volume/manager.py", line 168, in _create_volume\n
model_update = self.driver.create_volume(volume_ref)\n',
File "/opt/stack/cinder/cinder/openstack/common/lockutils.py", line 242, in inner\n
retval = f(*args, **kwargs)\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 892, in
create_volume\n    selected_pool, ldn = cli.get_poolnumber_and_ldnumber(pools, used_ldns,
max_ld_count)\n',
File "/usr/local/lib/python2.7/dist-packages/nec/volume.py", line 335, in
get_poolnumber_and_ldnumber\n
raise exception.Error(self._msg('\All LDNs are used.'))\n',
'Error: MStorageISCSIDriver: All LDNs are used.\n']
-----

```

- If the following exception stack is output to the Cinder log, the iSMCLI command failed. Check the explanation in <exit_status> and <stderr> and take appropriate measures.

```
-----  
MStorageISCSIDriver: `<iSMCLI command line>` failed. status=<exit_status>, out="<stdout>",  
err="<stderr>".  
-----
```

- Check whether the disk array has been started and iSMCLI is usable. If the following message is output to the Cinder log, connection with the storage device by using iSMCLI may have failed.

```
-----  
"[Errno 113] EHOSTUNREACH"  
-----
```

Q. Volumes cannot be attached.

A.1. Check whether nova services are running.

A.2. Check whether the environment has been properly built.

- If the following exception stack is output to the nova-compute log, there is a possibility that any of the following has occurred.
 - iSCSI connection may not have been established due to a route problem.
 - The initiator has not been registered in the LD Set.
 - The disk array has not been started.

Traceback (most recent call last):

```

File "/opt/stack/nova/nova/openstack/common/rpc/amqp.py", line 430, in _process_data
    rval = self.proxy.dispatch(ctxt, version, method, **args)
File "/opt/stack/nova/nova/openstack/common/rpc/dispatcher.py", line 133, in dispatch
    return getattr(proxyobj, method)(ctxt, **kwargs)
File "/opt/stack/nova/nova/exception.py", line 117, in wrapped
    temp_level, payload)
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/nova/nova/exception.py", line 94, in wrapped
    return f(self, context, *args, **kw)
File "/opt/stack/nova/nova/compute/manager.py", line 209, in decorated_function
    pass
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/nova/nova/compute/manager.py", line 195, in decorated_function
    return function(self, context, *args, **kwargs)
File "/opt/stack/nova/nova/compute/manager.py", line 237, in decorated_function
    e, sys.exc_info())
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/nova/nova/compute/manager.py", line 224, in decorated_function
    return function(self, context, *args, **kwargs)
File "/opt/stack/nova/nova/compute/manager.py", line 2910, in attach_volume
    context, instance, mountpoint)
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/nova/nova/compute/manager.py", line 2905, in attach_volume
    mountpoint, instance)
File "/opt/stack/nova/nova/compute/manager.py", line 2945, in _attach_volume

```

```

connector)
File "/usr/lib/python2.7/contextlib.py", line 24, in __exit__
    self.gen.next()
File "/opt/stack/nova/nova/compute/manager.py", line 2936, in _attach_volume
    mountpoint)
File "/opt/stack/nova/nova/virt/libvirt/driver.py", line 959, in attach_volume
    disk_info)
File "/opt/stack/nova/nova/virt/libvirt/driver.py", line 945, in volume_driver_method
    return method(connection_info, *args, **kwargs)
File "/opt/stack/nova/nova/openstack/common/lockutils.py", line 242, in inner
    retval = f(*args, **kwargs)
File "/opt/stack/nova/nova/virt/libvirt/volume.py", line 245, in connect_volume
    self._run_iscsiadm(iscsi_properties, ("--rescan",))
File "/opt/stack/nova/nova/virt/libvirt/volume.py", line 179, in _run_iscsiadm
    check_exit_code=check_exit_code)
File "/opt/stack/nova/nova/utils.py", line 239, in execute
    cmd=' '.join(cmd))
ProcessExecutionError: Unexpected error while running command.
Command: sudo nova-rootwrap /etc/nova/rootwrap.conf iscsiadm -m node -T <iqn> -p
<IP>:3260 --rescan
Exit code: 255
Stdout: ''
Stderr: 'iscsiadm: No portal found.\n'
-----

```

- When the retry interval of another component such as HAProxy is set to be shorter than the Cinder timeout (`rpc_response_timeout`), the request being processed may be retried. In this case, the request results in an error.

When a message similar to the following is reported in the Cinder log, it indicates attaching the same volume was repeated.

```

-----
ExitStatus:1
", err="iSM31091:[ addldsetld ]Specified LD is assigned to LD Set.
-----

```

In this case, set the retry interval of HAProxy to be longer than the Cinder timeout.

If an error due to retries by HAProxy occurs, the process status shown in the management information of Nova and Cinder continues to be incomplete and disagrees with the process status, complete, on NEC Storage. For the volume with the process error, you need to resolve the disagreement by directly operating NEC Storage.

For example, when attaching failed, directly log in to NEC Storage and run the `iSMcfg delldsetld` command to deassign the volume with error from the LD set.

Example:

```
iSMcfg delldsetld -ldset LX:OpenStack0 -ldname LX:XXXXX
```

Q. I want to delete the error volume.

- A.1. To delete the volume whose status is error, use the normal deletion procedure (for example, use the `cinder delete` command). To delete the volume whose status is `error_deleting`, use the cinder management command (`cinder-manage volume delete`) or `cinder force-delete` command.**

This page is intentionally left blank.