



Hewlett Packard
Enterprise

Linux Integration Guide

Legal Notices

Copyright © 2021 by Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Hewlett Packard Enterprise believes in being unconditionally inclusive. If terms in this document are recognized as offensive or noninclusive, they are used only for consistency within the product. When the product is updated to remove the terms, this document will be updated.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Publication Date

Wednesday December 15, 2021 10:06:02

Document ID

yle1622674372267

Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: support@nimblestorage.com

For all other general support contact information, go to <https://www.hpe.com/us/en/services/nimble-storage.html>.



Contents

HPE Storage Toolkit for Linux.....	6
Prerequisites.....	6
NORADATAMGR Prerequisites.....	6
Connection Manager Prerequisites.....	7
HPE Storage Toolkit fo Linux Backup Service Prerequisites.....	8
Nimbletune Prerequisites.....	8
Download the Toolkit for Linux Installation Package.....	8
Install HPE Storage Toolkit for Linux	8
Use Java on Host Instead of Packaged Java.....	9
Upgrades Between Major Releases.....	9
Toolkit for Linux Commands.....	10
HPE Storage Oracle Application Data Manager.....	11
The noradatamgr.conf File.....	11
ASM_DISKSTRING Settings.....	12
Accessing NORADATAMGR.....	12
NORADATAMGR Workflows.....	13
Describe an Instance.....	15
Describe a Diskgroup.....	16
Enable NPM for an instance.....	17
Disable NPM for an instance.....	18
Edit an Instance.....	18
Take a Snapshot of an Instance.....	19
List and Filter Snapshots of an Instance.....	21
List Pfile and Control File.....	24
Clone an Instance.....	25
Clone Refresh.....	30
Delete a Snapshot of an Instance.....	34
Destroy a Cloned Instance.....	35
Destroy an ASM Diskgroup.....	36
Overview of Using RMAN with NORADATAMGR	36
Add a NORADATAMGR Backup to an RMAN Recovery Catalog.....	38
Clean Up Cataloged Diskgroups and Volumes Used With RMAN.....	40
RAC Database Recovery.....	42
Oracle NPM and Linux Toolkit Backup Service.....	43
NPM Commands.....	43
Troubleshooting NORADATAMGR.....	44
Log Files.....	44

SQL Plus.....	44
Information for a noradatamgr Issue.....	44
ASM Disks Offline When Controller Failure Occurs.....	45
noradatamgr Command Fails after HPE Storage Toolkit for Linux Installation.....	45
Snapshot and Clone Workflows Fail if spfile not Specified.....	45
Failover causes ASM to mark disks offline or Oracle database fails.....	45
Bringing Up the RAC Nodes After an RMAN Recovery Operation.....	46

HPE Storage Connection Manager for Linux.....47

Configuration.....	47
ncmadm Workflows.....	47
DM Multipath Configuration.....	48
Considerations for Using the HPE Storage Connection Manager for Linux in Scale-out Mode.....	49
LVM Filter for the HPE Storage Connection Manager for Linux in Scale-Out Mode.....	49
iSCSI and FC Configuration.....	49
Block Device Settings.....	50
Discover Devices for iSCSI.....	50
Discover Devices for Fibre Channel.....	52
List Devices.....	52
iSCSI Topology Scenarios.....	52
iSCSI Connection Handling.....	53
SAN Boot Device Support.....	54
Create and Mount Filesystems.....	55
Major Kernel Upgrades and HPE Storage Connection Manager for Linux.....	58
Known Issues and Troubleshooting Tips.....	59
Information for a Linux Issue.....	59
HPE Storage Toolkit for Linux Fails During Reboot When /opt Mounted to Another Device.....	59
Old iSCSI Targets IQNs Still Seen after Volume Rename on the Array.....	60
iSCSI_TCP_CONN_CLOSE Errors.....	60
rescan-scsi-bus.sh Stops Responding or Takes a Long Time While Scanning Standby Paths.....	60
Long Boot Time With a Large Number of FC Volumes.....	60
Duplicate Devices Found Warning.....	61
Newly Added iSCSI Targets Not Discovered.....	61
System Enters Maintenance Mode on Host Reboot with iSCSI Volumes Mounted.....	61
hung_tasks_timeout Errors.....	61
RHEL 7.x: FC System Fails to Boot.....	62
RHEL 7.x: Fibre Channel Active Paths are not Discovered Promptly.....	62
During Controller Failover, multipathd Failure Messages Repeat.....	62
LVM Devices Missing from Output.....	62
Various LVM and iSCSI Commands Hang on /dev/mapper/mpath Devices.....	62
Devices Not Discovered by Multipath.....	63
System Drops into Emergency Mode.....	63
HPE Storage Toolkit for Linux Installer by Default Disables dm-switch Devices on Upgrades.....	63
HPE Storage Connection Manager for Linux with RHEL HA Configuration.....	63

The reservation_key Parameter.....	63
Volume Move Limitations in HA Environments.....	64
iSCSI Target Scope Conversion Utility.....	65
Features.....	65
Prerequisites.....	65
Change an iSCSI Target Scope.....	65
multipath.conf Settings.....	66

Nimbletune.....68

Nimbletune Commands.....	68
Change Nimbletune Recommendations.....	74

Block Reclamation.....75

Reclaiming Space on Linux.....	76
--------------------------------	----

Host Adjustments for Multiple Initiators Using a Single Subnet.....77

Determine the kernel and release version of the host.....	77
ARP adjustments for RHEL 6.x - 8.x, OEL 6.x - 8.x, CentOS 6.x - 8.x.....	78
ARP adjustments for Debian/Ubuntu 18.04 LTS Kernel 3.x - 4.x.....	78
Linux kernel parameter definitions.....	78

Helpful Information.....80

Alignment Requirements for Linux.....	80
---------------------------------------	----

HPE Storage Toolkit for Linux

This document describes the HPE Storage Toolkit for Linux version and its services. The toolkit can install the following services:

- HPE Storage Connection Manager
- HPE Storage Oracle Application Data Manager
- HPE Storage Backup Service
- HPE Host Tuning Utility
- The Config Collector Service

To use toolkit and its services, the following knowledge is assumed:

- Basic Linux system administration skills
- Basic iSCSI SAN knowledge
- Oracle ASM and database technology (to use HPE Storage Oracle Application Data Manager)

Note: The toolkit collects and sends host-based information to the array, once daily per array, where the information is archived for use by InfoSight predictive analytics. No proprietary data or personal information is collected. Only basic information about the host is collected, such as the hostname, host OS, whether it is running as a VM or in machine mode, and the MPIO configuration.

Prerequisites

To use HPE StorageToolkit for Linux and its components, you must have administrator access and a management connection from the host to the array. In addition, the following system requirements must be met:

Hardware requirements:

- 1GB of RAM or more
- 500MB to 1GB free disk space on the `/root` partition or on the `/opt/` partition if mounted separately
- Hostname set other than “localhost” if Application Data Manager is used

Software requirements vary depending upon what Linux toolkit features are installed. Refer to the following topics for more information:

- [Prerequisites and Support](#) on page 7 for HPE Storage Connection Manager
- [NORADATAMGR Prerequisites](#) on page 6 for HPE Storage Oracle Application Data Manager
- [HPE Storage Toolkit fo Linux Backup Service Prerequisites](#) on page 8 for the Linux toolkit Backup Service
- [Nimbletune Prerequisites](#) on page 8 for the Nimbletune utility

Recommendation: By default, some Linux operating systems mount `/opt` to a separate device. If your system does this, HPE Storage recommends that you install the Linux toolkit within the root device. If you don't, the Linux toolkit might not start as expected during a reboot because it cannot locate its dependencies. When `/opt/NimbleStorage` is mounted on a separate partition, it might not mount before the Linux toolkit service starts.

NORADATAMGR Prerequisites

To use NORADATAMGR and its components, the following prerequisites must be met:

Software requirements:

- Array OS 3.5 or higher
- Oracle 11.2.x.x and Oracle 12cR1 (12.1.x.x) or Oracle 12cR2 (12.2.x.x) or Oracle 18 or Oracle 19c.



Note: Regarding Oracle 11.2.0.4: In RHEL/CentOS/Oracle Linux 7.3, 7.4 and 7.5, the Clone a Database workflow fails due to Oracle bug 25638411. This bug has been fixed as part of Patch 19544733. Oracle patch "Patch 19544733: STRICT BEHAVIOR OF LINKER IN 4.8.2 VERSION OF GCC-BINUTILS" has to be applied to the GRID_HOME to be able to run the clone workflow on RHEL/CentOS/Oracle Linux 7.3, 7.4 or 7.5. Refer to the *HPE Storage Linux Toolkit Release Notes* for more information.

Note: The multi-tenant database feature added in 12cR2 consisting of CDB and PDBs is not supported by NORADATAMGR.

- Oracle Automatic Storage Management (ASM) or ASMLib
- RHEL 6.5 and above
- RHEL 7.1 and above
- Oracle Linux 6.5 and above
- CentOS 6.5 and above

Partition requirement:

- Use whole multipath Linux disks for ASMLib disks/Oracle ASM diskgroups.
- Do not create partitions on multipath Linux disks that will make up the ASMLib disks/Oracle ASM diskgroups.

Connection Manager Prerequisites

Prerequisites and Support

The array OS and host OS prerequisites for the HPE Storage Connection Manager for Linux are:

- Array OS 2.3.12 or higher
- RHEL 6.5 and above
- RHEL 7.1 and above
- CentOS 6.5 and above
- CentOS 7.1 and above
- Oracle Linux 6.5 and above
- Oracle Linux 7.1 and above
- Ubuntu LTS 14.04
- Ubuntu LTS 16.04
- Ubuntu LTS 18.04
- Ubuntu LTS 20.04
- SLES 12 and above

The HPE Storage Connection Manager for Linux is supported on the following protocols:

- iSCSI
- Fibre Channel

The HPE Storage Connection Manager for Linux does not support network bonding on iSCSI.

Host Prerequisites

Before you begin, ensure these packages are installed on the host.

- sg3_utils and sg3_utils-libs
- device-mapper-multipath
- libStorageMgmt
- iscsi-initiator-utils (for iSCSI deployments)

Port Requirement

Default REST port 5392 must not be blocked by firewalls or host IP tables.



HPE Storage Toolkit fo Linux Backup Service Prerequisites

The array OS and host OS prerequisites for the Linux Toolkit Backup Service are:

- Array OS 5.0.x or higher
- For replication, the downstream group must be running Array OS 4.x or higher

Nimbletune Prerequisites

Nimbletune is supported on any of the following host operating systems:

- RHEL 6.5 and above
- RHEL 7.1 and above
- CentOS 6.5 and above
- CentOS 7.1 and above
- Oracle Linux 6.5 and above
- Oracle Linux 7.1 and above
- Ubuntu LTS OS 14.04
- Ubuntu LTS OS 16.04
- Ubuntu LTS OS 18.04
- SLES 12 and above

Download the Toolkit for Linux Installation Package

The HPE Storage Toolkit for Linux is available from the Software Download section of HPE InfoSight. You can download the software directly from the HPE Storage Toolkit for Linux Toolkit page, or you can use the Anonymous Download Form. The instructions in this section tell you how to download the software from InfoSight.

Procedure

1. Go to <https://infosight.hpe.com> and log in.
2. Click **Resources** > **Nimble Storage** > **Software Downloads**.
3. In the Integration Kits pane, click **HPE Storage Toolkit for Linux**.
4. From the HPE Storage Toolkit for Linux page, select the current version under Downloads and save it to your Linux host. The download starts automatically.

Note: Each toolkit row also has a link to the Anonymous Download Form. When you select that link, you go to the Anonymous Download page, which contains the Toolkit for Linux End User License Agreement (EULA). After you check the box to accept the EULA, you can click the Download button on that page to get the software.

5. It is a good practice to download the latest *HPE Storage Toolkit for Linux Release Notes*. The Release Notes contain special notes for using the toolkit and information about known issues.

Install HPE Storage Toolkit for Linux

The HPE Storage Toolkit for Linux installer detects your host information and prompts you to install the components that are supported in your environment.

Before you begin

You must have the installation package stored locally on your host.

Note: When installing HPE Storage Toolkit for Linux, scale-out mode is not supported on Ubuntu and SUSE Linux OS.



Procedure

1. Make sure the file is executable and run the installer as root:

```
/path/to/download/HPE_Storage_toolkit_for_linux_x.x.x.YY
```

where x.x.x is the toolkit version number and YY is the toolkit build number.

If your system mounts `/opt` to a separate device, HPE recommends that you install toolkit within the root device. If you don't, `/opt/NimbleStorage` might not mount before the toolkit service starts, which could cause the toolkit service to fail during a reboot. You can use the **--install-dir /nimblestorage -ncm** option of the toolkit installer to change the directory.

Note: For automated installations, you can use the **--silent-mode** option. Use the **--help** option to see more installation options.

2. Accept the EULA.
3. Enter yes to continue with the toolkit installation.
4. If prompted, enter yes to install the HPE Storage Connection Manager.
5. If prompted, enter yes to install the Oracle plugin.

Note: You are only prompted if your environment supports the NORADATAMGR.

6. Enter yes to enable iSCSI services.
7. Verify the status of the plugins you installed.

```
nltadm --status
```

8. Add an array group to the HPE Storage Toolkit for Linux to enable the plugins to provision volumes.

```
nltadm --group --add --ip-address hostname or management IP address --username username--password password
```

Note: HPE recommends using the IP address as the API endpoint to avoid any unrelated connectivity issues with HPE Storage Toolkit for Linux.

9. Verify management communication between the host and the array.

```
nltadm --group --verify --ip-address management IP address
```

Use Java on Host Instead of Packaged Java

The Toolkit for Linux comes packaged with a Java JRE. The Toolkit for Linux clients (ncmadm and nltadm) and daemons use the packaged Java that comes with each Toolkit for Linux release. If vulnerabilities are found in Java, a new JRE with fixes gets updated into the next toolkit release. This means that you might have to wait for the next toolkit release to get an updated Java JRE.

If you prefer not to wait for the toolkit release to get an updated JRE, you can elect to use Java on the host instead of the Java that come packaged with the toolkit.

To use the Java installed on the host instead of using the packaged Java, install toolkit with the **--override-jre** option, as shown in the following example:

```
HPE_Storage_Toolkit_for_linux_4.0.0.0 --silent-mode --accept-eula --ncm --override-jre
```

Note: Make sure you export the proper JAVA_HOME before installing HPE Storage Toolkit for Linux 4.0 and above.

Upgrades Between Major Releases

You can upgrade the HPE Storage Toolkit for Linux by installing the new toolkit software.



If you are upgrading the toolkit from a major version, such as release 3.0.0, to another major version, such as release 4.0.0, you should always uninstall the current toolkit before installing the new toolkit. Otherwise, an outage occurs because the system cannot access the LUNs.

This is not an issue if you are updating from a minor toolkit release to another minor toolkit release.

Toolkit for Linux Commands

Use these commands to manage the Toolkit for Linux services (Connection Manager and NORADATAMGR).

```
nltadm --help
usage:

nltadm [--start <ncm|oracle|collector|backup-service> ]
      nltadm [--stop <ncm|oracle|collector|backup-service> ]
      nltadm [--enable <oracle|collector|backup-service> ]
      nltadm [--disable <oracle|collector|backup-service> ]
      nltadm [--status ]
      nltadm [--config {--dry-run} {--format <xml|json>}]
      nltadm [--version ]
      nltadm [--diag {CASE NUMBER} ]
      nltadm [--group --add --ip-address <GROUP MANAGEMENT IP> --username
<GROUP USERNAME> {--password <GROUP PASSWORD> } ]
      nltadm [--group --remove --ip-address <GROUP MANAGEMENT IP> ]
      nltadm [--group --verify --ip-address <GROUP MANAGEMENT IP> ]
      nltadm [--group --list ]
      nltadm [--help ]

      --add                               Add a Nimble Group
      --config                             Send host configuration data to the
array.
      --diag <CASE NUMBER>               Collect diagnostic dump for NLT compo-▶
nents
      --disable <SERVICE NAME>          Disable a service
      --dry-run                             Dump host configuration information
without sending to array.
      --enable <SERVICE NAME>           Enable a service
      --format <OUTPUT FORMAT>           Output format for host configuration
information. Should be only used with --dry-run option
      --group                             Nimble Group commands
      -h,--help                           Prints this help message
      --ip-address <GROUP MANAGEMENT IP> Nimble Group management IP address
      --list                               List Nimble Groups along with connection
status
      --password <GROUP PASSWORD>        Nimble Group password
      --remove                             Remove a Nimble group
      --start <SERVICE NAME>            Start a service
      --status                             Display the status of each service
      --stop <SERVICE NAME>             Stop a service
      --username <GROUP USERNAME>        Nimble Group username
      -v,--version                         Display the version of Linux Toolkit
      --verify                             Verify connectivity to a Nimble Group
```



HPE Storage Oracle Application Data Manager

The HPE Storage Oracle Application Data Manager enables you to take snapshots for backups and recovery points and to create clones of both datastores and volumes. NORADATAMGR simplifies the backup and cloning processes so that you do not need to understand the mapping between your database and the objects in the array.

NORADATAMGR, which uses a command line interface, is part of the HPE Storage Toolkit for Linux and works with Oracle Automatic Storage Management (ASM).

When you use HPE Storage Toolkit for Linux 2.4.0 or later, NORADATAMGR also supports Oracle Recovery Manager (RMAN). The NORADATAMGR **catalog-backup** command makes HPE Storage snapshot data available to RMAN for restore operations. This command is faster than actually copying the data and allows you use a snapshot of an HPE Storage database instance for restore operations.

You can use role-based access control (RBAC) with NORADATAMGR. RBAC is provided at the host level. You maintain a list of hosts that are allowed to access an instance.

The supported Oracle workflows can be performed on both local and remote hosts.

Note: NORADATAMGR on a single host can connect to only one array group. Multiple hosts with instances having the same SID connecting to the same group are not supported.

To use NORADATAMGR, you must be running either Oracle 11gR2, Oracle 12c, Oracle 18c, or Oracle 19c. In addition, Oracle ASM must be installed on a system running Red Hat Enterprise Linux 6.5 or higher. The HPE Storage Validated Configuration Matrix tool, which is online at <https://infosight.hpe.com/resources/nimble/validated-configuration-matrix>, contains information about requirements for using the Linux Toolkit with Oracle.

Note:

For NORADATAMGR to work with ASMLIB, you must verify or adjust the Disc Scan Ordering in Oracle ASM. To adjust the Disc Scan Ordering, ensure that Oracle ASMLib has the following parameters set in the `/etc/sysconfig/oracleasm` and `/etc/sysconfig/oracleasm-_dev_oracleasm` files (and make sure that these 2 files are always in sync):

- **ORACLEASM_SCANORDER** option is set to **dm-mpath**.
 - **ORACLEASM_SCANEXCLUDE** option is set to **sd**.
-

The noradatamgr.conf File

Below is the default `noradatamgr.conf` file with explanations of the parameters. When setting the user and group, ensure that the user is a valid OS user and that both the group and the user have `sysdba` privileges. If the specified user does not have the correct `sysdba` privileges, NORADATAMGR operations will fail.

Changes to `noradatamgr.conf` require a NLT service restart to take effect.

```
/opt/NimbleStorage/etc/noradatamgr.conf
#Oracle user and group, and db user
orcl.user.name=oracle
orcl.user.group=oinstall
orcl.db.user=sysdba

# Maximum time in seconds that the Nimble Oracle App Data Manager will wait
for a database to get into and out of hotbackup mode.
orcl.hotBackup.start.timeout=14400
orcl.hotBackup.end.timeout=15
```

```
# Snapshot name prefix if a snapshot needs to be taken before cloning an instance
nos.clone.snapshot.prefix=BaseFor

# NPM parameter for enabling hot-backup mode for all instances on the host
during scheduled snapshots
npm.hotbackup.enable=false;

# NPM parameter for enabling hot-backup mode for all instances on the host
during scheduled snapshots
npm.hotbackup.enable=false;
```

ASM_DISKSTRING Settings

HPE Storage recommends the following options for the ASM_DISKSTRING setting when you use HPE Storage Oracle Application Data Manager. This ensures proper discovery of devices by ASM for building ASM disk groups.

On udev systems, ASM_DISKSTRING should be set to `/dev/mapper`.

On asmlib systems, ASM_DISKSTRING should either be blank to default to ORCL: glob style disk paths, or set to `/dev/oracleasm/disks`.

Accessing NORADATAMGR

Use the following instructions to allow non-root users to execute `noradatamgr`:

To allow anyone in the DBA group to execute `noradatamgr`

```
chgrp dba /opt/NimbleStorage/etc
chmod g+rx /opt/NimbleStorage/etc
chgrp dba /opt/NimbleStorage/etc/client
chmod g+rx /opt/NimbleStorage/etc/client
chgrp dba /opt/NimbleStorage/bin
chmod g+x /opt/NimbleStorage/bin
chgrp dba /opt/NimbleStorage/etc/nora*
chmod g+r /opt/NimbleStorage/etc/nora*
chgrp dba /usr/bin/noradatamgr
chmod g+x /usr/bin/noradatamgr
chgrp dba /opt/NimbleStorage/etc/client/log4j.properties
chmod g+x /opt/NimbleStorage/etc/client/log4j.properties
```

To allow the Oracle user to execute `noradatamgr`

```
chown oracle /opt/NimbleStorage/etc
chmod u+rx /opt/NimbleStorage/etc
chown oracle /opt/NimbleStorage/etc/client
chmod u+rx /opt/NimbleStorage/etc/client
chown oracle /opt/NimbleStorage/bin
chmod u+x /opt/NimbleStorage/bin
chown oracle /opt/NimbleStorage/etc/nora*
chmod u+r /opt/NimbleStorage/etc/nora*
chown oracle /usr/bin/noradatamgr
chmod u+x /usr/bin/noradatamgr
chown oracle /opt/NimbleStorage/etc/client/log4j.properties
chmod u+x /opt/NimbleStorage/etc/client/log4j.properties
```



NORADATAMGR Workflows

The HPE Storage Oracle Application Data Manager (NORADATAMGR) uses the following command options in the supported NORADATAMGR workflows:

```
[root ~]# noradatamgr --help
usage:
noradatamgr [--describe [--instance {SID}] [--diskgroup {DISKGROUP NAME} ]

noradatamgr [--edit --instance {SID} --allow-hosts {HOSTNAME1,HOSTNAME2... |
ALL} ]

noradatamgr [--snapshot --instance {SID} --snapname {SNAPSHOT NAME} [--hot-
backup] [--replicate] ]

noradatamgr [--clone --instance {SID} --clone-name {DB NAME} [--snapname
{SNAPSHOT NAME}] [--clone-sid {SID}] [--clone-home
{ORACLE HOME}] [--inherit-pfile] [--override-pfile {ABSOLUTE PATH TO PFILE}]
[--folder {FOLDER NAME}] [--sync-repl]
[--diskgroups-only] ]

noradatamgr [--list-snapshots --instance {SID} [--verbose] [--npm] [--cli] ]

noradatamgr [--delete-snapshot --instance {SID} --snapname {SNAPSHOT NAME} [-
delete-replica] ]

noradatamgr [--get-pfile --instance {SID} --snapname {SNAPSHOT NAME} ]

noradatamgr [--get-cfile --instance {SID} --snapname {SNAPSHOT NAME} ]

noradatamgr [--destroy [--instance {SID}] [--diskgroup {DISKGROUP NAME}] [--
silent-mode] ]

noradatamgr [--enable-npm --instance {SID} ]

noradatamgr [--disable-npm --instance {SID} ]

noradatamgr [--catalog-backup --instance {SID} --snapname {SNAPSHOT NAME} --
diskgroup-prefix {DISKGROUP PREFIX} --rcat {RECOVERY
CATALOG INSTANCE} --rcat-user {USER WITH RECOVERY_CATALOG_OWNER ROLE} [--rcat-
passwd {PASSWORD FOR RCAT USER}] ]

noradatamgr [--uncatalog-backup --instance {SID} --rcat {RECOVERY CATALOG IN▶
STANCE} --rcat-user {USER WITH RECOVERY_CATALOG_OWNER
ROLE} [--rcat-passwd {PASSWORD FOR RCAT USER}] ]

noradatamgr [--clone-refresh --clone-sid {SID} [--refresh-snapname {SNAPSHOT
NAME FROM WHICH TO REFRESH}] [--silent-mode] ]

noradatamgr [--help ]
Nimble Storage Oracle App Data Manager
  --allow-hosts <HOSTNAME1[, HOSTNAME2 ...]> Comma separated list of hostnames
that a database instance can be cloned to.

  -c,--clone                               Clone a database instance.

  --catalog-backup                          Clone diskgroups from a snapshot
and catalog them to the recovery catalog so that they
are available as RMAN backups.

  --cli                                     Filter only NORADATAMGR CLI
```

```

triggered snapshots.

--clone-home <ORACLE HOME>                ORACLE HOME of the cloned instance.

--clone-name <DB NAME>                    Database name of the cloned in▶
stance.

--clone-refresh                            Refresh the cloned instance to
a newer snapshot from the parent instance.

--clone-sid <INSTANCE SID>                ORACLE SID of the cloned instance.

-d,--describe                             Describe the storage for a database
instance or an ASM diskgroup.

--delete-replica                           Delete the snapshot from remote
replication group as well.

--destroy                                  Destroy an ASM diskgroup.

--disable-npm                              Disable scheduled snapshots for
the given instance.

--diskgroup <DIKSGROUP NAME>              ASM diskgroup name.

--diskgroup-prefix <DISKGROUP PREFIX>     The prefix to be used for the
cloned diskgroups.

--diskgroups-only                          Only clone the diskgroups of the
specified instance, do not clone the instance itself.

-e,--edit                                  Edit the storage properties for
a database instance.

--enable-npm                               Configure the volume collection
of the instance to enable scheduled snapshots.

--folder <FOLDER>                         Name of the folder in which cloned
volumes must reside.

-h,--help                                  Display this help and exit.

--hot-backup                               Indicates the database must be
put in hot backup mode before a snapshot backup is
taken.

--inherit-pfile                             Inherit pfile parameters from
source instance.

--instance <ORACLE SID>                   Database Instance SID.

-l,--list-snapshots                        List the snapshot backups of a
database instance.

--no-rollback                              Do not rollback cloned disks and
diskgroups on failure of the clone workflow.

--npm                                       Filter only NPM schedule triggered
snapshots.

--override-pfile <OVERRIDE PFILE>        Absolute path to file containing

```

```

override values for pfile properties of a cloned
instance.

-p,--get-pfile          Print contents of the pfile from
a snapshot backup.

-r,--delete-snapshot   Destroy the snapshot backup for
a database instance.

--rcat <RCAT>          Net service name for the recovery
catalog.

--rcat-passwd <RCAT PASSWD> Password for the recovery catalog
user who has the RECOVERY_CATALOG_OWNER role granted.

--rcat-user <RCAT USER> Recovery catalog user who has
the RECOVERY_CATALOG_OWNER role granted.

--refresh-snapname <SNAPSHOT NAME FROM WHICH TO REFRESH> Name of the
snapshot backup on the parent instance which will be used to refresh the
clone.

--replicate            Replicate the snapshot.

-s,--snapshot          Create an application consistent
or crash consistent snapshot backup of a database
instance.

--silent-mode         Initiate workflow in a non-inter▶
active/silent mode without user intervention.

--snapname <SNAPSHOT NAME> Name of the snapshot backup.

--sync-repl           Enable synchronous replication
configuration on the cloned instance.

--uncatalog-backup    Cleanup RMAN backups and cloned
diskgroups generated by the `catalog-backup` command
for the instance.

--verbose             Verbose option.

-x,--get-cfile        Retrieve the control file from a
snapshot backup.

Collapse

```

Describe an Instance

You can obtain information about the storage footprint of an instance using the **describe --instance** command. Instances must be local, live database instances backed by ASM diskgroups that contain disks backed by volumes. The information you see in these cases includes the following:

- Diskgroup names
- Disk and device names
- HPE Storage volume names and sizes
- Hosts allowed to access the instance to perform operations such as cloning, or listing snapshots

If the instance is down, or not backed by diskgroups, or if the diskgroups do not contain disks backed by HPE Storage volumes, you see messages describing the following conditions:

- Instances that are down or not backed by ASM diskgroups
- Diskgroups that do not contain any HPE Storage disks (disks backed by HPE Storage volumes)

If Oracle database instance disk groups are backed by an HPE Storage array with array OS 5.x or later, the NPM status is displayed at the bottom of the describe output.

From Linux Toolkit V3.0 and higher, the describe instance CLI will also display if synchronous replication has been enabled for the volume collection containing the instance volumes.

If a group does not support synchronous replication (array OS version < 5.1.0.0), the sync replication status is shown as N/A.

Procedure

Describe an instance.

```
noradatamgr --describe --instance instance_name
```

Example

```
[root]# noradatamgr --describe --instance syncb
Diskgroup: ARCHDG
  Disk: /dev/oracleasm/disks/ARCH
        Device: dm-3
        Volume: pragati-test-oel610-b-arch
        Serial number: db342856b129460f6c9ce900c9bc2351
        Size: 120GB
Diskgroup: REDODG
  Disk: /dev/oracleasm/disks/REDO
        Device: dm-2
        Volume: pragati-test-oel610-b-redo
        Serial number: 3e43cf4629c77c5c6c9ce900c9bc2351
        Size: 100GB
Diskgroup: DATADG
  Disk: /dev/oracleasm/disks/DATA
        Device: dm-4
        Volume: pragati-test-oel610-b-data
        Serial number: 398eb7742041f4a06c9ce900c9bc2351
        Size: 110GB

Allowed hosts: syncrepl-rhel75-a.vlab.nimblestorage.com,pragati-test-oel610-b

NPM enabled: Yes
Agent Hostname: pragati-test-oel610-b.vlab.nimblestorage.com:9000

Cataloged diskgroups: None

Replication type: Synchronous
Synchronous replication state: in_sync
```

Describe a Diskgroup

You can obtain information about the storage footprint of any ASM diskgroup on a host using the **describe --diskgroup** command. The information you see in these cases includes the following:

- Disk and device names
- Volume names and sizes for all the disks in the diskgroup
- Diskgroups that do not contain any HPE Storage disks (disks backed by HPE Storage volumes)

If the instance is down, or not backed by diskgroups, or if the diskgroups do not contain disks backed by HPE Storage volumes, you see messages describing the following condition:

- Diskgroups that do not contain any HPE Storage disks (disks backed by HPE Storage volumes)

Procedure

Describe a diskgroup.

```
noradatamgr --describe --diskgroup diskgroup_name
```

Example

```
[root]# noradatamgr --describe --diskgroup DATADG
DATADG
  Disk: /dev/oracleasm/disks/DATA1
        Device: dm-0
        Volume: rh68db-data1
        Serial number: c0505f3bfe7fc7fe6c9ce900f4f8b199
        Size: 100GB
  Disk: /dev/oracleasm/disks/DATA2
        Device: dm-1
        Volume: rh68db-data2
        Serial number: 336a1fae8363f4956c9ce900f4f8b199
        Size: 100GB
```

Enable NPM for an instance

This command configures the volume collection and adds the application specific metadata to the volume collection and the volumes.

Before you begin

You must be on array OS 5.0.x and higher.

Procedure

Enable the NPM for an instance.

```
noradatamgr --enable-npm --instance SID
```

Example

Enable NPM for an instance

```
[root]# noradatamgr --instance xyz123 --enable-npm
Success: Scheduled snapshots have been enabled for instance xyz123.
Please create the required schedule on volume collection rac-scan1 to start
the Oracle scheduled snapshots.

[root]# noradatamgr --instance xyz123 --describe
Diskgroup: REDODG
  Disk: /dev/oracleasm/disks/REDO1
        Device: dm-1
        Volume: rac-scan1-redo1
        Serial number: 4556477f760e919f6c9ce9002c92b39c
        Size: 30GB
Diskgroup: FRADG
  Disk: /dev/oracleasm/disks/FRA1
        Device: dm-5
        Volume: rac-scan1-fra1
        Serial number: a37f89039236cb1f6c9ce9002c92b39c
        Size: 120GB
Diskgroup: DATADG
  Disk: /dev/oracleasm/disks/DATA1
        Device: dm-4
```



```
Volume: rac-scan1-data1
Serial number: e7b4816607f7b0146c9ce9002c92b39c
Size: 100GB
```

```
Allowed hosts: ALL
NPM enabled: Yes
```

Disable NPM for an instance

Before you begin

You must be on array OS 5.0.x and higher.

Procedure

Disable NPM for an instance.

```
noradatamgr --disable-npm --instance SID
```

Example

Disable NPM for an instance

```
[root]# noradatamgr --instance xyz123 --disable-npm
Success: Scheduled snapshots have been disabled for instance xyz123.
Please delete the schedules on volume collection rac-scan1.

[root]# noradatamgr --instance xyz123 --describe
Diskgroup: REDODG
  Disk: /dev/oracleasm/disks/REDO1
    Device: dm-1
    Volume: rac-scan1-redo1
    Serial number: 4556477f760e919f6c9ce9002c92b39c
    Size: 30GB
Diskgroup: FRADG
  Disk: /dev/oracleasm/disks/FRA1
    Device: dm-5
    Volume: rac-scan1-fra1
    Serial number: a37f89039236cb1f6c9ce9002c92b39c
    Size: 120GB
Diskgroup: DATADG
  Disk: /dev/oracleasm/disks/DATA1
    Device: dm-4
    Volume: rac-scan1-data1
    Serial number: e7b4816607f7b0146c9ce9002c92b39c
    Size: 100GB

Allowed hosts: ALL
NPM enabled: No
```

What to do next

Note: If the HPE Storage Toolkit for Linux is uninstalled or group information is removed from toolkit, then NPM will automatically be disabled for all database instances.

Edit an Instance

You can edit a local live instance to modify the list of hosts that are allowed to perform cloning and snapshot listing operations on that instance.



If some or all of the data, and online redo volumes of the instance are not in a volume collection when you run the **edit** command, a volume collection is created automatically. Data or online log volumes cannot be part of two different volume collections.

The **--allow-hosts** command is used to place access control on a local live instance at a host level, by specifying the host names which are allowed to access the instance. The following guidelines apply:

- You can use a comma-separated list of host names, or "all" to allow all hosts to access the instance. When listing the host names, you must list the name as it is specified when running the **hostname** command on the host to be added.
- The name of the host on which the edit workflow is run, will always be added to the list of allowed hosts.
- No host name validation is performed. You must enter a valid host name to ensure that operations (such as listing snaps and clones) are successful.
- If you change a host name for any reason, you must run the edit workflow again using the updated list of host names.
- If you want to add a new host to the list, you must run the edit workflow again listing all the host names to be added, including the new host name.

Procedure

Edit an instance.

```
noradatamgr --edit --instance instance_name --allow-hosts hosts_list
```

Example

```
[root]# noradatamgr --edit --instance rh68db --allow-hosts host-rhell,host-rhel2
Allowed hosts set to: host-rhell,host-rhel2
Success: Storage properties for database instance rh68db updated.
[root@host-rhell ~]# noradatamgr --describe --instance rh68db
Diskgroup: ARCHDG
    Disk: /dev/oracleasm/disks/ARCH1
        Device: dm-2
        Volume: rh68db-arch1
        Serial number: 4dc46ccf29cb4a216c9ce900f4f8b199
        Size: 100GB
Diskgroup: REDODG
    Disk: /dev/oracleasm/disks/REDO1
        Device: dm-3
        Volume: rh68db-redo1
        Serial number: 62271db7e2578d7e6c9ce900f4f8b199
        Size: 50GB
Diskgroup: DATADG
    Disk: /dev/oracleasm/disks/DATA1
        Device: dm-0
        Volume: rh68db-data1
        Serial number: c0505f3bfe7fc7fe6c9ce900f4f8b199
        Size: 100GB
    Disk: /dev/oracleasm/disks/DATA2
        Device: dm-1
        Volume: rh68db-data2
        Serial number: 336a1fae8363f4956c9ce900f4f8b199
        Size: 100GB

Allowed hosts: host-rhell,host-rhel2
```

Take a Snapshot of an Instance

You can use the NORADATAMGR to take a snapshot of a local live instance from both a standalone an Oracle Real Application Clusters (RAC) setup. Archive log mode must be enabled on the host before you take a snapshot of a RAC.

Note: Before specifying that the snapshot be replicated, ensure that replication has been set on the array on the volume collection that contains the data and online redo log volumes of the instance.

For taking a snapshot of a local live instance, the **snapshot** command places the data and online log volumes in a single volume collection (which is created if one does not already exist). You can optionally specify that the instance be placed in hot backup mode before taking the snapshot, or that the snapshot be replicated to a downstream group. To place the instance in hot backup mode, you must enable archive log mode on the host.

A snapshot of an Oracle instance taken by NORADATAMGR is a snapshot collection of the volume collection that contains the data and log volumes backing the instance. This snapshot can be used at a later time to create a clone of the instance. All snapshots taken by NORADATAMGR are application consistent, and can provide instance-level recovery on an instance cloned from the snapshot. Point-in-time recovery is not supported.

A snapshot of any instance in an Oracle RAC cluster is a snapshot collection of the entire RAC database and contains all the Oracle metadata necessary to recover the database, including the data and redo logs. The snapshot can also be used to clone the RAC database.

The NORADATAMGR tool mandates that a snapshot of a RAC instance is always taken with hot-backup mode enabled. The tool will try to get a RAC instance in hot-backup mode even if that option is not specified while taking a snapshot.

Regarding synchronous replication, also note the following:

- If synchronous replication has been setup on the volume collection, using the **--replicate** flag will create synchronously replicated snapshot collection.
- The snapshot name will be created with the user specified name on the upstream pool and appended with ".R" on the downstream pool.
- If the **--replicate** option is not used, a snapshot is created on the current upstream pool for synchronous replication enabled volume collection.

Before you begin

Procedure

Take a snapshot of an instance.

```
noradatamgr --snapshot --instance instance_name --snapname snapshot_name [--hot-backup] [--replicate]
```

Example

Take a Snapshot

```
[root]# noradatamgr --snapshot --instance rh68db --snapname rh68.snap.603
Success: Snapshot backup rh68.snap.603 completed.
```

Take a Snapshot and Replicate

```
[root]# noradatamgr --snapshot --instance rh68db --snapname rh68db.snap.repl
--replicate
Success: Snapshot backup rh68db.snap.repl completed.
[root]# noradatamgr --list-snapshot --instance rh68db --verbose
Snapshot Name: rh68db.snap.repl taken at 16-10-24 14:47:14
Instance: rh68db
Snapshot can be used for cloning rh68db: Yes
Hot backup mode enabled: No
Replication status: complete
Database version: 11.2.0.4.0
Host OS Distribution: redhat
Host OS Version: 6.8
```

Take a Snapshot in Hot Backup Mode

```
SQL> archive log list;
Database log mode           Archive Mode
Automatic archival         Enabled
Archive destination        +ARCHDGG
```

```

Oldest online log sequence      103008
Next log sequence to archive    103010
Current log sequence            103010
SQL>

[root]# noradatamgr --snapshot --instance rh68db --snapname rh68db.snap.hotbackup
--hot-backup-mode
Putting instance rh68db in hot backup mode...
Success: Snapshot backup rh68db.snap.hotbackup completed.
Taking instance rh68db out of hot backup mode...
[root@hiqa-sys-rhell ~]#

[root]# noradatamgr --list-snapshot --instance rh68db --verbose
Snapshot Name: rh68db.snap.hotbackup taken at 16-10-24 14:55:37
Instance: rh68db
Snapshot can be used for cloning rh68db: Yes
Hot backup mode enabled: Yes
Replication status: N/A
Database version: 11.2.0.4.0
Host OS Distribution: redhat
Host OS Version: 6.8

```

Create synchronously replicated snapshots

```

[root]# noradatamgr --snapshot --snapname syncb.snap --instance syncb --replicate
Success: Snapshot backup syncb.snap completed.

[root]# noradatamgr --list-snapshots --instance syncb | grep syncb.snap
-----+-----+-----
Snap Name          Taken at          Pool Name          Cloning   Created
                    possible using
-----+-----+-----
syncb.snap         19-03-07 12:23   default            Yes        CLI
syncb.snap.R       19-03-07 12:23   default-group-array5  Yes        CLI

```

List and Filter Snapshots of an Instance

You can list snapshots of an instance that were taken using the NORADATAMGR **snapshot** command or through the toolkit Backup Service schedule. The list can contain snapshots taken from both local and remote instances (remote instances are those which are not running on the host from which the **list-snapshots** command is invoked). Snapshots taken of instances that are shut down are also listed.

You can filter snapshots by those taken only through NPM or only with the CLI.

Note: The **list-snapshots** command will fail if the host does not appear in the list of allowed hosts, or if the instance does not have the correct storage configuration (data and log volumes of the instance are in one volume collection) and the required metadata is not tagged to the volume collection.

Also note the following:

- For Linux Toolkit V3.0 and higher, the pool name is also displayed while listing the snapshots for an instance. This is particularly useful for distinguishing between upstream snapshots and replicas in the case of synchronously replicated snapshot collections.
- If verbose mode is used, then the pool name will show up only if the snapshot collection was created on a synchronously replicated setup; otherwise there will not be any change other than manually triggered or scheduled based periodic snapshots.
- For tabular views, the column - pool name will be added if the array group is running OS V5.1.0.0 or above. Snapshot collections created on volume collections that do not have synchronous replication set up will also list the pool name on which they are created.

Procedure

List snapshots of an instance.

noradatamgr --list-snapshots --instance *instance_name* [--npm, --cli] --verbose

Example

List Snapshots

```
[root]# noradatamgr --list-snapshots --instance syncb
-----+-----+-----+-----+-----+
Snap Name      Taken at  Pool Name      Cloning  Created
                possible using
-----+-----+-----+-----+-----+
prag-test-...  19-03-06  default-group-array5  Yes      NPM
prag-test-...  19-03-06  default            Yes      NPM
prag-test-...  19-03-07  default-group-array5  Yes      NPM
prag-test ...  19-03-07  default            Yes      NPM
syncb.snap1.R  19-03-07  default-group-array5  Yes      CLI
syncb.snap1    19-03-07  default            Yes      CLI
```

With **--verbose** option

```
[root]# noradatamgr --list-snapshots --instance syncb --verbose
Snapshot Name: pragati-test-oel610-Schedule-0-2019-03-06::20:00:00.000.R taken
at 19-03-06 20:00:45
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
Replication status: N/A
Pool name: default-group-sjc-array586
Created using: NPM Schedule
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10
Snapshot Name: pragati-test-oel610-Schedule-0-2019-03-06::20:00:00.000 taken
at 19-03-06 20:00:45
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
Replication status: N/A
Pool name: default
Created using: NPM Schedule
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10
Snapshot Name: pragati-test-oel610-Schedule-0-2019-03-07::08:00:00.000.R taken
at 19-03-07 08:00:35
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
Replication status: N/A
Pool name: default-group-sjc-array586
Created using: NPM Schedule
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10
Snapshot Name: pragati-test-oel610-Schedule-0-2019-03-07::08:00:00.000 taken
at 19-03-07 08:00:35
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
```

```

Replication status: N/A
Pool name: default
Created using: NPM Schedule
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10
Snapshot Name: syncb.snap1.R taken at 19-03-07 13:10:56
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
Replication status: N/A
Pool name: default-group-sjc-array586
Created using: NORADATAMGR CLI
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10
Snapshot Name: syncb.snap1 taken at 19-03-07 13:10:56
Instance: syncb
Snapshot can be used for cloning syncb: Yes
Hot backup mode enabled: No
Replication status: N/A
Pool name: default
Created using: NORADATAMGR CLI
Database version: 12.2.0.1.0
Host OS Distribution: oracle
Host OS Version: 6.10

```

List Snapshots (from Remote Host)

```

(Note: instance rh68db is running on host-rhell)
[root@host-rhel2]# noradatamgr --list-snapshots --instance rh68db
-----+-----+-----+-----
--
Snap Name                Taken at                Instance                Usable for cloning
-----+-----+-----+-----
--
rh68db.hot.repl.102      16-10-24 17:36         rh68db                  Yes
rh68.snap.603           16-10-24 12:01         rh68db                  Yes
rh68db.snap.601         16-10-21 10:36         rh68db                  Yes
rh68db.snap.600         16-10-20 14:48         rh68db                  Yes
rh68db.snap.repl        16-10-24 14:47         rh68db                  Yes
rh68db.snap.hotbackup   16-10-24 14:55         rh68db                  Yes
rh68db.hot.repl.101     16-10-24 17:13         rh68db                  Yes

```

List Snapshots With Filters

```

[root]# noradatamgr --instance xyz123 --list-snapshot --npm
-----+-----+-----+-----+-----
--
Snap Name                Taken at                Instance                Cloning    Created
-----+-----+-----+-----+-----
                                possible using
--
rac-scan1-30-min-2017-10-10::14... 17-10-10 14:00         rac1db2                Yes        NPM
rac-scan1-30-min-2017-10-10::13... 17-10-10 13:00         rac1db2                Yes        NPM
rac-scan1-30-min-2017-10-10::12... 17-10-10 12:30         rac1db2                Yes        NPM

[root]# noradatamgr --instance xyz123 --list-snapshot --cli
-----+-----+-----+-----+-----
--
Snap Name                Taken at                Instance                Cloning    Created

```

					possible using
rac1db2.cli.snap.2	17-10-09	15:04	rac1db2	Yes	CLI
rac1db2.cli.snap.3	17-10-09	15:04	rac1db2	Yes	CLI
rac1db2.cli.snap.4	17-10-09	15:05	rac1db2	Yes	CLI
rac1db2.cli.snap.1	17-10-09	15:03	rac1db2	Yes	CLI
rac1db2.cli.snap.21	17-10-09	18:39	rac1db2	Yes	CLI

List Pfile and Control File

Use these commands to check the pfile and control file of a database instance as captured in a snapshot taken by NORADATAMGR. If a database instance is cloned from this snapshot, the clone inherits the parameters from these files. You can use the **--override-pfile** option while cloning to specify the pfile parameter values.

Procedure

List a pfile or a control file.

```
noradatamgr --get-pfile --instance instance_name --snapname snapshot_name
```

```
noradatamgr --get-cfile --instance instance_name --snapname snapshot_name
```

Example

Retrieve a pfile from a snapshot

```
[root]# noradatamgr --get-pfile --instance rh68db --snapname rh68db.snap.repl
Pfile contents:
rh68db.__db_cache_size=35567697920
rh68db.__java_pool_size=805306368
rh68db.__large_pool_size=268435456
rh68db.__oracle_base='/u01/app/base'#ORACLE_BASE set from environment
rh68db.__pga_aggregate_target=13555990528
rh68db.__sga_target=40667971584
rh68db.__shared_io_pool_size=0
rh68db.__shared_pool_size=3489660928
rh68db.__streams_pool_size=268435456
*.audit_file_dest='/u01/app/base/admin/rh68db/adump'
*.audit_trail='db'
*.compatible='11.2.0.4.0'
*.control_files='+REDODG/rh68db/controlfile/current.256.921347795'
*.db_block_size=8192
*.db_create_file_dest='+DATADG'
*.db_create_online_log_dest_1='+REDODG'
*.db_domain=''
*.db_name='rh68db'
*.diagnostic_dest='/u01/app/base'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=rh68dbXDB) '
*.log_archive_dest_1='LOCATION=+ARCHDG'
*.log_archive_format='%t_%s_%r.dbf'
*.open_cursors=300
*.pga_aggregate_target=13535019008
*.processes=150
*.remote_login_passwordfile='EXCLUSIVE'
*.sga_target=40607154176
*.undo_tablespace='UNDOTBS1'
```

Retrieve a cfile from a snapshot

```
[root]# noradatamgr --get-cfile --instance fcra1db1 --snapname snap.1
Control file contents:
```



```

STARTUP NOMOUNT
CREATE CONTROLFILE REUSE DATABASE "FCRAC1DB" NORESETLOGS  ARCHIVELOG
  MAXLOGFILES 192
  MAXLOGMEMBERS 3
  MAXDATAFILES 1024
  MAXINSTANCES 32
  MAXLOGHISTORY 4672
LOGFILE
  GROUP 19 '+REDODG/fcrac1db/onlineelog/group_19.266.957110493'  SIZE 120M
BLOCKSIZE 512,
  GROUP 20 '+REDODG/fcrac1db/onlineelog/group_20.264.957110499'  SIZE 120M
BLOCKSIZE 512,
  GROUP 21 '+REDODG/fcrac1db/onlineelog/group_21.268.957110505'  SIZE 120M
BLOCKSIZE 512,
  GROUP 22 '+REDODG/fcrac1db/onlineelog/group_22.267.957110509'  SIZE 120M
BLOCKSIZE 512,
  GROUP 23 '+REDODG/fcrac1db/onlineelog/group_23.265.957110513'  SIZE 120M
BLOCKSIZE 512,
  GROUP 24 '+REDODG/fcrac1db/onlineelog/group_24.263.957110547'  SIZE 120M
BLOCKSIZE 512
DATAFILE
  '+DATADG/fcrac1db/datafile/system.262.955105763',
  '+DATADG/fcrac1db/datafile/sysaux.260.955105763',
  '+DATADG/fcrac1db/datafile/undotbs1.261.955105763',
  '+DATADG/fcrac1db/datafile/users.259.955105763',
  '+DATADG/fcrac1db/datafile/undotbs2.257.955105853',
  '+DATADG/fcrac1db/datafile/soe.263.955389549',
  '+DATADG/fcrac1db/datafile/undotbs3.264.956339675'
CHARACTER SET WE8MSWIN1252

```

Clone an Instance

The clone workflow can be used for cloning a local live database instance or cloning from a snapshot of the database instance that has already been taken.

An instance cloned from a snapshot of a RAC instance will always be a stand-alone instance. The DBA should manually create a cluster from this cloned instance if required.

When cloning a local live instance, a new snapshot of the instance is taken. This snapshot has a default prefix of BaseFor in its name. This prefix is configurable in `/opt/NimbleStorage/etc/noradatamgr.conf`. Changes to `noradatamgr.conf` require a toolkit service restart to take effect.

To provide a snapshot name that has already been taken, the `list-snapshots` workflow can be invoked. The required snapshot can then be used to clone the instance. The instance to be cloned can be local or remote. If the source instance to be cloned is based on a remote host, its volumes must reside on the same group that the host from which the CLI is invoked is configured to use. Also, this host (from which the CLI is invoked) must also appear in the list of allowed hosts of the remote instance, for it to be able to clone the instance.

You can also clone only the diskgroups of an instance (without creating a cloned instance containing the cloned diskgroups). This is useful when you want to restore just a single file from a particular snapshot.

Regarding synchronous replication, note the following:

- A new option, **--sync-repl** has been added for the HPE Storage Toolkit V3.0 and higher for a cloned instance to enable synchronous replication on the volume collection of the cloned instance. This option is only supported on array OS version V 5.1 and higher.
- The replication partner is not automatically set and must be set later by the storage administrator.
- The clone CLI output will ask for the replication partner to be set for the volume collection.

Procedure

Clone an instance.

noradatamgr --clone --instance *instance_name* **--clone-name** *clone_name* [**--snapname** *snapshot_name*] [**--clone-sid** *clone_sid_name*] [**--clone-home** *ORACLE_HOME*] [**--inherit-pfile**][**--override-pfile** *absolute_path_to_pfile*] [**--folder** *folder_name*] [**--diskgroups-only**] [**--sync-repl**]

- **--clone-home** option: The clone workflow will use this as the ORACLE_HOME for the cloned instance. If you do not use the **--clone-home** command and a local instance is being cloned, the cloned instance's ORACLE_HOME value is set to that of the source instance. If you do not use the **--clone-home** command and a remote instance is being cloned, the cloned instance's ORACLE_HOME value is set to the value of the environment variable \$ORACLE_HOME.
- **--inherit-pfile** option: This option allows the cloned instance to inherit all the pfile parameters (except the basic ones that are needed to bring the cloned instance up) of the source instance captured in the snapshot. Do not use this option if there is a parameter in the source instance pfile that is incompatible and prevents the cloned instance from starting.
- **--override-pfile** option: This allows you to override the pfile for the cloned instance. If you do not use the **--override-pfile** command, a pfile is generated for the cloned instance. If you use the **--override-pfile** command but specify an incorrect pfile value, cloning will fail.

Example

Clone a database instance from local host

```
(Instance rh68db is running on host-rhell)
[root@host-rhell ~]# noradatamgr --clone --instance rh68db --clone-name
clonedDB --snapname rh68.snap.603 --clone-home /u01/app/base/product/11.2.0/db▶
home_1/
Initiating clone ...
[#####] 100%
Success: Cloned instance rh68db to clonedDB.
Diskgroup: CLONEDDBDATADG
    Disk: /dev/oracleasm/disks/CLONEDDB0002
        Device: dm-8
        Volume: clonedDB-DATA2
        Serial number: ee11f455cc0e40376c9ce900f4f8b199
        Size: 100GB
    Disk: /dev/oracleasm/disks/CLONEDDB0001
        Device: dm-4
        Volume: clonedDB-DATA1
        Serial number: 51e1732dbf0abf326c9ce900f4f8b199
        Size: 100GB
Diskgroup: CLONEDDBLOGDG
    Disk: /dev/oracleasm/disks/CLONEDDB
        Device: dm-10
        Volume: clonedDB-LOG1
        Serial number: 8b999dffa7607ffb6c9ce900f4f8b199
        Size: 50GB

Allowed hosts: host-rhell

[root@host-rhell ~]#
[root@host-rhell ~]# ps aux|grep pmon
oracle  16983  0.0  0.0 39960408 28976 ?        Ss   13:19   0:00
ora_pmon_clonedDB
root    18993  0.0  0.0 103324   868 pts/0    S+   13:21   0:00 grep pmon
grid   28472  0.0  0.0 1340504 26864 ?        Ss   Oct20   0:31 asm_pmon_+ASM
oracle 29160  0.0  0.0 39960684 20904 ?        Ss   Oct20   0:34 ora_pmon_rh68db
[root@host-rhell ~]#
[grid@host-rhell ~]$ asmcmd lsdg
State   Type   Rebal  Sector  Block       AU   Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED  EXTERN  N         512    4096   1048576   102400   14554
0         14554          0         N  ARCHDG/
```

MOUNTED	NORMAL	N	512	4096	1048576	204800	89780
0		44890		0		N	CLONEDDBDATADG/
MOUNTED	EXTERN	N	512	4096	1048576	51200	50835
0		50835		0		N	CLONEDDBLOGDG/
MOUNTED	NORMAL	N	512	4096	1048576	204800	89782
0		44891		0		N	DATADG/
MOUNTED	EXTERN	N	512	4096	1048576	10236	10177
0		10177		0		N	OCRDG/
MOUNTED	EXTERN	N	512	4096	1048576	51200	50956
0		50956		0		N	REDODG/

Clone a database instance from a remote host

(Instance rh68db is running on host-rhell1 and we clone this instance on host-rhel2. Please ensure rh68db on host-rhell1 allows host-rhel2 to list snapshots or clone.)

```
[root@host-rhel2 ~]# noradatamgr --clone --instance rh68db --clone-name db603
--snapname rh68.snap.603 --clone-home /u01/app/base/product/11.2.0/dbhome_1
Initiating clone ...
```

```
[#####] 100%
```

```
Success: Cloned instance rh68db to db603.
```

```
Diskgroup: DB603DATADG
```

```
  Disk: /dev/mapper/mpathy
```

```
    Device: dm-26
```

```
    Volume: db603-DATA2
```

```
    Serial number: 02fcdccbfed7ee536c9ce900f4f8b199
```

```
    Size: 100GB
```

```
  Disk: /dev/mapper/mpathz
```

```
    Device: dm-27
```

```
    Volume: db603-DATA1
```

```
    Serial number: 7542d8330bd19cae6c9ce900f4f8b199
```

```
    Size: 100GB
```

```
Diskgroup: DB603LOGDG
```

```
  Disk: /dev/mapper/mpathaa
```

```
    Device: dm-28
```

```
    Volume: db603-LOG1
```

```
    Serial number: 7fe12c89e12a72926c9ce900f4f8b199
```

```
    Size: 50GB
```

```
Allowed hosts: host-rhel2
```

```
[root@host-rhel2 ~]# ps aux|grep pmon
```

```
root      5035  0.0  0.0 112644  956 pts/0    S+   18:09   0:00 grep --col▶
or=auto pmon
grid      9908  0.0  0.0 1343280 24088 ?      Ss   Oct20   0:35 asm_pmon_+ASM
oracle    11735  0.0  0.0 4442960 16880 ?      Ss   Oct20   0:37 ora_pmon_rh72db
oracle    26029  0.0  0.0 39963196 24896 ?      Ss   18:08   0:00 ora_pmon_db603
```

Clone diskgroups only

```
ASM diskgroups before cloning:
```

```
[grid@host-rhell1 ~]$ asmcmd lsdg
```

State	Type	Rebal	Sector	Block	AU	Total_MB	Free_MB
Req_mir_free_MB	Usable_file_MB	Offline_disks	Voting_files	Name			
MOUNTED	EXTERN	N	512	4096	1048576	102400	14554
0		14554		0		N	ARCHDG/
MOUNTED	NORMAL	N	512	4096	1048576	204800	89782
0		44891		0		N	DATADG/
MOUNTED	EXTERN	N	512	4096	1048576	10236	10177
0		10177		0		N	OCRDG/
MOUNTED	EXTERN	N	512	4096	1048576	51200	50956

```

0          50956          0          N REDODG/
[grid@host-rhell ~]$

Only clone the diskgroups of the specified instance
[root@host-rhell ~]# noradatamgr --clone --instance rh68db --clone-name destDB
--snapname rh68.snap.603 --diskgroups-only
Initiating clone ...
[#####] 100%
Success: Cloning diskgroups of instance rh68db completed.
DESTDBLOGDG
    Disk: /dev/oracleasm/disks/DESTDB
           Device: dm-11
           Volume: destDB-LOG1
           Serial number: 90e8b1d1b0e605b76c9ce900f4f8b199
           Size: 50GB
DESTDBDATADG
    Disk: /dev/oracleasm/disks/DESTDB0002
           Device: dm-12
           Volume: destDB-DATA1
           Serial number: 0dac7bca494ad93d6c9ce900f4f8b199
           Size: 100GB
    Disk: /dev/oracleasm/disks/DESTDB0001
           Device: dm-13
           Volume: destDB-DATA2
           Serial number: 9add48220356e99c6c9ce900f4f8b199
           Size: 100GB
[root@host-rhell ~]#

ASM diskgroups after cloning:
[grid@host-rhell ~]$ asmcmd lsdg
State Type Rebal Sector Block AU Total_MB Free_MB
Req_mir_free_MB Usable_file_MB Offline_disks Voting_files Name
MOUNTED EXTERN N 512 4096 1048576 102400 14554
0 14554 0 N ARCHDG/
MOUNTED NORMAL N 512 4096 1048576 204800 89782
0 44891 0 N DATADG/
MOUNTED NORMAL N 512 4096 1048576 204800 89782
0 44891 0 N DESTDBDATADG/
MOUNTED EXTERN N 512 4096 1048576 51200 50956
0 50956 0 N DESTDBLOGDG/
MOUNTED EXTERN N 512 4096 1048576 10236 10177
0 10177 0 N OCRDG/
MOUNTED EXTERN N 512 4096 1048576 51200 50956
0 50956 0 N REDODG/

```

The processes value was set to 150 for instance rh68db when snapshot was taken.

```

[root@host-rhell 14]# noradatamgr --get-pfile --instance rh68db --snapname
rh68db.snap.110
Pfile contents:
rh68db.__db_cache_size=35567697920
rh68db.__java_pool_size=805306368
rh68db.__large_pool_size=268435456
rh68db.__oracle_base='/u01/app/base'#ORACLE_BASE set from environment
rh68db.__pga_aggregate_target=13555990528
rh68db.__sga_target=40667971584
rh68db.__shared_io_pool_size=0
rh68db.__shared_pool_size=3489660928
rh68db.__streams_pool_size=268435456
*.audit_file_dest='/u01/app/base/admin/rh68db/adump'
*.audit_trail='db'
*.compatible='11.2.0.4.0'

```

```
*.control_files='+REDODG/rh68db/controlfile/current.256.921347795'
*.db_block_size=8192
*.db_create_file_dest='+DATADG'
*.db_create_online_log_dest_1='+REDODG'
*.db_domain=''
*.db_name='rh68db'
*.diagnostic_dest='/u01/app/base'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=rh68dbXDB)'
*.log_archive_dest_1='LOCATION='+ARCHDGDG'
*.log_archive_format='%t_%s_%r.dbf'
*.open_cursors=300
*.pga_aggregate_target=13535019008
  *.processes=150
*.remote_login_passwordfile='EXCLUSIVE'
*.sga_target=40607154176
*.undo_tablespace='UNDOTBS1'
```

Create a text file on the host and set the processes value to 300

```
[root@host-rhell 14]# cat /tmp/my-pfile-new-value.txt
*.processes=300
[root@host-rhell 14]#
```

Clone a new instance from rh68db snapshot and override pfile

```
[root@host-rhell 14]# noradatamgr --clone --instance rh68db --clone-name rh68c110
--snapname rh68db.snap.110 --override-pfile /tmp/my-pfile-new-value.txt
Initiating clone ...
[#####] 100%
Success: Cloned instance rh68db to rh68c110.
Diskgroup: RH68C110LOGDG
  Disk: /dev/oracleasm/disks/RH68C110
  Device: dm-35
  Volume: rh68c110-LOG1
  Serial number: 109895d6f7fd35486c9ce900f4f8b199
  Size: 50GB
Diskgroup: RH68C110DATADG
  Disk: /dev/oracleasm/disks/RH68C1100002
  Device: dm-34
  Volume: rh68c110-DATA2
  Serial number: 0b91b7feee77e1856c9ce900f4f8b199
  Size: 100GB
  Disk: /dev/oracleasm/disks/RH68C1100001
  Device: dm-33
  Volume: rh68c110-DATA1
  Serial number: 6e9f5eef661518d36c9ce900f4f8b199
  Size: 100GB

Allowed hosts: host-rhell
```

Verify processes value in cloned instance parameter file

```
[root@host-rhell 14]# su - oracle
[oracle@hiqa-sys-rhell ~]$ . oraenv
ORACLE_SID = [rh68db] ? rh68c110
The Oracle base has been set to /u01/app/base
[oracle@hiqa-sys-rhell ~]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Tue Nov 15 14:41:16 2016

Copyright (c) 1982, 2013, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Automatic Storage Management, OLAP, Data Mining
and Real Application Testing options
```

```
SQL> show parameter processes
```

NAME	TYPE	VALUE
aq_tm_processes	integer	1
db_writer_processes	integer	3
gcs_server_processes	integer	0
global_txn_processes	integer	1
job_queue_processes	integer	1000
log_archive_max_processes	integer	4
processes	integer	300

```
SQL>
```

Clone CLI with --sync-repl option

```
[root]# noradatamgr -clone --instance syncb --clone-name clone --sync-repl
Initiating clone ...
Initiating snapshot backup BaseForclone-ccb05223-7a0f-41ed-a3d4-fd51289ba345
for instance syncb...
Success: Snapshot backup BaseForclone-ccb05223-7a0f-41ed-a3d4-fd51289ba345
completed.
[#####] 100%
Success: Cloned instance syncb to clone.
Diskgroup: CLONELOGDG
    Disk: /dev/oracleasm/disks/CLONE
        Device: dm-7
        Volume: clone-LOG1
        Serial number: 4fae2c707ee1cff06c9ce900c9bc2351
        Size: 100GB
Diskgroup: CLONEDATADG
    Disk: /dev/oracleasm/disks/CLONE0001
        Device: dm-6
        Volume: clone-DATA1
        Serial number: 74e7d0288066442b6c9ce900c9bc2351
        Size: 110GB

Allowed hosts: pragati-test-oel610-b

NPM enabled: No

Cataloged diskgroups: None

Replication type: Synchronous
Synchronous replication state: not_applicable

Volume collection has been created with synchronous replication. Please set
the replication partner for this volume collection on the array group.
```

Clone Refresh

The Clone Refresh feature enables you to refresh a cloned Oracle database instance to a more recent point in time. This feature enables you to refresh a test/dev database instance that has been cloned from the production database to a more recent version of the data. This feature provides the ability to refresh a test/dev cloned instance to a more recent snapshot from the source instance.



The feature employs a new CLI option, **clone-refresh**, which is used to refresh an existing cloned instance. The feature essentially stitches together the Destroy Instance and Clone Instance workflows into a single workflow to refresh an existing cloned instance.

If the cloned instance has synchronous replication enabled, then the refreshed clone will also have synchronous replication enabled for the volume collection. The replication partner will need to be set manually after the clone-refresh completes. The **clone-refresh** option does not provide the option for inheriting or overriding parameters from the pfile of the source instance. These options are present in the **clone** command. Even if the cloned instance has been originally created with any of the clone options, they are not available in the **clone-refresh** option. In instances where these options are required, your DBA should manually clean up the cloned instance and create a new clone with the required options. Note that the **clone-refresh** option is disruptive as it deletes the clone instance's volumes along with the snapshots.

Before using the Clone Refresh feature, keep the following caveats in mind:

- The **clone-refresh** option will only work on standalone cloned instances.
- The cloned instance needs to be running for the **clone-refresh** option to proceed.
- The cloned instance will be cleaned up as part of the **clone-refresh** option, so it should not be in use when running **clone-refresh**.
- After a clone refresh, all the clone instance's snapshots created prior to running the **clone-refresh** option will be lost. Consequently, after refreshing a clone, you should immediately create a snapshot to back up your data.
- The CLI will accept only the cloned instance's sid and an optional snapshot name from the source instance to which the clone should be refreshed.
- This CLI option needs to run on the host on which the clone-sid is running.
- The snapshot name from which the clone should be refreshed is optional if the source instance from which the clone is built is running locally on the host. The tool will create a new snapshot of the source instance and base the refresh from this snapshot. However, if the source instance is remote, then the refresh snapshot name is mandatory.
- The refresh snapshot must be more recent than the one from which the cloned instance has been created.

Procedure

Clone refresh

```
noradatamgr --clone-refresh --clone-sid clone_instance_name --refresh-snapname snapshot_name [ --silent-mode]
```

- **clone-sid**: The cloned instance SID.
- **refresh-snapname**: The snapshot name of the source instance to which the cloned instance should be updated or refreshed.
- **silent-mode**: If this option is specified, the CLI will not prompt for confirmation to delete the instance and the diskgroup.

Example

With **--refresh-snapname**

```
[root]# noradatamgr --clone-refresh --clone-sid CLN1MAR --refresh-snapname
PRAG2019.MAR7.snap
The cloned instance and the underlying ASM diskgroups will be deleted and then
recreated.
Do you really want to proceed with destroying the instance?(yes/no)yes
Instance CLN1MAR has been cloned from source sid PRAG2019.
Starting tear down of instance CLN1MAR...
Instance CLN1MAR deleted.
Diskgroup CLN1MARDATADG deleted.
Diskgroup CLN1MARLOGDG deleted.
Success: Instance CLN1MAR and its backing volumes cleaned up successfully.
Rebuilding clone from source instance PRAG2019...
Initiating clone ...
[ ] 0%
[ ] 0%
[ ] 0%
```

```

[ ] 0%
[#### ] 9%
[#### ] 9%
[#### ] 9%
[#### ] 9%
[##### ] 18%
[##### ] 18%
[##### ] 18%
[##### ] 18%
[##### ] 27%
[##### ] 27%
[##### ] 27%
[##### ] 27%
[##### ] 36%
[##### ] 36%
[##### ] 36%
[##### ] 36%
[##### ] 45%
[##### ] 45%
[##### ] 45%
[##### ] 45%
[##### ] 54%
[##### ] 54%
[##### ] 54%
[##### ] 54%
[##### ] 63%
[##### ] 63%
[##### ] 63%
[##### ] 63%
[##### ] 72%
[##### ] 72%
[##### ] 72%
[##### ] 72%
[##### ] 81%
[##### ] 81%
[##### ] 81%
[##### ] 81%
[##### ] 90%
[##### ] 90%
[##### ] 90%
[##### ] 90%
[##### ] 100%

```

Success: Cloned instance PRAG2019 to CLN1MAR.

Diskgroup: CLN1MARDATADG

```

Disk: /dev/oracleasm/disks/CLN1MAR0001
Device: dm-14
Volume: CLN1MAR-DATA1
Serial number: 7102f354b295a0ff6c9ce9009a22694a
Size: 110GB

```

Diskgroup: CLN1MARLOGDG

```

Disk: /dev/oracleasm/disks/CLN1MAR
Device: dm-15
Volume: CLN1MAR-LOG1
Serial number: 62b278301f5c39da6c9ce9009a22694a
Size: 120GB

```

Allowed hosts: hiqa-rhel2

NPM enabled: No

Cataloged diskgroups: None

Replication type: Synchronous

Synchronous replication state: not_applicable

Volume collection has been created with synchronous replication. Please set the replication partner for this volume collection on the array group.

Without **--refresh-snapname**

```
[root]# noradatamgr --clone-refresh --clone-sid CLN1MAR
The cloned instance and the underlying ASM diskgroups will be deleted and then
recreated.
Do you really want to proceed with destroying the instance?(yes/no)yes
Instance CLN1MAR has been cloned from source sid PRAG2019.
Starting tear down of instance CLN1MAR...
Instance CLN1MAR deleted.
Diskgroup CLN1MARDATADG deleted.
Diskgroup CLN1MARLOGDG deleted.
Success: Instance CLN1MAR and its backing volumes cleaned up successfully.
Rebuilding clone from source instance PRAG2019...
Initiating clone ...
Initiating snapshot backup BaseForCLN1MAR-2002f035-d329-445e-9fc3-6b0191ea7408
for instance PRAG2019...
Success: Snapshot backup BaseForCLN1MAR-2002f035-d329-445e-9fc3-6b0191ea7408
completed.
[ ] 0%
[ ] 0%
[ ] 0%
[ ] 0%
[#### ] 9%
[#### ] 9%
[#### ] 9%
[#### ] 9%
[##### ] 18%
[##### ] 18%
[##### ] 18%
[##### ] 18%
[##### ] 27%
[##### ] 27%
[##### ] 27%
[##### ] 27%
[##### ] 36%
[##### ] 36%
[##### ] 36%
[##### ] 36%
[##### ] 45%
[##### ] 45%
[##### ] 45%
[##### ] 45%
[##### ] 54%
[##### ] 54%
[##### ] 54%
[##### ] 54%
[##### ] 63%
[##### ] 63%
[##### ] 63%
[##### ] 63%
[##### ] 72%
[##### ] 72%
[##### ] 72%
[##### ] 72%
[##### ] 81%
[##### ] 81%
[##### ] 81%
```

```

[#####] 81%
[#####] 90%
[#####] 90%
[#####] 90%
[#####] 90%
[#####] 100%
Success: Cloned instance PRAG2019 to CLN1MAR.
Diskgroup: CLN1MARDATADG
    Disk: /dev/oracleasm/disks/CLN1MAR0001
        Device: dm-14
        Volume: CLN1MAR-DATA1
        Serial number: 2ddd002233d5c6b76c9ce9009a22694a
        Size: 110GB
Diskgroup: CLN1MARLOGDG
    Disk: /dev/oracleasm/disks/CLN1MAR
        Device: dm-15
        Volume: CLN1MAR-LOG1
        Serial number: fe70bca676a125206c9ce9009a22694a
        Size: 120GB

Allowed hosts: hiqa-rhel2

NPM enabled: No

Cataloged diskgroups: None

Replication type: Synchronous
Synchronous replication state: not_applicable

Volume collection has been created with synchronous replication. Please set
the replication partner for this volume collection on the array group.

```

Delete a Snapshot of an Instance

Use this command to delete a snapshot of an instance. If this snapshot was created using the **--replica** option in the snapshot CLI, you can also use this command to delete the replicated snapshot, if one exists.

For synchronous replication enabled setups, note the following options:

- You can use the **--delete-replica** option to delete the snapshot along with the replica.
- You can delete a synchronously replicated snapshot from a pool by specifying the snapshot name. The snapshot on the replica has a different name than the upstream snapshot (appended by ".R"). Consequently, running the **--delete-snapshot** command with the replica name will delete only the replica without deleting the upstream snapshot.

Procedure

Delete a snapshot of an instance.

```
noradatamgr --delete-snapshot --instance instance_name --snapname snapshot_name [--delete-replica]
```

Example

Delete Snapshots of a Given Instance

```

[root]# noradatamgr --list-snapshot --instance rh68db
-----+-----+-----+-----
--
Snap Name                Taken at          Instance         Usable for cloning
-----+-----+-----+-----
--
rh68.snap.603            16-10-24 12:01   rh68db           Yes
rh68db.snap.601         16-10-21 10:36   rh68db           Yes

```



```
rh68db.snap.600          16-10-20 14:48  rh68db      Yes
rh68db.snap.602          16-10-21 15:21  rh68db      Yes
```

Delete a Specified Snapshot

```
[root]# noradatamgr --delete-snapshot --instance rh68db --snapname
rh68db.snap.602
Success: Snapshot rh68db.snap.602 deleted.
```

Verify the Specified Snapshot is Removed

```
[root]# noradatamgr --list-snapshot --instance rh68db
-----+-----+-----+-----
--
Snap Name                Taken at                Instance                Usable for cloning
-----+-----+-----+-----
--
rh68.snap.603            16-10-24 12:01         rh68db                  Yes
rh68db.snap.601         16-10-21 10:36         rh68db                  Yes
rh68db.snap.600         16-10-20 14:48         rh68db                  Yes
```

With **--delete-replica** option

```
[root]# noradatamgr --delete-snapshot --instance syncb --snapname syncb.snap1
--delete-replica
Synchronously replicated snapshot replica deleted successfully.
Success: Snapshot syncb.snap1 deleted.
```

Without **--replicate** option

```
[root]# noradatamgr --list-snapshots --instance syncb | grep syncb.snap2
syncb.snap2          19-03-07 12:28  default
Yes      CLI
syncb.snap2.R        19-03-07 12:28  default-group-sjc-array586Yes
CLI
[root]# noradatamgr --delete-snapshot --instance syncb --snapname syncb.snap2
Success: Snapshot syncb.snap2 deleted.
[root]# noradatamgr --list-snapshots --instance syncb | grep syncb.snap2
syncb.snap2.R        19-03-07 12:28  default-group-sjc-array586Yes
CLI
```

Destroy a Cloned Instance

This workflow destroys a cloned database instance with the specified SID. The database to be destroyed must be based on cloned volumes.

The workflow completes these activities:

- 1 Deletes the database instance
- 2 Deletes the underlying diskgroups in the database instance
- 3 Attempts to delete the volume collection associated with the volumes that back the disks for the instance

Procedure

Destroy a cloned database instance.

```
noradatamgr --destroy --instance instance_name
```

Example

Destroy a cloned database instance.

```
[root]# noradatamgr --destroy --instance DESTDB
Instance DESTDB deleted.
Do you really want to proceed with destroying the instance/diskgroup?(yes/no)yes
Diskgroup DESTDBDATADG deleted.
Diskgroup DESTDBLOGDG deleted.
Success: Instance DESTDB and its backing volumes cleaned up successfully.
```

Destroy an ASM Diskgroup

To destroy a diskgroup, the Oracle instance using that diskgroup must be shut down, and no active processes can be reading from or writing to the diskgroup. You can only destroy a diskgroup based on HPE Storage volumes that have been cloned.

Note that for synchronously replicated volumes, both upstream and downstream volumes will be cleaned up while destroying a diskgroup / database instance.

Procedure

Destroy an ASM diskgroup.

```
noradatamgr --destroy --diskgroup diskgroup_name
```

Example

Destroy an ASM Diskgroup

```
[root]# noradatamgr --destroy --diskgroup DESTDBLOGDG
Do you really want to proceed with destroying the instance/diskgroup?(yes/no)yes
Success: Diskgroup DESTDBLOGDG deleted.
```

Verify the specified ASM Diskgroup is removed

```
[grid]$ asmcmd lsdg
State      Type      Rebal   Sector  Block      AU      Total_MB  Free_MB
Req_mir   free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED   EXTERN   N        512     4096     1048576  102400    14554
0         14554
MOUNTED   NORMAL   N        512     4096     1048576  204800    89782
0         44891
MOUNTED   NORMAL   N        512     4096     1048576  204800    89782
0         44891
MOUNTED   EXTERN   N        512     4096     1048576  10236     10177
0         10177
MOUNTED   EXTERN   N        512     4096     1048576  51200     50956
0         50956
N        REDODG/
```

Overview of Using RMAN with NORADATAMGR

You can use the HPE Storage Oracle Application Data Manager (NORADATAMGR) to catalog the backup instances you created with NORADATAMGR so that they are listed as valid Oracle Recovery Manager (RMAN) backups. Then you can use RMAN to restore these backups.

Starting with HPE Toolkit for Linux 2.4.0, NORADATAMGR includes the **catalog-backup** command and the **uncatalog-backup** command. The **catalog-backup** command clones the diskgroups in the database instance snapshot captured by NORADATAMGR and registers them in the RMAN recovery catalog. RMAN can now access the data. After you use RMAN to restore the data, you must run the **uncatalog-backup** command. This command removes the cloned diskgroups from the recovery catalog, deletes the diskgroups, and cleans up the volumes.

You can only restore from one instance snapshot at a time. You must run the **uncatalog-backup** command after each restore operation to remove that instance from the recovery catalog. After the cleanup operation, you can restore data from another NORADATAMGR snapshot by using the **catalog-backup** command to add cloned diskgroups from another NORADATAMGR snapshot to the recovery catalog.

You always execute the **catalog-backup** and **uncatalog-backup** commands on the host where the target backup instance resides.

Note: You cannot recover an instance by using replicated downstream snapshots from a replication partner. You can, however, clone an instance from a replicated snapshot.

The following is a high-level overview of the steps required to enable RMAN to create a backup from NORADATAMGR data:

Before you begin

You must have the following:

- A recovery catalog. The NORADATAMGR target database backup instance must be registered with the recovery catalog.
- Login credentials and the correct role-based access control (RBAC) role are required by the recovery catalog. Normally, these are the operating system login credentials and the RBAC role RECOVERY_CATALOG_OWNER.
- Archive log mode should be enabled for the instance. RMAN requires this mode for recovery operations.
- The clock on the host (or server) should be synchronized with the HPE Storage group. If the NTP daemon has been disabled on the host, and the clocks on the host and the HPE Storage group are out-of-sync, you must make sure you choose the snapshot that contains the required changes for the restore operation. The creation-time for the snapshot is based on the HPE Storage group clock.
- Toolkit 2.4.0 or later installed on the host.

Note: Unless controlfile autobackup is enabled in RMAN, snapshots created using earlier versions of the HPE Storage Toolkit NORADATAMGR might not contain an RMAN controlfile backup. The controlfile is required for a successful recovery.

Procedure

1. Use NORADATAMGR to create a backup snapshot.
2. Locate the HPE Storage snapshot that has the backup data you need by executing the **noradatamgr --list-snapshot** command.
3. Confirm that none of the diskgroups are cataloged with the recovery catalog by executing the **noradatamgr --describe output** command.
4. If there no cataloged diskgroups, create an RMAN backup copy from the NORADATAMGR snapshot by executing the **noradatamgr --catalog-backup** command.

The **catalog-backup** command clones the diskgroups in the snapshot you selected and adds them to RMAN as a valid backup copy. It uses the format {PREFIX}-{DGTYPE(data/log)} to name the cloned diskgroups.

5. Restore the files using RMAN.

For information about using RMAN, refer to your Oracle documentation. You will need to connect the target database where the RMAN recovery catalog resides by executing RMAN commands.

Make sure you use any time up until the snapshot creation time in the UNTIL TIME line for the RMAN recovery. Doing this ensures that RMAN uses the cataloged diskgroup cloned from the snapshot for the recovery operation. After restoring to the snapshot, RMAN uses the archived redo logs to recover data up to the time that is provided in the RMAN command.

The following is an example of the type of RMAN commands you might need to execute to restore a backup instance called `rcat@hpenimble.hpe.com:1522/CATDB`:

```
$ rman target / catalog rcat/rcat@hpenimble.hpe.com:1522/CATDB
RMAN> shutdown immediate
RMAN> startup mount;
RMAN> run
```

```

    {
      SET UNTIL TIME "TO_DATE('2018 Apr 11 14:33','yyyy mon dd hh24:mi')";
      restore database;
      recover database;
    }
  RMAN> alter database open resetlogs

```

- Remove the cloned diskgroups from the RMAN recovery catalog, delete the diskgroups, and clean up the underlying volumes by running the **noradatamgr --uncatalog-backup** command on the target database host.

Note: You must always use the **uncatalog-backup** command to remove the current cataloged instance snapshot from the RMAN recovery catalog before you can recover another NORADATAMGR instance snapshot.

- Confirm that none of the diskgroups are cataloged with the recovery catalog by executing the **noradatamgr --describe output** command.
- If you want to use RMAN to recover another NORADATAMGR instance snapshot, repeat steps 1 through 6.

Add a NORADATAMGR Backup to an RMAN Recovery Catalog

The **catalog-backup** command included with HPE Storage Oracle Application Data Manager (NORADATAMGR) allows you to use Oracle Recovery Manager (RMAN) to create backups of data stored in snapshots created with NORADATAMGR. The command clones the diskgroups located in the NORADATAMGR snapshot and catalogs them so that they show up in the RMAN recovery catalog.

Note: This command is available with HPE Storage Toolkit for Linux 2.4.0 and later.

Before you run any RMAN commands, you must run the **catalog-backup** command on the host where the target database resides.

The **catalog-backup** command uses the format:

```

noradatamgr --catalog-backup --instance {target_db_SID} --snapname {snapshot_to_recover_from} --diskgroup-prefix {prefix_for_naming_cloned_diskgroups} --rcat {net_service_name_for_recovery_catalog} --rcat-user {recovery_catalog_user} [--rcat-passwd {password_for_rcat_user}]

```

If the RMAN recovery catalog is on the host where the target database is running, you can use the recovery catalog name as the value for the **--rcat** option. If the recovery catalog is on a remote host, you must use a net service name for the parameter.

For example, if you have a backup instance called PRODDB that has a listener running on port 1521, the net service name would be `node01.<company-domain.com>:1521/PRODDB`.

If you do not provide a value for the **--rcat_passwd** option, you will be prompted to enter a password when you execute the command.

To identify the cloned diskgroups, you must use the following format as the name:

```
{PREFIX}-{DGTYPE(data/log)}
```

Each backup instance can have only one set of cataloged diskgroups that can be used for the restore operation. If you need to use another NORADATAMGR snapshot as part of the restore operation, you must run the NORADATAMGR **uncatalog-backup** command to delete the existing cataloged diskgroups and then re-run the **catalog-backup** command with the other NORADATAMGR snapshot.

You must execute the **catalog-backup** command on the host where the target backup instance resides.

Before you begin

You must have the following:

- A snapshot backup or clone created using NORADATAMGR.
- A recovery catalog. The NORADATAMGR target database backup instance must be registered with the recovery catalog.

- Login credentials and the correct role-based access control (RBAC) role are required by the recovery catalog. Normally, these are the operating system login credentials and the RBAC role RECOVERY_CATALOG_OWNER.
- Toolkit 2.4.0 or later installed on the host.

Note: Unless controlfile autobackup is enabled in RMAN, snapshots created using earlier versions of the toolkit NORADATAMGR might not contain an RMAN controlfile backup. The controlfile is required for a successful recovery.

Procedure

1. Check to see if there are any cataloged diskgroups for that target instance by running the command **noradatamgr --describe --instance <instance_name>**.

This example uses proddb as the target instance.

```
# noradatamgr --instance proddb -d
Diskgroup: ARCHDG
    Disk: /dev/oracleasm/disks/ARCH
        Device: dm-3
        Volume: prod-arch
        Serial number: 6024de7dc66a60f46c9ce8004c92b39c
        Size: 120GB

Diskgroup: REDODG
    Disk: /dev/oracleasm/disks/REDO
        Device: dm-2
        Volume: prod-redo
        Serial number: 6db13e1bf84f48036c9ce8004c92b39c
        Size: 100GB

Diskgroup: DATADG
    Disk: /dev/oracleasm/disks/DATA
        Device: dm-5
        Volume: prod-data
        Serial number: 999cb6128f05af506c9ce8004c92b39c
        Size: 110GB

Allowed hosts: ALL
NPM enabled: No
Cataloged diskgroups: None
```

2. If there are no cataloged diskgroups, run **noradatamgr --catalog-backup** to catalog the diskgroup.

Note: If there are cataloged diskgroups, run **noradatamgr --uncatalog-backup** to remove them. Then run **noradatamgr --catalog-backup**.

```
# noradatamgr --instance proddb --catalog-backup
--snapname my.snapshot.4 --diskgroup-prefix MS4 --rcat tester.vlab.hpstor▶
age.com:1521/CATDB
--rcat-user rtest --rcat-passwd rtest
[#####] 100%
Successfully completed cataloging backups from snapshot my.snapshot.4.
Cloned and cataloged the following diskgroups: MS4LOGDG, MS4DATADG
#
```

3. You can now complete the restore operation using RMAN. For instructions about using RMAN, see your Oracle documentation.

Clean Up Cataloged Diskgroups and Volumes Used With RMAN

After you use Oracle Recovery Manager (RMAN) to recover data from an HPE Storage Oracle Application Data Manager (NORADATAMGR) snapshot, you should clean up the cataloged diskgroups and underlying HPE Storage volumes. The NORADATAMGR **uncatalog-backup** command removes the cloned, cataloged diskgroups that were used for the RMAN restore operation, deletes the diskgroups, and cleans up the volumes.

Note: This command is available with HPE Storage Toolkit for Linux 2.4.0 and later.

This command uses the format:

```
noradatamgr --uncatalog-backup --instance {target_db_SID} --rcat {net_service_name_for_recovery_catalog} --rcat-user {recovery_catalog_user} [--rcat-passwd {password_for_rcat_user}]
```

If the RMAN recovery catalog is on the target database host, you can use the recovery catalog name as the parameter for the **--rcat** option of the **uncatalog-backup** command. If the recovery catalog is on a remote host, you must use a net service name for the parameter.

For example, if you have a backup instance called PRODDB that has a listener running on port 1521, the net service name would be `node01.<company-domain.com>:1521/PRODDB`.

You do not need to include the diskgroup prefix that you created with the **catalog-backup** command because you can have only one set of cataloged NORADATAMGR diskgroups in the recovery catalog at one time. This command uses the instance name to locate the diskgroups.

If you do not provide a value for the **--rcat_passwd** option, you will be prompted to enter a password.

You must execute the **uncatalog-backup** command after each RMAN restore operation that uses data from a NORADATAMGR backup instance. This is because you can only restore from one NORADATAMGR instance backup at a time.

You must execute the **uncatalog-backup** command on the host where the target backup instance resides.

Before you begin

You must have the following:

- A recovery catalog. The NORADATAMGR target database backup instance must be registered with the recovery catalog.
- Login credentials and the correct role-based access control (RBAC) role are required by the recovery catalog. Normally, these are the operating system login credentials and the RBAC role RECOVERY_CATALOG_OWNER.
- HPE Storage Toolkit for Linux 2.4.0 or later installed on the host.

Note: Unless controlfile autobackup is enabled in RMAN, snapshots created using earlier versions of HPE Storage Toolkit for Linux NORADATAMGR might not contain an RMAN controlfile backup. The controlfile is required for a successful recovery.

You should only execute this command after you have performed an RMAN restore operation using an instance snapshot created by NORADATAMGR.

Procedure

1. Check to see if there are any diskgroups cataloged for the specified database.

Note: This example uses the instance rhtest and locates two cataloged diskgroups that use the prefix "MS4".

```
# noradatamgr --instance rhtest -d
Diskgroup: RHTESTARCHDG
  Disk: /dev/oracleasm/disks/RHTESTARCH
        Device: dm-3
        Volume: rhtest-arch
        Serial number: 6024md6nj66a60f46c9ce9002c92b39c
        Size: 120GB
```



```

Diskgroup: RHTESTDATADG
  Disk: /dev/oracleasm/disks/RHTESTDATA
    Device: dm-2
    Volume: rhtest-data
    Serial number: 5md13njb74f48036c9ce9002c92b59c
    Size: 100GB

Diskgroup: RHTESTREDODG
  Disk: /dev/oracleasm/disks/RHTESTREDO
    Device: dm-5
    Volume: rhtest-redo
    Serial number: 533md6128f93nj506c9ce9002c92b39c
    Size: 110GB

Allowed hosts: ALL
NPM enabled: No
Cataloged diskgroups: MS4DATADG,MS4LOGDG
#

```

2. Run --uncatalog-backup:

```

# noradatamgr --uncatalog-backup --instance proddb
--rcat tester.vlab.hpstorage.com:1521/CATDB --rcat-user rtest --rcat-passwd
rtest
Uncataloged RMAN backups for datafile and controlfile on cloned diskgroups.

Diskgroup MS4DATADG deleted.
Diskgroup MS4LOGDG deleted.
Success: Completed cleaning up cataloged backup and cloned diskgroups for
instance rhtest.

```

3. Verify that there is no longer any diskgroup that is cataloged.

```

# noradatamgr --instance rhtest -d
Diskgroup: RHTESTARCHDG
  Disk: /dev/oracleasm/disks/RHTESTARCH
    Device: dm-3
    Volume: rhtest-arch
    Serial number: 6024md6nj66a60f46c9ce9002c92b39c
    Size: 120GB

Diskgroup: RHTESTDATADG
  Disk: /dev/oracleasm/disks/RHTESTDATA
    Device: dm-2
    Volume: rhtest-data
    Serial number: 5md13njb74f48036c9ce9002c92b59c
    Size: 100GB

Diskgroup: RHTESTREDODG
  Disk: /dev/oracleasm/disks/RHTESTREDO
    Device: dm-5
    Volume: rhtest-redo
    Serial number: 533md6128f93nj506c9ce9002c92b39c
    Size: 110GB

Allowed hosts: ALL
NPM enabled: No
Cataloged diskgroups: None
#

```

RAC Database Recovery

If you have an Oracle Real Application Clusters (RAC) database, you can use the following procedure to recover an HPE Storage Oracle Application Data Manager (NORADATAMGR) backup instance of it using Oracle Recovery Manager (RMAN). You must add the database instance snapshot that you want to use for the recovery operation to the RMAN catalog.



Important: You must execute all the commands for this procedure from the same node.

Before you begin

You must have the following:

- A recovery catalog. The NORADATAMGR target database backup instance must be registered with the recovery catalog.
- Login credentials and the correct role-based access control (RBAC) role are required by the recovery catalog. Normally, these are the operating system login credentials and the RBAC role RECOVERY_CATALOG_OWNER.
- Archive log mode should be enabled for the instance. RMAN requires this mode for recovery operations.
- The clock on the host (or server) should be synchronized with the HPE Storage group. If the NTP daemon has been disabled on the host, and the clocks on the host and the HPE Storage group are out-of-sync, you must make sure you choose the snapshot that contains the required changes for the restore operation. The creation-time for the snapshot is based on the HPE Storage group clock.
- HPE Storage Toolkit for Linux 2.4.0 or later installed on the host.

Note: Unless controlfile autobackup is enabled in RMAN, snapshots created using earlier versions of toolkit NORADATAMGR might not contain an RMAN controlfile backup. The controlfile is required for a successful recovery.

Procedure

1. Run the **noradatamgr --catalog-backup** command to add the snapshot of the RAC database to the RMAN catalog. This is the snapshot you will use for the RMAN backup.
2. Convert the RAC database to a stand-alone database instance by setting the **CLUSTER_DATABASE** option to **false**.
3. Shutdown your system and then restart it and mount the database.
4. Restore the database using RMAN.

For information about using RMAN, refer to your Oracle documentation. You must connect the target database where the RMAN recovery catalog resides by executing RMAN commands.

Make sure you use any time up until the snapshot creation time in the **SET UNTIL TIME** line for the RMAN recovery. Doing this ensures that RMAN uses the cataloged diskgroup cloned from the snapshot for the recovery operation. After restoring to the snapshot, RMAN uses the archived redo logs to recover data up to the time that is provided in the RMAN command.

The following is an example of the type of RMAN commands you might need to execute to restore a backup instance called `rcat@hpenimble.hpe.com:1522/TESTDB`:

```
$ rman target / catalog rcat/rcat@hpenimble.hpe.com:1522/TESTDB
RMAN> shutdown immediate
RMAN> startup mount;
RMAN> run
  {
    SET UNTIL TIME "TO_DATE('2018 Apr 11 14:33','yyyy mon dd hh24:mi')";
    restore database;
    recover database;
  }
RMAN> alter database open resetlogs
```

5. When the RMAN recovery of the database is complete, remove the cloned diskgroups from the RMAN recovery catalog, delete the diskgroups, and clean up the underlying volumes by running the **noradatamgr --uncatalog-backup** command line on the target database host.

This command also cleans up the entries for the cataloged diskgroups from the crsctl and srvctl configuration. If these configurations are not cleaned up, the rest of the nodes in the RAC environment cannot be started. For information about manually cleaning up these configurations and restarting the RAC nodes, see [Bringing Up the RAC Nodes After an RMAN Recovery Operation](#) on page 46.

Note: You must always use the **uncatalog-backup** command to remove the current cataloged instance snapshot from the RMAN recovery catalog before you can recover another NORADATAMGR instance snapshot.

6. Change the **CLUSTER_DATABASE** option back to **true** and shut down the database.
7. Start database by using the **srvctl start database** command.
8. Confirm that the database has been started on all the RAC nodes.

Oracle NPM and Linux Toolkit Backup Service

Oracle NPM and the HPE Storage Toolkit for Linux Backup Service are two additional modules in the Linux Toolkit framework.



Important: Replication of database snapshots taken by Oracle NPM is not supported unless the downstream array in the replication partnership has array OS 4.x or a later version installed. If the downstream array does not meet this requirement, the replicated snapshots cannot be used to clone the database instance using the NoradataMgr tool.

Oracle NPM provides the ability to create Oracle-aware scheduled snapshots for an instance without need for doing this manually using the NORADATAMGR snapshot CLI. This makes use of the schedules that are present on a volume collection in the storage group.

The Backup Service can be started and stopped by running the **nltdm start** and **stop** commands. This service must be in the running state to use Oracle NPM. The status of the server can be viewed by running the **nltdm --status** command.

The Backup Service requires a group to be added in HPE Storage Toolkit for Linux for it to start. The port on which it starts can be controlled through the following configuration in `/etc/nlt.conf`:
`nlt.daemon.backup-service.port=9000`. By default, it starts on port 9000.

Hotbackup mode is not used by default for NPM scheduled snapshots, except in the case of a RAC instance. If hot-backup mode needs to be enabled for instances on the host, the following entry needs to be added to `/opt/NimbleStorage/etc/noradatamgr.conf`:

```
npm.hotbackup.enable=true
```

Any change in the `conf` files requires a restart of the Linux Toolkit service to take effect.

Requirements for enabling NPM for a particular database instance:

- Linux Toolkit Backup Service should be in the running state.
This can be verified by running the **nltdm --status** command.
- The array OS version of the group added to HPE Storage Toolkit for Linux should be 5.0.x or later.
- Running `hostname --fqdn` should display the full hostname along with the domain name (FQDN).

The group will try to access the host on this hostname.

NPM Commands

Use these commands to manage the Oracle NPM Service.

- To enable NPM, use

```
noradatamgr --enable-npm --instance {SID}
```

For more information, refer to [Enable NPM for an instance](#) on page 17.



- To disable NPM, use

```
noradatamgr --disable-npm --instance {SID}
```

For more information, refer to [Disable NPM for an instance](#) on page 18.

- To check the status of whether NPM has been enabled for an instance, use

```
noradatamgr --describe --instance {SID}
```

For more information, refer to [Describe an Instance](#) on page 15.

- To filter for snapshots that have been triggered by NPM, use

```
noradatamgr -list --instance <SID> --npm
```

For more information, refer to [List and Filter Snapshots of an Instance](#) on page 21.

Troubleshooting NORADATAMGR

Use the following sections to help troubleshoot any issues you might encounter:

Log Files

-NLT logs:

- /opt/NimbleStorage/log/{nlt.log + noradatamgr.log}

-Oracle logs:

- ASM alert log @ \$ORACLE_BASE/diag/asm/+asm/+ASM/trace/alert_+ASM.log
- Instance alert log @ \$ORACLE_BASE/diag/rdbms/<INSTANCE NAME>/<SID>/alert/alert_<SID>.log
- ASMLib log @ /var/log/oracleasm

SQL Plus

-Oracle SqlPlus:

- Connect to the +ASM instance using:
 - -sqlplus / as sysasm
 - -Views of interest: v\$ASM_DISKGROUP, v\$ASM_DISK
- Connect to the db instance using:
 - -sqlplus / as sysdba

Information for a noradatamgr Issue

Use the following as a guide when gathering information to report a troubleshooting issue related to noradatamgr:

-Get the ASM and storage footprint of the instance:

- ```
noradatamgr --describe --instance [SID]
```

-Get the group information:

- ```
nltadm --group --list
```

-Get information about the storage objects for the instance and their metadata:



- Get authentication token

```
curl -k -X POST -d '{"data":{"username":"<name>", "password":
"<password>"}}' https://<management-ip>:5392/v1/tokens
```

- Get information about volumes/volume collections/snapshot collections of the instance

```
curl -k -H "X-Auth-Token:<auth-token>" -X
GET https://xx.xx.xxx.xx:5392/v1/volume/<id>
curl -k -H "X-Auth-Token:<auth-token>" -X
GET https://xx.xx.xxx.xx:5392/v1/volume_collection/<id>
```

ASM Disks Offline When Controller Failure Occurs

On controller failover, an ASM diskgroup goes offline and the following error messages are seen in the ASM alert.log:

```
WARNING: Waited 15 secs for write IO to PST disk 0 in group 1.
WARNING: Waited 15 secs for write IO to PST disk 0 in group 1.
WARNING: Waited 15 secs for write IO to PST disk 0 in group 2.
NOTE: process _b000_+asm (4515) initiating offline of disk 0.3915944537 (DATA1)
with mask 0x7e in group 2
NOTE: process _b000_+asm (4515) initiating offline of disk 1.3915944538 (DATA2)
with mask 0x7e in group 2
```

This is a known issue with Oracle. Change the `_asm_hbeatiowait` from 15 to 120 seconds to avoid this. Run the following command in an ASM instance to set the desired value for `_asm_hbeatiowait`:

```
alter system set "_asm_hbeatiowait"=120 scope=spfile sid='*';
```

And then restart the ASM instance / crs, to take new parameter value in effect.

noradatamgr Command Fails after HPE Storage Toolkit for Linux Installation

A `noradatamgr` command fails soon after HPE Storage Toolkit for Linux installation with the following error message:

```
ERROR CliClient: - The required files were not found. Make sure service is
running using nltadm --status
```

Also, the Oracle-App-Data-Manager service is shown as “STOPPED” in `nltadm --status`.

This error message is expected when array information is not added into Toolkit for Linux. Add group information using the `nltadm --group --add --ip-address <IP> --username <admin user>` command, and start the service using the `nltadm --start oracle` command. The `noradatamgr` command should now work.

Snapshot and Clone Workflows Fail if spfile not Specified

An spfile must be specified for each DB instance for snapshot and clone workflows. It gives the following error in the log:

```
ORA-01565: error in identifying file '?/dbs/spfile@.ora'
ORA-27037: unable to obtain file status
```

An spfile must be used. To create an spfile from memory, use the following command:

```
SQL> create pfile='/tmp/initFGTEST.ora.5881b834-b914-4df2-8da5-7a46e4abd87e'
```

Failover causes ASM to mark disks offline or Oracle database fails

Failover on the array causes ASM to mark disks offline and/or become inaccessible for a very short period of time, or the Oracle database fails to start due to ASM not being able to locate or find its disks.



If any of the above issues occur, verify that Oracle ASMlib has the correct parameters set in the `/etc/sysconfig/oracleasm` file. For information, see [HPE Storage Oracle Application Data Manager](#) on page 11.

Bringing Up the RAC Nodes After an RMAN Recovery Operation

If you are working with Oracle Real Application Clusters (RAC), you might have an issue where the RAC nodes are not able to start after you perform a recovery operation using Oracle Recovery Manager (RMAN). This can happen when entries from the cataloged disk groups appear in the `crsctl` and `srvctl` repositories.

In a RAC recovery operation, cloned diskgroups are added to the `crsctl` and `srvctl` repositories after you run the **catalog-backup** command and the RMAN recovery commands to recover the database instance.

If the RAC nodes are not starting after the recovery operation finishes, you must remove these entries manually.

Procedure

1. Confirm that the cloned diskgroups were added to the `crsctl` repository by executing the **crsctl status res -t** command.
2. If there are diskgroups in the `crsctl` repository, remove them by executing the **crsctl delete resource 'ora.<DG-NAME>.dg' -f** command.
 <DB-NAME> is the name of the diskgroup.
3. Check to see whether any diskgroups were added to the `srvctl` configuration dependency list for the target RAC database by executing the **srvctl config database -d <RAC-DB-NAME>** command.
 <RAC-DB-NAME> is the name of the RAC database.
4. If there are diskgroups in the dependency list, update the `srvctl` configuration to remove them by executing the **srvctl modify database -d <RAC-DB-NAME> -a '<COMMA-SEPARATED-LIST-DISKGROUPS>'** command.
 <RAC-DB-NAME> is the name of the RAC database.
 <COMMA-SEPARATED-LIST-DISKGROUPS> contains the names of the diskgroups remaining in the dependency list of the RAC database. You must separate the names of the diskgroups using commas. Do not include any diskgroups that were cataloged and cloned.
 The nodes in the RAC instance should now be able to come back up.
5. If you did not run the **uncatalog-backup** command after the RMAN recovery operation, you should run it now.



HPE Storage Connection Manager for Linux

This section covers the configuration and use of HPE Storage Connection Manager for Linux on the Linux operating system. You can use the connection manager to:

- Manage multipath connections (at the I/O region level) to volumes striped across multiple arrays.
- Migrate volumes between storage pools.
- Merge two storage pools within a group.
- Automatically manage iSCSI and multipath configuration.
- Provide flexibility to mount volumes based on volume names.
- Configure block device level settings for optimal performance.
- Ensure that connection redundancy is always maintained to the array.

Configuration

The HPE Storage Connection Manager for Linux installation involves both automatic and manual configuration. After you complete the installation, you might want to verify some of the settings.

The basic configuration includes the following tasks, which are explained later in this chapter:

- DM multipath setup. The connection manager handles this automatically.
- Selecting the scale-out mode. If you select scale-out mode, you also need to change the Logical Volume Management (LVM) filters.
- Protocol configuration. The connection manager configures settings based on best practice recommendations for the iSCSI and Fibre Channel protocols. You will need to use the **ncmadm** command to discover devices for the protocol you are using. You can also use it when you are working with iSCSI topology scenarios.
- Block device settings. By default, these settings are disabled. You can enable them.
- iSCSI connection handling. Information in the connection manager `/etc/ncm.conf` file controls the number of sessions that are created for each array per each available target portal in a group.
- Setting up support for SAN boot devices. The connection manager rebuilds the `initramfs` file under the `/boot` directory.
- Creating and mounting filesystems. You can use the **ncmadm** command to list the available devices.

ncmadm Workflows

One of the tools you use when configuring and verifying your HPE Storage Connection Manager for Linux installation is the **ncmadm** command. You can use **ncmadm --help** to get information about the workflows available with this command.

```
ncmadm --help
usage:
  ncmadm [--list]
  ncmadm [--longlist]
  ncmadm [--info] <DEVICE>
  ncmadm [--rename]
  ncmadm [--diag {CASE_NUMBER}]
  ncmadm [--rescan {DISCOVERY_IP}]
  ncmadm [--fcrescan]
  ncmadm [--refresh <DEVICE>]
  ncmadm [--logout <TARGET_IQN|all>]
  ncmadm [--remove <DM_DEVICE|all>] [--force]
  ncmadm [--resize {DM_DEVICE}]
  ncmadm [--unbind <IFACE>]
  ncmadm [--bind NIC [--iface IFACE]]
  ncmadm [--config --iscsi configure|validate]
```



```

ncmadm [--config --mpath configure|validate]
ncmadm [--config --target_scope <group|volume>
        --device <device path|all>]
-b,--bind <NIC>          Bind network interface (for example
                        eth2) to iSCSI
-c,--config              Check and apply recommended best
                        practice settings
-d,--diag <CASE_NUMBER> Collect NCM diagnostic data
-device,--device <arg>  Volume name whose target scope being
                        changed.
-e,--rescan <DISCOVERY_IP> Discover, login and create new
                        Nimble devices
-fcrescan,--fcrescan    Rescan FC devices
-force,--force          Attempt to forcefully remove the
                        device even when in use.
-h,--help               ncmadm usage information
-I,--iface <IFACE>     Create iface. If not specified, it
                        will be generated automatically
                        during bind
-info,--info <DEVICE>  Display information about the
                        specified Nimble Device
-iscsi,--iscsi          Check and apply Nimble recommended
                        iSCSI settings.
-l,--list                List Nimble devices.
-ll,--longlist           List Nimble devices.
-mpath,--mpath           Check and apply Nimble recommended
                        multipath settings.
-o,--logout <TARGET>   iSCSI logout on the specified target
                        name (iqn). This will not cleanup
                        the dm device stack. Use 'all' as
                        target to logout from ALL Nimble
                        targets.
-r,--refresh <DEVICE>  Refresh the specified Nimble device
                        table and map
-remove,--remove <DM_DEVICE> Remove the specified Nimble device
                        with complete dm stack. Use 'all'
                        option to cleanup ALL Nimble
                        devices.
-rename,--rename         Refresh volume names from array. Use
                        this if volume name is changed on
                        array.
-resize,--resize <DM_DEVICE> Trigger rescan on the host to pick
                        up the new volume size after
                        resizing of the volume on the array.
-target_scope,--target_scope <arg> Converts iscsi target scope of a
                        specified volume
-u,--unbind <IFACE>    Unbind network interface and remove
                        specified iface

```

DM Multipath Configuration

The connection manager automatically applies the recommended configuration for HPE Storage devices in `/etc/multipath.conf` for both RHEL 6.x and 7.x OS versions.

After the storage connection installation, use the following command to verify and configure multipath settings for both RHEL 6.x and 7.x OS versions:

ncmadm --config --mpath validate

If the settings are not correct, enter the following command to correct them:

ncmadm --config --mpath configure

HPE Storage recommends enabling `user_friendly_names` in the defaults section. Refer to [multipath.conf Settings](#) on page 66 for example settings.

Considerations for Using the HPE Storage Connection Manager for Linux in Scale-out Mode

During HPE Storage Toolkit for Linux installations, you can choose scale-out mode for the connection manager. In silent-install mode, you can use the `--enable-scaleout` option to do this. By default, scale-out mode is disabled. Scale-out mode for the connection manager creates additional `/dev/dm` and `/dev/mapper` devices managed by the `dm-switch` kernel module, present in RHEL 6.7 and 7.1 onwards. Scale-out mode is recommended only if you have a volume pool striped across multiple arrays in a group. For single array use cases, this is not necessary. For a striped volume pool, scale-out mode (`dm-switch` devices), ensures that no performance overhead is encountered by properly routing I/O to the correct array in the pool based on the I/O region.

The following considerations and limitations must be taken into account before you enable scale-out mode:

- 1 Applications and filesystems must be reconfigured to make use of these new devices rather than traditional `/dev/mapper/mpath` devices.
- 2 A host reboot is necessary when installing using this mode.
- 3 If Oracle DB is installed with ASMLIB packages, then scale-out mode is not supported. Also if aliases are configured for device names in `/etc/multipath.conf` then scale-out mode is not recommended.
- 4 Scale-out mode is not supported in OpenStack and synchronous replication environments.
- 5 Scale-out mode is not supported on Ubuntu and SUSE Linux OS.

LVM Filter for the HPE Storage Connection Manager for Linux in Scale-Out Mode

If you install HPE Storage Connection Manager for Linux with with scale-out mode (`dm-switch`) devices enabled, then you must change the Logical Volume Manager (LVM) filter to use these new devices.

You can add or modify the following filters in `/etc/lvm/lvm.conf` to use LVM devices with the connection manager.

After you add them, you should run the `lvdisplay/pvdisplay` commands to check and correct any parsing errors (such as missing quotes) that are reported in the output before you proceed.

```
filter = ["r|^/dev/nimblestorage/lower_tier_devices/.*/"]
global_filter = ["r|^/dev/nimblestorage/lower_tier_devices/.*/"]
preferred_names = [ "^/dev/nimblestorage/", "^/dev/mapper/mpath", "^/dev/[hs]d"]
```

When you work with LVM filters, keep the following in mind:

- If LVM filters are not set up properly with the connection manager, I/O forwarding can be seen on the array, and LVM commands display "duplicate device found" errors.
- For SAN boot systems, if any changes are made to `lvm.conf` after the connection manager installation, you must sync the changes with `lvm.conf` in `initramfs` using the following command:

```
dracut --add "multipath ncm" -f /boot/initramfs-'uname -r'-ncm.img
```

iSCSI and FC Configuration

During installation, the HPE Storage Connection Manager for Linux automatically configures settings depending upon whether you have an iSCSI stack or a Fibre Channel stack.

For an iSCSI stack, the connection manager configures the following, recommended settings in the `/etc/iscsi/iscsid.conf` file:

- `noop_out_interval = 5`
- `noop_out_timeout = 10`
- `replacement_timeout = 10`
- `queue_depth = 64`

After you the connection manager installation completes, you can use the **ncmadm** command to verify the settings or change them. The command has the following format:

ncmadm --config --iscsi validate | configure

If you change the settings, you must log out and log back in before they take effect.

For a Fibre Channel stack, the connection manager automatically configures the following recommended transport timeout settings in the `/etc/multipath.conf` file:

- `dev_loss_tmo=infinity`
- `fast_io_fail_tmo=5`

Block Device Settings

The connection manager will copy a udev rule to assist in tuning block settings for devices (under `/etc/udev/rules.d/99-nimble-tune.rules`). Based on the application needs, you can enable the following settings:

Note: These settings are disabled by default.

- `max_sectors_kb` set to `max_hw_sectors_kb`
- `nr_requests` set to 512
- `scheduler` set to "noop"

Discover Devices for iSCSI

The first time HPE Storage Connection Manager for Linux discovers iSCSI devices, it uses the iSCSI discovery IP address. After the initial discovery is done, you no longer need to specify the discovery address.

If you add volumes to an initiator group on an array later, you can use the command **ncmadm --rescan** on the host. The connection manager automatically scans and logs into newly mapped volumes.

Note: You might want to verify what is configured on your HPE Storage system before performing this task. You can use the following commands to provide information about your system:

- **vol --list**
 - **vol --info ***
 - **initiatorgrp --list**
 - **initiatorgrp --info ***
-

Procedure

1. Get the iSCSI discovery IP address of the array that you configured. You can use the Network Configuration screen in the GUI to do this.





Network Configuration

[Network Configurations](#) | [View](#)

Group
Subnets
Interfaces
Diagnostics

Management IP
Used for the Web UI, CLI and replication. Resides on management subnet and floats across all "Mgmt only" and "Mgmt + Data" interfaces on that subnet.

IP Address	Network	Netmask
10.18.236.71	10.18.236.0	255.255.252.0

iSCSI Host Connection Method

Manual ?
 Automatic ?
 Enable rebalancing ?

Discovery IP Addresses

Subnet Label	Discovery IP	Netmask
data1	172.18.212.84	255.255.252.0
data2	172.18.216.76	255.255.252.0
mgmt-data		255.255.252.0

Routes

Default Gateway 10.18.236.1

Network	Netmask	Gateway

Static Routes

Note: The *GUI Administration Guide* and the *CLI Administration Guide* contain information about working with HPE Storage arrays.

- Perform the discovery using the `ncmadm --rescan <discovery_ip>` command.

Note: You only need to include the iSCSI discovery IP address the first time you discover devices. You can perform all successive discovery operations on the host by entering the `ncmadm --rescan` command without the discovery address.

The following example of an initial discovery uses an iSCSI discovery IP address of 172.18.212.84.

```
-bash-4.1# ncmadm --rescan 172.18.212.84
Discovery IP :
172.18.212.84
List of discovered devices:
target:iqn.2007-11.com.nimblestorage:linux-vol-3-
v5406993b0c2187fd.0000025d.58271ec6 portal:172.18.216.76:3260,2460

Login to the new devices:
target:iqn.2007-11.com.nimblestorage:linux-vol-3-
v5406993b0c2187fd.0000025d.58271ec6 portal:172.18.216.76:3260,2460

Rescan initiated. Use ncmadm -l to list Nimble device
```

Discover Devices for Fibre Channel

Procedure

1. Use the **ncmadm --fcrescan** command to discover the Fibre Channel-based volumes.

```
[root@~]# ncmadm --fcrescan
Fc rescan initiated
```

2. Use the **ncmadm --list** command to display the discovered devices.

Note: You can also follow the Red Hat documentation to discover new volumes on the host using `rescan-scsi-bus.sh`. However, connection manager-based rescan is optimized to discover active paths prior to standby paths, which makes active-path discovery faster.

List Devices

Procedure

1. Use the **ncmadm -l** command to list the discovered devices.

```
[root@t_rhel7u0x64 ~]# ncmadm -l
size=100G uuid=2d8874e3b8ac7d5fe6c9ce9002ae1012c mount-device: /dev/nimblestor▶
age/linux-vol-2d8874e3b8ac7d5fe6c9ce9002ae1012c
target: iqn.2007-11.com.nimblestorage:linux-vol-
v3e28972c12297606.00000003.2c01e12a connection_mode=automatic
` + array_id=1 device=/dev/dm-0 sync_replicated=no
  |- 7:0:0:0 sdb 8:16 session=5 sess_st=LOGGED_IN dev_st=running
  iface=default tgtportal=172.16.39.66 data_ip=172.16.39.65
  |- 5:0:0:0 sdd 8:48 session=3 sess_st=LOGGED_IN dev_st=running
  iface=default tgtportal=172.16.231.66 data_ip=172.16.231.65
  |- 6:0:0:0 sde 8:64 session=4 sess_st=LOGGED_IN dev_st=running
  iface=default tgtportal=172.16.231.66 data_ip=172.16.231.65
  |- 8:0:0:0 sdf 8:80 session=6 sess_st=LOGGED_IN dev_st=running
  iface=default tgtportal=172.16.39.66 data_ip=172.16.39.65
```

2. Use the **ncmadm -ll** command to list the discovered devices along with its faulty paths.

```
[root@t-rhel8u0x64 ~]# ncmadm -ll
size=1.0G uuid=219a9af8dbd9262876c9ce9003af4c32f mount-device: NA
target: NA connection_mode=NA
` + array_id=-1 device=/dev/dm-2 sync_replicated=no
  |- 11:0:0:0 sdb 8:16 session=9 sess_st=LOGGED_IN dm_st=failed chk_st=faulty
  dev_st=offline iface=default tgtportal=172.16.226.92 data_ip=172.16.226.92
  |- 12:0:0:0 sdc 8:32 session=10 sess_st=LOGGED_IN dm_st=failed chk_st=faulty
  dev_st=offline iface=default tgtportal=172.16.34.92 data_ip=172.16.34.92
```

iSCSI Topology Scenarios

Single Subnet Topology

If the host has a single subnet, redundancy can be achieved by creating logical iSCSI interfaces as shown in the figure below. Logical iSCSI interfaces can be created using the command:

ncmadm -b eth0


where `eth0` is the data interface on the host.



Note: Use iSCSI port binding only if all Ethernet ports you intend to use are configured with a single IP subnet and broadcast domain.


Regarding sessions:

- New sessions will use these interfaces.
- For existing sessions to use these interfaces, you must log out and then rescan.
- If additional interfaces are added when sessions are already present on other interfaces, the connection manager automatically creates connections on the missing interfaces during the next interval.



Multiple Subnet Topology

Ethernet interfaces are configured with different IP subnets (eth0 and eth1) as shown in the figure below; therefore, don't create iSCSI interfaces. Sessions are created on the default interface on both target portals (192.168.1.10 and 192.168.2.10 in the figure below).



iSCSI Connection Handling

The following parameters have been added to `/etc/nmc.conf` to control the number of sessions to be created per array, per available target portal in a group, in automatic connection mode. HPE Storage Connection Manager for Linux attempts to create connections through each interface for single subnet scenarios. If multiple target portals are available, then the connection manager will try to make connections to each one of them.

For static (manual) connection mode, the connection manager will not dynamically manage iSCSI sessions.

A service restart of the connection manager is necessary if any of these values are changed manually.

```
#Maximum number of sessions to be created for each array in Nimble Group per
iSCSI Portal (Maximum <=8)
max_sessions_per_array=8
#Minimum number of sessions to be created for each array in Nimble Group per
```

```
iSCSI Portal (Minimum >=1)
min_sessions_per_array=2
```

Note:

For SAN boot systems, you must rebuild the `initramfs` file after any changes to `ncm.conf`. Use the following command that is appropriate for your system.

If the connection manager is installed with scale-out mode disabled, the command is:

```
dracut --add "multipath" --force /boot/initramfs-`uname -r`-ncm.img
```

If scale-out mode is enabled (multi-array group), the command is:

```
dracut --add "multipath ncm" --force /boot/initramfs-`uname -r`-ncm.img
```

where ``uname -r`` returns the current running kernel version.

SAN Boot Device Support

It is not recommended to create a root/boot device on a volume that is striped across multiple arrays in a group. If a striped volume is used, then I/O performance degradation may occur.

New initramfs Creation and GRUB Updates

As part of the installation, HPE Storage Connection Manager for Linux rebuilds the `initramfs` file under the `/boot` directory with naming format

```
initramfs-`uname -r`-ncm.img
```

where ``uname -r`` is the current running kernel version.

This is required to keep `/etc/multipath.conf`, `/etc/lvm/lvm.conf`, and other connection manager files in sync with `initramfs` and to apply correct configuration parameters during the early boot process.

After the connection manager installation has completed, a new menu entry is added to the GRUB configuration file. This entry contains the newly created `initramfs` file.

An example `grub.conf` entry is shown below:

```
title Red Hat Enterprise Linux 6 with NCM (2.6.32-573.el6.x86_64)
  root (hd0,0)
  kernel /tboot.gz logging=vga,serial,memory
  module /vmlinuz-2.6.32-573.el6.x86_64 ro root=UUID=893e4a15-8ac2-
    41dc-94aa-0d7ff13cc82a intel_iommu=on amd_iommu=on rd_NO_LUKS
    KEYBOARDTYPE=pc KEYTABLE=us LANG=en_US.UTF-8 rd_NO_MD
    SYSFONT=latacyrheb-sun16 crashkernel=auto rd_NO_LVM rd_NO_DM
    rhgb quiet rhgb quiet
  module /initramfs-2.6.32-573.el6.x86_64-ncm.img
```

Any additional changes to `/etc/multipath.conf` or `/etc/lvm.conf` after the connection manager installation requires you to manually update the `initramfs` created by the connection manager. The command for this is as follows:

```
dracut --add multipath --f /boot/initramfs-`uname -r`-ncm.img
```

If the connection manager was installed in scale-out mode, use the command below:

```
dracut --add "multipath ncm" --f /boot/initramfs-`uname -r`-ncm.img
```

Note: The connection manager makes the newly added entry the default selection during boot. Therefore, the connection manager configuration parameters are applied by default during the early boot process.

Create and Mount Filesystems

Create file systems on top of the mount-device listed in the output from command **ncmadm -l**. The symbolic links to devices on which filesystems can be created are listed in the `/dev/nimblestorage` directory.

```
-bash-4.1# ls -l /dev/nimblestorage/
total 0
lrwxrwxrwx. 1 root root 8 Mar 14 15:03 linux-vol-1-
255a3eb706da66cc76c9ce900c61e2758 -> ../dm-12
lrwxrwxrwx. 1 root root 8 Mar 14 15:03 linux-vol-2-
2b83629bac96756356c9ce900c61e2758 -> ../dm-10
lrwxrwxrwx. 1 root root 7 Mar 14 15:03 linux-vol-3-
2e7f4333b869bc3d36c9ce900c61e2758 -> ../dm-6
lrwxrwxrwx. 1 root root 7 Mar 14 15:03 linux-vol-4-22b726298fb5bd▶
bc16c9ce900c61e2758 -> ../dm-3
drwxr-xr-x. 2 root root 120 Mar 9 17:52 lower_tier_devices
```

Create and mount file systems as shown below.

```
-bash-4.1# mkfs.ext4 /dev/nimblestorage/linux-vol-1-
255a3eb706da66cc76c9ce900c61e2758
mke2fs 1.41.12 (17-May-2010)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621440 blocks
131072 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2684354560
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 27 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
-bash-4.1#

-bash-4.1# mkdir /mnt/linux-vol-1
-bash-4.1# mount /dev/nimblestorage/linux-vol-1-255a3eb706da66cc76c9ce900c61e2758
/mnt/linux-vol-1/
-bash-4.1#

-bash-4.1# mount
/dev/mapper/255a3eb706da66cc76c9ce900c61e2758 on /mnt/linux-vol-1 type ext4
(rw)
-bash-4.1#
```

Mandatory mount options include the following:

- Specify the **_netdev** option in `/etc/fstab` to mount iSCSI based volumes. For example:

```
/dev/nimblestorage/linux-vol-1-255a3eb706da66cc76c9ce900c61e2758 /mnt/linux-
vol-1 ext4 defaults,rw,_netdev 0 0
```

- For RHEL 7.x hosts, also specify the **nofail** option to prevent boot issues due to temporary mount failures. For example:

```
/dev/nimblestorage/linux-vol-1-255a3eb706da66cc76c9ce900c61e2758 mnt/linux-  
vol-1 ext3 defaults, __netdev, nofail 0 0
```

Note: Red Hat recommends that filesystems always be mounted by UUID. You can use volume names for ease of use, but you can mount filesystems either way.

Note: Even if filesystems are mounted using `/dev/nimblestorage/` symbolic link names, the **mount** command will only display `/dev/mapper/UUID` switch device names or `/dev/mapper/mpath` device names. This is expected.

LVM Usage

To create logical volume management (LVM) on top of a volume, use the following steps:

Procedure

1. Make sure LVM filters are set up as mentioned in [LVM Filter for the HPE Storage Connection Manager for Linux in Scale-Out Mode](#) on page 49.
2. Create physical volumes as below.

```
-bash-4.1# pvcreate /dev/nimblestorage/linux-vol-2-  
2b83629bac96756356c9ce900c61e2758  
Physical volume "/dev/nimblestorage/linux-vol-2-  
2b83629bac96756356c9ce900c61e2758" successfully created  
-bash-4.1#  
  
-bash-4.1# pvcreate /dev/nimblestorage/linux-vol-3-  
2e7f4333b869bc3d36c9ce900c61e2758  
Physical volume "/dev/nimblestorage/linux-vol-3-  
2e7f4333b869bc3d36c9ce900c61e2758" successfully created  
-bash-4.1#  
  
-bash-4.1# pvdisplay  
--- Physical volume ---  
PV Name                /dev/sda2  
VG Name                vg_rkumarlinuxredhat67  
PV Size                15.51 GiB / not usable 3.00 MiB  
Allocatable           yes (but full)  
PE Size                4.00 MiB  
Total PE              3970  
Free PE               0  
Allocated PE          3970  
PV UUID               xJazhZ-J1rH-uYs5-nSmf-xMdB-Lt2D-CJOVkr  
  
"/dev/nimblestorage/linux-vol-3-2e7f4333b869bc3d36c9ce900c61e2758" is a new  
physical volume of "10.00 GiB"  
--- NEW Physical volume ---  
PV Name                /dev/nimblestorage/linux-vol-3-  
2e7f4333b869bc3d36c9ce900c61e2758  
VG Name                  
PV Size                10.00 GiB  
Allocatable           NO  
PE Size                0  
Total PE              0  
Free PE               0  
Allocated PE          0  
PV UUID               AMrsb2-VrYM-k9hM-Qb3I-AEAh-1UeL-AnYZ9G
```



```

"/dev/nimblestorage/linux-vol-2-2b83629bac96756356c9ce900c61e2758" is a new
physical volume of "10.00 GiB"
--- NEW Physical volume ---
PV Name                /dev/nimblestorage/linux-vol-2-
2b83629bac96756356c9ce900c61e2758
VG Name
PV Size                10.00 GiB
Allocatable           NO
PE Size               0
Total PE              0
Free PE               0
Allocated PE          0
PV UUID               0uxL1b-vu7T-1Cfh-tYJQ-VcHQ-IIh3-xpo3Jw

```

3. Create a volume group using the physical volumes created above.

```

-bash-4.1# vgcreate nimble_volume_group /dev/nimblestorage/linux-vol-2-
2b83629bac96756356c9ce900c61e2758 /dev/nimblestorage/linux-vol-3-
2e7f4333b869bc3d36c9ce900c61e2758
Volume group "nimble_volume_group" successfully created
-bash-4.1#

```

```

-bash-4.1# vgdisplay
--- Volume group ---
VG Name                nimble_volume_group
System ID
Format                 lvm2
Metadata Areas         2
Metadata Sequence No  1
VG Access              read/write
VG Status              resizable
MAX LV                 0
Cur LV                0
Open LV                0
Max PV                 0
Cur PV                2
Act PV                 2
VG Size                19.99 GiB
PE Size                4.00 MiB
Total PE               5118
Alloc PE / Size        0 / 0
Free PE / Size         5118 / 19.99 GiB
VG UUID                ad08aT-SEcM-P7kI-vSrG-EVNX-Myzu-JCLnFU

```

4. Create logical volumes in the volume group created above.

```

-bash-4.1# lvcreate -L10G -n nimble_logical_volume nimble_volume_group
Logical volume "nimble_logical_volume" created.

-bash-4.1# lvdisplay
--- Logical volume ---
LV Path                /dev/nimble_volume_group/nimble_logical_volume
LV Name                nimble_logical_volume
VG Name                nimble_volume_group
LV UUID                prV5hg-5roS-QMmk-8r45-d23m-kren-6MNB2y
LV Write Access        read/write
LV Creation host, time rkumar-linux-redhat67, 2016-03-14 15:31:24 -0700
LV Status               available

```

```
# open          0
LV Size        10.00 GiB
Current LE     2560
Segments      2
Allocation     inherit
Read ahead sectors
- currently set to 256
Block device   253:5
```

5. Create file systems on a logical volume and mount them.

```
-bash-4.1# mkfs.ext4 /dev/nimble_volume_group/nimble_logical_volume
mke2fs 1.41.12 (17-May-2010)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621440 blocks
131072 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2684354560
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 21 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
-bash-4.1#
```

```
-bash-4.1#
-bash-4.1# mkdir /mnt/nimble_logical_volume
-bash-4.1#
-bash-4.1# mount /dev/nimble_volume_group/nimble_logical_volume /mnt/nimble_logi▶
cal_volume
-bash-4.1#

-bash-4.1# mount
/dev/mapper/255a3eb706da66cc76c9ce900c61e2758 on /mnt/linux-vol-1 type ext4
(rw)
/dev/mapper/nimble_volume_group-nimble_logical_volume on /mnt/nimble_logi▶
cal_volume type ext4 (rw)
-bash-4.1#
```

Note: For iSCSI devices, pass the `_netdev` and `nofail` options during mount.

Major Kernel Upgrades and HPE Storage Connection Manager for Linux

The connection manager installation copies different sets of files for both RHEL 6.x and 7.x OS versions. Therefore, when the system is upgraded from RHEL 6.x to 7.x, the connection manager must be reinstalled as well.



After reinstalling the connection manager, reboot the system for proper device management. If the system is not rebooted, the connection manager will not discover the devices and application I/O might fail.

Removing a LUN from the Host

Before removing (also known as unmapping) the LUN from the host, follow the Red Hat recommendation for steps to properly shut down applications, unmount filesystems and remove devices from LVM if used. Red Hat recommended steps can be found in the [Red Hat Storage Administration Guide](#)

Use the **ncmadm --remove** command to clean up the complete device stack, and for iSCSI to log out from the target. If the device removal is successful, the LUN can be unmapped from the initiator group.

Refer to the **ncmadm** man page for more details on the **--remove** option.

Known Issues and Troubleshooting Tips

This section describes the most common issues. Refer to the HPE Storage Connection Manager for Linux release notes for more information and troubleshooting steps, if any of these issues are observed.

Information for a Linux Issue

Use the following as a guide when gathering information to report a troubleshooting issue related to Linux:

-Verify dm-multipath devices, UUID not populated, orphan paths, etc.

- `multipath -ll`
- `multipathd show paths format "%i %o %w %d %t %C"`
- `dmsetup table | egrep "mpath|switch"`
- `multipath.conf` settings

-Verify iSCSI connections if paths don't appear under multipath.

- `iscsiadm -m session -P3`
- `lsscsi -gld`
- Verify if port binding is enabled with dual subnet network configurations

-Collect sosreport and NLT diag logs.

- `sosreport`
- `nltadm --diag`

HPE Storage Toolkit for Linux Fails During Reboot When /opt Mounted to Another Device

By default, some Linux operating systems mount `/opt` to a different device than the one where the HPE Storage Toolkit for Linux is installed. As a result, HPE Storage Toolkit for Linux cannot locate its dependencies and fails during a reboot. This happens because `/opt/NimbleStorage` does not mount before the HPE Storage Toolkit for Linux service starts.



If your system mounts `/opt` to a separate device, HPE recommends that you install HPE Storage Toolkit for Linux within the root device (`/opt/NimbleStorage/`). Use the `--install-dir` option with HPE Storage Toolkit for Linux installer to specify a different installation directory. You could enter the following command:

```
HPE_Storage_toolkit_for_linux --silent-mode --install-dir /nimblestorage --ncm
```

You can also use this command to install other HPE Storage Toolkit for Linux features. For example, you might enter:

```
HPE_Storage_toolkit_for_linux --silent-mode --install-dir /nimblestorage --ncm --oracle --docker --flexvolume
```

Old iSCSI Targets IQNs Still Seen after Volume Rename on the Array

Connections to old iSCSI target IQN are still seen after volume is renamed on the array.

To resolve this issue, you must rename the iSCSI-based volume.

Procedure

1. Offline the volume and rename.
2. Log out of the iSCSI target using `ncmadm --logout <old target iqn>`
3. Set the volume online.
4. Do a rescan on the host using `ncmadm --rescan`.
5. Run `ncmadm --rename` to refresh the names.

iSCSI_TCP_CONN_CLOSE Errors

iSCSI_TCP-CONN_CLOSE errors are displayed in a system with a large number of iSCSI volumes.

```
Jan 26 16:04:49 hiqa-sys-rhel2 kernel: connection480:0: detected conn error (1020) ISCSI_ERR_TCP_CONN_CLOSE
```

These errors might be caused by `iscsid` using duplicate iSCSI session IDs (isids) if there are more than 256 sessions. This will cause the target to disconnect existing sessions and create new ones. For more information, refer to https://bugzilla.redhat.com/show_bug.cgi?id=1309488

- This issue is fixed in the RHEL 6.x version of `iscsi-initiator-utils-6.2.0.873-18.el6`. For RHEL 7.x, this issue is not fixed at the time of this documentation.

rescan-scsi-bus.sh Stops Responding or Takes a Long Time While Scanning Standby Paths

`rescan-scsi-bus.sh` has a `testonlineO` function which issues TUR and retry messages 8 times, with an interval of 1 second sleep. This causes a delay when scanning a large number of volumes if standby paths are present. Issuing TUR on standby paths returns a 2/4/0b check condition. This condition is not handled in `rescan-scsi-bus.sh` and therefore loops in `testonlineO`.

This is a known issue and raised with Red Hat, at https://bugzilla.redhat.com/show_bug.cgi?id=1299708.

- To work around the issue, comment out the “sleep 1” statement in the `testonlineO` function.

Long Boot Time With a Large Number of FC Volumes

Under some conditions, it may take a long time to boot if there are a large number of FC volumes mapped to the host.

If standby paths are first discovered on the SCSI bus during boot, any SCSI discovery commands on device-mapper devices will cause repeated “reinstated” messages and failed events. This might slow down the entire boot process, as `multipathd` is busy processing failed events. For more information, refer to https://bugzilla.redhat.com/show_bug.cgi?id=1291406

- A fix is available for RHEL 7.x in `device-mapper-multipath-0.4.9-88.el7`.



Duplicate Devices Found Warning

LVM commands output warnings stating that duplicate devices were found . For example:

```
Found duplicate PV LXUG8bh0Xbnvt1pyUxBj1BRDFzQjqw1x: using /dev/nimblestor▶
age/linux-vol-2d8874e3b8ac7d5fe6c9ce9002ae1012c not /dev/mapper/mpatha

Found duplicate PV LXUG8bh0Xbnvt1pyUxBj1BRDFzQjqw1x: using /dev/mapper/mpathb
not /dev/nimblestorage/linux-vol-2d8874e3b8ac7d5fe6c9ce9002ae1012c
```

This is not an error condition, but a warning indicating multiple device mapper devices appear with the same LVM metadata (for example, dm-switch, dm-multipath).

- By default, the HPE Storage Connection Manager for Linux will apply filters to ignore all dm-mpath devices if the dm-switch module is supported on the host. Refer to [LVM Filter for the HPE Storage Connection Manager for Linux in Scale-Out Mode](#) on page 49 to apply the correct LVM filters in `/etc/lvm/lvm.conf`.

Newly Added iSCSI Targets Not Discovered

The connection manager does not discover newly added iSCSI targets on rescan based on the discovery IP address.

- Verify whether iSCSI port binding is used in a dual subnet IP address configuration. iSCSI port binding (iface) is only recommended in single IP address subnet configurations.
- Verify network settings, and whether the target discovery IP address can be pinged from the host.
- If Jumbo frames (MTU 9000 bytes) are enabled on the host, make sure they are enabled end to end in the network all the way to the target.

System Enters Maintenance Mode on Host Reboot with iSCSI Volumes Mounted

On reboot, the RHEL 7.x iSCSI system goes into maintenance mode.

- Make sure the **_netdev** option is added in `/etc/fstab` to allow the network to come up before the file systems are mounted.
- Add the **nofail** option and manually mount all the file systems using the command **mount -a** after system reboot.
- Reboot the system.

hung_tasks_timeout Errors

During controller upgrades, a controller might take more than 150 seconds with large configurations, and hung_tasks_timeout errors might appear in `/var/log/messages`. For example:

```
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: INFO: task jbd2/sdb-8:2321
blocked for more than 150 seconds.
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel:      Not tainted 2.6.32-
431.17.1.el6.x86_64 #1
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: "echo 0 > /proc/sys/ker▶
nel/hung_task_timeout_secs" disables this message.
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: jbd2/sdb-8    D 0000000000000000
0 2321      2 0x00000000
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: ffff8802360d9c20
0000000000000046 ffff8802360d9bb0 ffffffff81083e2d
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: ffff8802360d9b90
ffffffff810149b9 ffff8802360d9bd0 ffffffff810a6d01
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: ffff880239585098
ffff8802360d9fd8 00000000000000fbc8 ffff880239585098
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: Call Trace:
May 4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: [<ffffffff81083e2d>] ?
del_timer+0x7d/0xe0
```

```
May  4 11:11:00 rtp-fuji-qa31-2-rhel65-6 kernel: [<ffffffff810149b9>] ?
read_tsc+0x9/0x20
```

- You can safely ignore these messages. The connection manager will configure `multipathd` to queue I/Os for more than 150 seconds. Therefore, these messages will appear during controller reboot/upgrade scenarios. I/O will continue once connectivity is restored to the controller.

RHEL 7.x: FC System Fails to Boot

The RHEL 7.x: FC System fails to boot and drops into an emergency shell after failing to mount `/boot` on a multipath device.

When `multipathd` creates a new multipath device, no more paths can be added until the `udev` event changes for the multipath device being created, even if the device was created with no usable paths. Therefore, `udev` stops responding while trying to get information on the device, and bootup can timeout.

The issue is resolved in update of `device-mapper-multipath-0.4.9-95.el7`. Update to this package using the **yum update device-mapper-multipath** command.

RHEL 7.x: Fibre Channel Active Paths are not Discovered Promptly

The RHEL 7.x: FC System fails to discover active paths in a timely manner.

When `multipathd` creates a new multipath device, no more paths can be added until the `udev` event changes for the multipath device being created, even if the device was created with no usable paths. Therefore, `udev` stops responding while trying to get information on the device. Until `multipathd` times out, no active paths can be added to the device.

The issue is resolved in update of `device-mapper-multipath-0.4.9-95.el7`. Update to this package using the **yum update device-mapper-multipath** command.

During Controller Failover, multipathd Failure Messages Repeat

On Fibre Channel systems, the following message can appear repeatedly in `/var/log/messages`:

```
Nov 23 15:10:36 hitdev-rhel67 multipathd: 8:96: reinstated
Nov 23 15:10:36 hitdev-rhel67 multipathd: mpathal: queue_if_no_path enabled
Nov 23 15:10:36 hitdev-rhel67 multipathd: mpathal: Recovered to normal mode
Nov 23 15:10:36 hitdev-rhel67 multipathd: mpathal: remaining active paths: 1
Nov 23 15:10:36 hitdev-rhel67 kernel: sd 8:0:5:1: alua: port group 02 state S
```

The issue is fixed in `device-mapper-multipath-0.4.9-88.el7` on RHEL 7.x and `device-mapper-multipath-0.4.9-94.el6` for RHEL 6.x versions.

LVM Devices Missing from Output

Logical volume management (LVM) devices on non-HPE Storage volumes are missing from **lvdisplay/pvdisplay** output.

- Verify the LVM filters in `/etc/lvm/lvm.conf` using the information in [LVM Filter for the HPE Storage Connection Manager for Linux in Scale-Out Mode](#) on page 49.

Various LVM and iSCSI Commands Hang on /dev/mapper/mpath Devices

By default, the connection manager enables infinite queuing on `dm-multipath` devices.

If all paths of a multipath device are removed or logged out, there may be empty multipath maps on the host. Empty maps can be verified using the **multipath -ll** command. Empty multipath maps with queuing enabled (that is, `queue_if_no_path` enabled), will cause any command issued to hang until at least one path is restored.

- If the intention is to remove the device completely, then use the **ncmadm --remove /dev/dm-x** command to disable queuing and properly clean up the device. If that fails, manually disable infinite queuing using the command **dmsetup**

message mpathx 0 fail_if_no_path where *mpathx* is replaced with the actual device name from the `multipath -ll` output. Contact support to recover the paths if the LUN was not unmapped from the host and the I/O is hung.

Devices Not Discovered by Multipath

Verify that none of the HPE Storage devices are blacklisted in `multipath.conf`. Common errors of this sort can happen with the patterns listed below. Avoid these in the blacklist section:

```
blacklist {
    wwid *
    devnode sd*
    device {
        vendor *
        product *
    }
}
```

- Correct the blacklist section and run `multipathd reconfigure` to apply the changes.

System Drops into Emergency Mode

On a RHEL 6.8 SAN boot system, sometimes if the `/boot/initramfs-<kernel version>.img` image is rebuilt after adding or removing volumes to the host using `dracut`, the system fails to boot and drops into emergency mode.

Always make a backup of the default `initramfs` image on the SAN boot system before overwriting it. (For example, `cp /boot/initramfs-`uname -r`.img /boot/initramfs-`uname -r`.img.bak`).

The HPE Storage Toolkit for Linux always creates a new `initramfs` image on installation with `-ncm` appended to the name, and the entry is added to the GRUB menu list. If boot is failing, select the default image name from the GRUB menu during boot. (Press 'e' during boot loader).

HPE Storage Toolkit for Linux Installer by Default Disables dm-switch Devices on Upgrades

The HPE Storage Toolkit for Linux installer by default disables the creation of `dm-switch` devices to handle striped volumes. If `dm-switch` devices were previously enabled (either by the connection manager 1.0 or by HPE Storage Toolkit for Linux 2.0 with the `--enable-scaleout` option), they will be disabled on a fresh install if the default options are used, and the following warning message is displayed:

```
WARNING: dm-switch(scale-out) devices are in use, please revert LVM rules and
reboot host to disable scale-out configuration.
```

Reboot the host after HPE Storage Toolkit for Linux installation to cleanup `dm-switch` devices. If the volume is striped across multiple arrays and `dm-switch` devices are needed, re-install HPE Storage Toolkit for Linux with `--ncm` and `--enable-scaleout` options. This will ensure that the scale-out host configuration is not modified.

HPE Storage Connection Manager for Linux with RHEL HA Configuration

This section describes additional configuration steps and limitations when using the connection manager on RHEL HA nodes. Refer to the following Red Hat cluster configuration guides for detailed steps to configure an HA cluster and fencing methods:

- [Red Hat Enterprise Linux 6 Cluster Administration Guide](#)
- [Red Hat Enterprise Linux 7 High Availability Add-On Administration Guide](#)

The reservation_key Parameter

To install the connection manager in an HA environment, choose `fence-mpath` instead of `fence-scsi` agent if an SCSI PR reservation fencing method is used.



Add the **reservation_key** option to the "Nimble" device section in `/etc/multipath.conf`.

For example, to add key "0x123abc":

```
devices {
  device {
    vendor      "Nimble"
    product     "Server"
    .
    .
    reservation_key "0x123abc"
  }
}
```

After this change, run `multipathd reconfigure` for the change to take effect.

This must be done on each node in the cluster, with a node-specific reservation key. This reservation key must match with the corresponding entry in the `/etc/cluster/cluster.conf` file for the node.

Note: I/O errors with status **reservation_conflict** can occur, if the **reservation_key** option is not added to `multipath.conf`. This is because paths can be added and deleted dynamically, and the `multipathd` process will not automatically register keys on those new paths without the **reservation_key** configuration option.

Note: RHEL 6.x and 7.x versions do not support the APTPL (Active Persist Through Power Loss) flag for SCSI PR reservations. Therefore, SCSI PR registrations/reservations will be lost on controller upgrades or reboots. This might also lead to reservation conflict I/O errors.

Volume Move Limitations in HA Environments

If a volume is moved from one storage pool to another storage pool within a group, new `/dev/mapper/mpath` devices will be created for destination arrays. Once the volume move is complete, `mpath` devices representing source arrays will become stale. The connection manager will remove these devices once the volume move is complete.

However, `/etc/cluster/cluster.conf` will still reference these old `/dev/mapper/mpath` names, and cluster fence services might fail during startup. Therefore, once the volume move is complete, all the stale `mpath` devices must be removed from `/etc/cluster/cluster.conf` and new entries can be added based on `dm-*` devices from the `ncmadm -l` output.

For example, before a volume move, `lower_tier` `mpath` device:

```
/dev/mapper/mpatha
<fencedevices>
  <fencedevice agent="fence_mpath" debug="/var/log/fence_mpath.log" de▶
vices="/dev/mapper/mpatha" name="nimble-mpath-fencing"/>
</fencedevices>
```

After volume move, `lower_tier` `mpath` device:

```
./dev/mapper/mpathb
<fencedevices>
  <fencedevice agent="fence_mpath" debug="/var/log/fence_mpath.log" de▶
vices="/dev/mapper/mpathb" name="nimble-mpath-fencing"/>
</fencedevices>
```

If `/etc/cluster/cluster.conf` still has references to the source array `/dev/mapper/mpatha`, cluster services might fail during startup. To prevent this from happening, add the updated `mpath` names and remove the old ones from `/etc/cluster/cluster.conf` on all the nodes. Sync the config file to all nodes in the cluster, and restart cluster services as necessary for the new changes to take effect.



iSCSI Target Scope Conversion Utility

The HPE Storage Toolkit for Linux 4.0 includes an iSCSI target scope conversion utility that simplifies the iSCSI target scope conversion.

A new option has been added to ncmadm CLI that enables you to configure iSCSI target scopes to either group scope or volume scope. This option enables you to convert an iSCSI target scope with just one command.

Features

You can use the conversion utility to accomplish the following tasks:

- Change the iSCSI target scope to a group scope from volume scope or to a volume scope from the group scope.
- Change the iSCSI target scope of a specific device or all iSCSI devices of a specific group (seen from the initiator from where you issue the command).

Prerequisites

Before using the conversion utility, make sure of the following:

- You have added the IP address of the group where the iSCSI target resides to the HPE Storage Toolkit for Linux database, and the group is accessible to the initiator from where you convert the iSCSI target scope.
- The multipath config matches HPE Storage best practice configuration.
- The device is not busy when you attempt to convert the iSCSI target scope. If the device is serving data continuously, you must stop the I/O, unmount the device, and then perform the conversion.
- The target is not being accessed by other initiators, other than the one in which you issue the command to convert the iSCSI target scope
- The target scope conversion to group scope will not work if safe_logout is set to "yes" and the device is mounted. Therefore, make sure safe_logout is not set to "yes."

Change an iSCSI Target Scope

To change an iSCSI target scope, run the following command:

```
ncmadm [--config --target_scope <group|volume> --device <device path|all>]
```

Note: The target scope conversion to group scope is a disruptive operation.

Target Scope Conversion Messages

When changing an iSCSI target scope, you might see the following information or error messages:

Running basic checks

Information message indicating that before attempting a target scope configuration change, the HPE Storage Toolkit for Linux is running basic checks, such as group access, multipath configuration and others.

Failed to connect to group

Error message indicating that a group is not reachable from the host where you are attempting an iSCSI target scope configuration.

NCM not supported on this OS version.

Error message indicating that connection manager is not supported on a given OS version.

multipath.conf settings are NOT configured according to Nimble best practice recommendation

Error message indicating that the multipath configuration is not recommended by HPE Storage.

Volume not present on the group configured, use nltadm command to add group

Error message indicating that a volume is not hosted on the group that is configured in the HPE Storage Toolkit for Linux database.



Configuring target scope for volume <volume name>

Information message indicating the informing volume name that is operated upon.

Target scope change not required for this volume <volume name>

Information message indicating that the specified volume does not require a target scope change because the scope of the volume is the scope that is expected. Example: you attempt to change scope of volume to group and the specified volume is already in the group scope.

Target scope change skipping for " <volume name> "as it is connected to multiple initiators

Information message indicating that the specified volume has connections from initiators other than the one from where you are attempting to change the iSCSI target scope.

Volume/s not ready OR Target scope change not required

Information message indicating that all or a specified volume is either not ready for a target scope change or a target scope change is not required.

Target scope configured as per the request. Starting volume reinstate

Information message indicating that the target scope configuration completed and the volumes are being brought back up..

Paths for volume <volume name> taking more time to come online, try ncmadm -l later to list Nimble device

Information message indicating that a volume is taking time to come online. You can issue the **ncmadm -l** command to re-check after some time has elapsed to see if volume has come online.

Target scope config completed successfully for volume <volume name>

Information message indicating that target scope configuration completed successfully.

This process might take time.. Please wait...

Information message indicating that the Target scope configuration might take some time to complete.

Target scope config failed. Reason: Initiators failed to logout.

Error message indicating that the target scope configuration failed because the initiators failed to log out. This message is displayed when you attempt to convert the target scope to group and the device is busy.

multipath.conf Settings

This section can be used as a reference to validate `/etc/multipath.conf` settings. The connection manager automatically configures a "device" section for HPE Storage devices as part of installation. For example:

```
defaults {
    user_friendly_names yes
    find_multipaths      no
}
blacklist {
    devnode "^ (ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    device {
        vendor  ".*"
        product ".*"
    }
}
blacklist_exceptions {
    device {
        vendor  "Nimble"
        product "Server"
    }
}
devices {
    device {
        vendor          "Nimble"
        product         "Server"
        path_grouping_policy group_by_prio
    }
}
```

```

    prio                "alua"
    hardware_handler    "1 alua"
    path_selector       "service-time 0"
    path_checker        tur
    no_path_retry       30
    failback            immediate
    fast_io_fail_tmo    5
    dev_loss_tmo        infinity
    rr_min_io_rq        1
    rr_weight           uniform
}
}

```

The above blacklist and blacklist_exceptions configuration causes multipathd to blacklist all non-HPE Storage devices. If there are other vendor devices connected to the host, remove these entries and manually blacklist just the local disk using wwid <local disk wwid> under the blacklist section.

You must also replace the sample wwid information with information for the local disk on your system.

After making changes to the `/etc/multipath.conf` file, you must run the **multipathd reconfigure** command for the changes to take effect.

You should also disable `user_friendly_names` for SLES based systems as per SUSE recommendations for multipath settings or where native multipathing is used.



Nimbletune

Nimbletune is a Linux host tuning utility that is installed with the HPE Storage Toolkit for Linux. Nimbletune can be used to automate various Linux OS settings for optimal performance with arrays. Nimbletune enables you to list all current settings and provides options to configure the settings to.

After you have installed the Linux toolkit, you can locate the Nimbletune files in the `/opt/NimbleStorage/nimbletune` directory.

There are five categories of settings:

- 1 Filesystem
- 2 Disk
- 3 Multipath
- 4 iSCSI
- 5 Fibre Channel (FC)

Note:

The default Recommendations are optimal for most workloads. However, you may choose to change some of the values to ones that are more suitable to your environment. Refer to [Change Nimbletune Recommendations](#) on page 74 for more information.

Filesystem and FC recommendations need to be applied manually. Nimbletune lists only the recommended settings for these two categories. Refer to the host OS administration or HBA user guide to change these settings. Run Nimbletune with the **--verbose** option for more information

There is no need to reboot the host when changes are made to Nimbletune recommendations, and the changes persist through subsequent reboots.

Nimbletune Commands

Examples of Nimbletune commands are documented in this section. Each example contains tabular output with the following headings:

- Category: The category of the recommendation. The categories are filesystem, disk, multipath, iSCSI, and FC.
- Device: Indicates whether the recommendation applies to a particular device
- Parameter: Name of the parameter
- Value: Current value of the parameter
- Recommendation: Recommended value for the parameter
- Status: Current status of the parameter; not-recommended indicates that the current setting needs to be changed.
- Severity: Severity level of the parameter. Nimbletune supports three levels of severity:
 - 1 Critical: This is a mandatory setting, which, if not applied may cause data loss.
 - 2 Warning: Functionality is not affected, but this setting can cause performance degradation.
 - 3 Info: A low-priority setting that might improve performance.

Get Nimbletune Help

Enter the following command to display help for the Nimbletune utility:

```
# nimbletune --help
```



Example

```
#nimbletune --help
Nimble Linux Tuning Utility

Usage:

nimbletune --get [-category {filesystem | multipath | disk | iscsi | fc | all}]
  [-status {recommended | not-recommended | all}] [-severity {critical | warning
  | info | all}] [-verbose] [-json] [-xml]

nimbletune --set [-category {multipath | disk | iscsi | all}]

Options:
-c, -category      Recommendation category {filesystem | multipath | disk
  | iscsi | fc | All}. (Optional)

-st, -status       Recommendation status {recommended | not-recommended
  | All}. (Optional)

-sev, -severity    Recommendation severity {critical | warning | info |
  All}. (Optional)

-json             JSON output of recommendations. (Optional)

-xml              XML output of recommendations. (Optional)

-global           If true, settings will be configured globally at host
  level wherever applicable(eg iscsid.conf),default:false

-verbose          Verbose output. (Optional)

-v, -version       Display version of the tool. (Optional)
```

Note:

- The **-json** option displays output in JSON format.
- The **-xml** option displays output in XML format.
- The **-verbose** option displays verbose output.

Get All System Recommendations

Enter **--get** with no options to display all Nimbletune recommendations grouped by category.

nimbletune --get**Example**

```
# nimbletune --get

Recommendations for disk:
+-----+
+
+-----+-----+-----+-----+-----+-----+
|Category|Device| Parameter      | Value  |Recommend| Status      |Severity|
+-----+-----+-----+-----+-----+-----+
+
| disk   | all   | add_random     | 0      | 0        | recommended |warning|
| disk   | all   | rq_affinity    | 1      | 2        | not-recommended|warning|
| disk   | all   | scheduler      | deadline| noop     | not-recommended|warning|
```

```
| disk      | all  | rotational  | 0      | 0      | recommended  | warning |
| disk      | all  | nr_requests  | 128    | 512    | not-recommended | warning |
| disk      | all  | max_sectors_k | 512    | 4096   | not-recommended | warning |
| disk      | all  | read_ahead_kb | 4096   | 128    | not-recommended | warning |
+-----+
```

Recommendations for iscsi:

```
+-----+
| Cat  | Dev | Parameter          | Value  | Recommend | Status      | Severity |
+-----+
| iscsi| all | login_timeout     | 15     | 15        | recommended | warning  |
| iscsi| all | noop_out_timeout  | 10     | 10        | recommended | warning  |
| iscsi| all | cmds_max          | 128    | 256       | not recommended | info    |
| iscsi| all | queue_depth       | 64     | 64        | recommended | warning  |
| iscsi| all | nr_sessions       | 1      | 1         | recommended | info    |
| iscsi| all | startup           | automatic | automatic | recommended | warning  |
| iscsi| all | replacement_timeout | 10    | 10        | recommended | warning  |
| iscsi| all | iscsi_port_binding | enabled | disabled  | not recommended | critical |
+-----+
```

Recommendations for multipath:

```
+-----+
| Category | Device | Parameter          | Value  | Recommendation | Status |
| Sev      |        |                    |        |                 |        |
+-----+
| multipath| all   | fast_io_fail_tmo  | 5      | 5              | recom▶
mend|critical|
| multipath| all   | path_selector     | "service-time 0" | "service-time 0" | recom▶
mend|warning|
| multipath| all   | vendor            | "Nimble" | "Nimble"       | recom▶
mend|critical|
| multipath| all   | path_checker      | tur     | tur            | recom▶
mend|critical|
| multipath| all   | product           | "Server" | "Server"       | recom▶
mend|critical|
| multipath| all   | hardware_handler  | "1 alua" | "1 alua"       | recom▶
mend|critical|
| multipath| all   | failback          | immediate | immediate      | recom▶
mend|warning|
| multipath| all   | dev_loss_tmo     | infinity | infinity       | recom▶
mend|critical|
| multipath| all   | prio              | "alua"  | alua           | recom▶
mend|critical|
```

```

| multipath|all |no_path_retry      |30          |30          |recom►
mend|critical|

| multipath|all |path_grouping_policy|group_by_prio |group_by_prio |recom►
mend|critical|
+-----+
-----

Recommendations for fc:
+-----+
|Category| Device| Parameter      |Value |Recommend| Status          |Severity|
+-----+
|fc      | all  | ql2xmaxqdepth| 32   | 64       |not-recommended | warning|
+-----+

```

Get Recommendations Filtered by Category

Get recommendations filtered by **--category** to view current settings. The example gets recommendations for disk settings.

```
# nimbletune --get --category disk
```

Example

```

# nimbletune --get --category disk
Recommendations for disk:
+-----+
+-----+
| Category | Device | Parameter      | Value | Recommendation | Severity|
+-----+
| disk     | all   | add_random     | 1     | 0              | warning|
| disk     | all   | rq_affinity    | 1     | 2              | warning|
| disk     | all   | scheduler      | cfq   | noop           | warning|
| disk     | all   | rotational     | 0     | 0              | warning|
| disk     | all   | nr_requests    | 128   | 512            | warning|
| disk     | all   | max_sectors_kb| 512   | 4096           | warning|
+-----+
+-----+

```

Get Recommendations Filtered by Status

Get recommendations filtered by **--status** to view current settings. The example returns all not-recommended settings.

```
# nimbletune --get --status not-recommended
```

Example

```

# nimbletune --get --status not-recommended
Recommendations for disk:
+-----+
+-----+
| Category | Device | Parameter      | Value | Recommendation | Severity|
+-----+
+-----+

```



```

| disk      | all      | add_random      | 1      | 0
|           | not-recommended | warning|
| disk      | all      | rq_affinity     | 1      | 2
|           | not-recommended | warning|
| disk      | all      | scheduler       | cfq    | noop
|           | not-recommended | warning|
| disk      | all      | nr_requests     | 128    | 512
|           | not-recommended | warning|
| disk      | all      | max_sectors_kb  | 512    | 4096
|           | not-recommended | warning|
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+

Recommendations for fc:
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+
| Category | Device  | Parameter          | Value  | Recommendation▶
|          | Status  | |Severity|
+-----+-----+-----+-----+-----+
| fc       | all     | ql2xmaxqdepth     | 32     | 64
|          | not-recommended | warning|
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+

```

Get Recommendations Filtered by Category and Status

Get recommendations filtered by **--category** and **--status** to view current settings. The example returns all not-recommended settings for disk.

```
# nimbletune --get --category disk --status not-recommended
```

Example

```

# nimbletune --get --category disk --status not-recommended
Recommendations for disk:
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+
| Category | Device  | Parameter          | Value  | Recommendation▶
|          | Status  | |Severity|
+-----+-----+-----+-----+-----+
| disk     | all     | add_random      | 1      | 0
|           | not-recommended | warning|
| disk     | all     | rq_affinity     | 1      | 2
|           | not-recommended | warning|
| disk     | all     | scheduler       | cfq    | noop
|           | not-recommended | warning|
| disk     | all     | nr_requests     | 128    | 512
|           | not-recommended | warning|
| disk     | all     | max_sectors_kb  | 512    | 4096
|           | not-recommended | warning|
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+

```

Apply Recommendations by Category

Use the **--set** command to apply recommendations by category. The example sets the recommendations for disk.



nimbletune --set --category disk**Example**

```
# nimbletune --set --category disk
Successfully applied disk recommendations
```

The following command sets the recommendations for iscsi at volume level:

nimbletune --set -category iscsi**Example**

```
# nimbletune --set -category iscsi
Successfully applied iscsi recommendations
```

The following command sets the recommendations for iscsi at host level:

nimbletune --set -category iscsi -global true**Example**

```
# nimbletune --set -category iscsi -global true
Successfully applied iscsi recommendations
```

Note: Once the settings has been applied at the host level, using `nimbletune --set -category iscsi` applies the settings at volume level, but does not revert the changes that has been made at the host level when the `nimbletune --set -category iscsi -global true` command was used. If you need the previous settings or new settings at host level, you must manually edit `iscsid.conf`.

Verify Recommendations Set by Category

Verify recommendations by category. The example verifies the settings for the disk category.

nimbletune --get --category disk**Example**

```
# nimbletune --get --category disk
Recommendations for disk:
+-----+-----+-----+-----+-----+
| Category | Device   | Parameter           | Value | Recommendation |
| dation   | Status  | |Severity|         |       |                |
+-----+-----+-----+-----+-----+
| disk     | all     | add_random         | 0     | 0              | | |
|          | recommended | | warning|         |       |                |
| disk     | all     | rq_affinity        | 2     | 2              |
|          | recommended | | warning|         |       |                |
| disk     | all     | scheduler           | noop  | noop           |
|          | recommended | | warning|         |       |                |
| disk     | all     | rotational         | 0     | 0              |
|          | recommended | | warning|         |       |                |
| disk     | all     | nr_requests        | 512   | 512            |
|          | recommended | | warning|         |       |                |
| disk     | all     | max_sectors_kb     | 4096  | 4096           |
|          | recommended | | warning|         |       |                |
+-----+-----+-----+-----+-----+
```



Change Nimbletune Recommendations

Nimbletune uses a common JSON configuration file to validate host settings. You can change the recommendation values by directly editing the file.

Procedure

1. Edit the configuration file.
vim /opt/NimbleStorage/Nimbletune/config.json
2. Change the recommendation value for any parameter as needed.
3. Save and exit the file.
4. Verify that the recommendations have been changed as expected.

nimbletune --get



Block Reclamation

Space usage monitoring on a SAN is different from how space usage is monitored within a host's file system. A SAN reports free space in terms of how many blocks have not been written to. These blocks are sometimes referred to as "clean blocks". The number of clean blocks is multiplied by the block size to provide a more user-friendly space usage figure.

In contrast, host file systems report free space in terms of the total capacity of a datastore or volume, less the sum of all files within the file system. When a file is deleted, free space is instantly increased within the host file system. However, in the majority of cases, deleting files on the host does not automatically notify the SAN that those blocks can be freed up, since the physical block remains in place after the deletion; in such cases, only the file system metadata is updated. This leads to a discrepancy between how much free space is being reported within the file system and how much free space is being reported on the SAN. This discrepancy is not limited to arrays as all block-storage SANs that use thin provisioning have the same space discrepancy issue.

To work around this space discrepancy issue, Windows, VMware, and Linux file systems have implemented a feature that notifies the SAN to free up blocks that are no longer in use by the host file system. This feature is called block unmap or SCSI unmap.

Example:

Suppose 4 TB of data are written onto a volume mounted on a Windows 2008 R2 NTFS host, and then 2 TB of data are deleted. When files are deleted, the data blocks remain in place and the file system headers are updated with info that the blocks are not in use and are available; however, the array continues to see the blocks in use as data is still physically present on the volume. So, unless the host OS supports SCSI unmap to inform the underlying storage target of the freed up space on the file system, the storage continues to report that the data is still in use. In most cases, this is not a problem because the host file system will eventually reuse the deleted blocks for new data and the underlying storage will not report an increase in utilized space; however, this could become a problem if the array is becoming full and the space is needed for other volumes / snapshots.

Note: For these features to work optimally, you must be running array OS version 1.4.7.0 or later. SCSI unmap is supported in array OS 1.4.4.0 and later releases.

File systems support two methods for informing storage of vacated free blocks on the file system: online (periodic discards) and batched discards. In both methods, unused blocks are returned to the underlying SAN storage by overwriting the unused file system blocks with zeroes.

Arrays support "zero page detection" to reclaim previously provisioned space. The array will detect the zeroes written by the host file system and reduce the reported storage space used on the fly.

Periodic (online) discards happen automatically; that is, no scheduled run is required to reclaim unused space on the array. On the other hand, Batched discards require that a user manually run a tool or command to reclaim unused space.

Examples of file systems that support online discards:

- NTFS (Windows Server 2012, array OS 1.4.4.0 and higher)
- VMFS6 (VMware ESXi 6.5 and higher)
- VMFS5 (VMware ESXi 5.0, removed in 5.0 Update 1 due to performance overhead)
- ext4 (version 2.6.27-5185-g8a0aba7 onwards)

File systems that support batched discards:

- NTFS (Windows Server 2003-2008 R2, through use of "sdelete -z" utility)
- VMFS5 (VMware ESXi 5.0, Update 1 onwards, through use of "vmkfstools" utility)
- ext4 (Linux / v2.6.36-rc6-35-g7360d17)
- ext3 (Linux / v2.6.37-11-g9c52749)
- xfs (Linux / v2.6.37-rc4-63-ga46db60)





Important: Reclaiming free space under any operating system is write I/O intensive as the host OS writes zeroes to the unused blocks. Consequently, batched discard runs should be scheduled to run when writes to the host volumes are minimal.

Reclaiming Space on Linux

Before you can reclaim space on Linux, the Linux kernel running on the host must be one of the following at minimum:

- ext4 (Linux / v2.6.36-rc6-35-g7360d17)
- ext3 (Linux / v2.6.37-11-g9c52749)
- xfs (Linux / v2.6.37-rc4-63-ga46db60)

Online and automatic discards can be enabled for ext4 partitions using the **discard mount** option.

```
mount -o discard /dev/sdc /mnt/test
```

The discards are disabled by default under Linux/ext4 due to the additional I/O overheads that are incurred. Refer to the man pages of your Linux distribution for further details.

You can invoke manual block reclamation (batched discards) using the `fstrim` and `dd` utilities (included with most Linux distributions in the `util-linux-ng` package) if the volume is already mounted without discard options. When using the `dd` utility, run the following command from the volume file system folder:

```
dd if=/dev/zero of=/mnt/path_to_nimble_volume/balloonfile bs=1M count=XYZ
```

where `XYZ` is derived from the free space reported on the host.

For example, if the file system is showing 195GB used out of 1TB, then the balloon file size needs to be $1\text{TB} - 195\text{GB} = 805\text{GB}$, hence $XYZ = (805 * 1024 = 824320)$ since it is in units of megabytes (because `bs` is 1M). The `dd` command will create a file called `balloonfile` with zeros in 1 MB blocks with the remaining free space of the volume. Blocks that consist of all zeros become free blocks. After the file `dd` command has completed, the balloon file that was created can be deleted.

Note: If the volume on the array is over provisioned, do not attempt to manually reclaim the total leftover space of the volume, but rather reclaim up to the estimated compressed volume size, and before reaching 85% of possible used space on the array.



Host Adjustments for Multiple Initiators Using a Single Subnet

This chapter explains the potential changes which may need to be made to the Linux operating system ARP function in order to correctly connect iSCSI MPIO configurations when using direct attached volumes from the HPE Nimble Array with Linux host. Changes can be made at the system level for all Ethernet adapters, or at the individual adapter level. These changes are most commonly used with a single data subnet and multiple iSCSI adapters for the host.

Determine the kernel and release version of the host

The method to use when changing the Linux operating system ARP function depends on the kernel and release version of the host.

Procedure

1. For releases prior to 6.x, use the following commands:

uname -a to determine the kernel version:

```
[root@example ~]# uname -a
Linux example.world.org 4.18.0-80.el8.x86_64 #1 SMP Tue Jun 4 09:19:46 UTC
2019 x86_64 x86_64 x86_64 GNU/Linux
```

cat /etc/*release to determine the release version:

```
[root@example ~]# cat /etc/*release
CentOS Linux release 8.0.1905 (Core)
NAME="CentOS Linux"
VERSION="8 (Core)"
ID="centos"
ID_LIKE="rhel fedora"
VERSION_ID="8"
PLATFORM_ID="platform:el8"
PRETTY_NAME="CentOS Linux 8 (Core)"
ANSI_COLOR="0;31"
CPE_NAME="cpe:/o:centos:centos:8"
HOME_URL="https://www.centos.org/"
BUG_REPORT_URL="https://bugs.centos.org/"
CENTOS_MANTISBT_PROJECT="CentOS-8"
CENTOS_MANTISBT_PROJECT_VERSION="8"
REDHAT_SUPPORT_PRODUCT="centos"
REDHAT_SUPPORT_PRODUCT_VERSION="8"
CentOS Linux release 8.0.1905 (Core)
CentOS Linux release 8.0.1905 (Core)
```

2. For release 6.x and later, use the **hostnamectl** command to determine both the kernel and release versions.

```
[root@example ~]# hostnamectl
Static hostname: example.world.org
Icon name: computer-server
Chassis: server
Machine ID: 4e54f4ef3ase35fef546af6d13f1ef31e5
Boot ID: 7a466435e0c056df4sd54f3d1f13ew3
Operating System: CentOS Linux 8 (Core)
CPE OS Name: cpe:/o:centos:centos:8
Kernel: Linux 4.18.0-80.el8.x86_64
Architecture: x86-64
```



What to do next

Based on the results, complete the appropriate procedure later in this chapter:

- [ARP adjustments for RHEL 6.x - 8.x, OEL 6.x - 8.x, CentOS 6.x - 8.x](#) on page 78
- [ARP adjustments for Debian/Ubuntu 18.04 LTS Kernel 3.x - 4.x](#) on page 78

For more information on the parameters used in these procedures, see [Linux kernel parameter definitions](#) on page 78.

ARP adjustments for RHEL 6.x - 8.x, OEL 6.x - 8.x, CentOS 6.x - 8.x**Procedure**

1. To make adjustments at the system-wide level, add the following to `/etc/sysctl.conf`:

```
net.ipv4.conf.all.arp_ignore=1
net.ipv4.conf.all.arp_announce=2
net.ipv4.conf.all.rp_filter=2
```

To make adjustments at the individual adapter level, add the following to `/etc/sysctl.conf` for each interface used for iSCSI traffic:

```
net.ipv4.conf.ethX.arp_ignore=1
net.ipv4.conf.ethX.arp_announce=2
net.ipv4.conf.ethX.rp_filter=2
```

2. Restart `sysctl`:
`sysctl -p`

ARP adjustments for Debian/Ubuntu 18.04 LTS Kernel 3.x - 4.x**Procedure**

1. Make the following change to `sysctl.conf`:

At the system-wide level, add the following:

```
net.ipv4.conf.all.rp_filter=2
```

At the individual adapter level, add the following:

```
net.ipv4.conf.ethX.rp_filter=2
```

2. To determine if the filtering is occurring, add the following to `sysctl.conf`:

```
sysctl -w net.ipv4.conf.all.log_martians=1
```

And, then review `/var/log/messages` or the kernel `syslog`.

Linux kernel parameter definitions**arp_ignore**

(Integer) Defines different modes for replying to received ARP requests that resolve local target IP addresses:

0 - (Default): Reply for any local target IP address, configured on any interface.



- 1 - Reply only if the target IP address is a local address configured on the incoming interface.
- 2 - Reply only if the target IP address is a local address configured on the incoming interface and both with the sender's IP address are part from same subnet on this interface.
- 3 - Do not reply for local addresses configured with the scope host, only resolutions for global and link addresses are replied.
- 4-7 - Reserved.
- 8 - Do not reply for all local addresses.

```
net.ipv4.conf.eth0.arp_ignore=1
net.ipv4.conf.eth1.arp_ignore=1
```

arp_filter

(Integer) Enables or disables the restriction to allow or disallow an interface on subnet "A" to answer for all interfaces on subnet "A."

Enabling `arp_filter` disables the ability to allow a single interface to answer for multiple interfaces on the same subnet.

```
net.ipv4.conf.eth0.arp_filter=1
net.ipv4.conf.eth1.arp_filter=1
```

arp_announce

(Integer) Defines different restriction levels for announcing the local source IP address from IP packets in ARP requests sent on the interface:

0 - (default) Use any local address, configured on any interface

1 - Try to avoid local addresses that are not in the target's subnet for this interface. This mode is useful when target hosts reachable via this interface require the source IP address in ARP requests to be part of their logical network configured on the receiving interface. When generating the request, check all subnets that include the target IP and preserve the source address if it is from such a subnet. If there is no such subnet, select the source address according to the rules for level 2.

2 - Always use the best local address for this target. In this mode, ignore the source address in the IP packet and try to select the local address preferred for communication with the target host. The local address is selected by looking for the primary IP addresses on all subnets on the outgoing interface that include the target IP address. If no suitable local address is found, select the first local address on the outgoing interface or on all other interfaces, with the expectation to receive a reply for the request and even sometimes no matter the source IP address announced.

```
net.ipv4.conf.eth0.arp_announce=1
net.ipv4.conf.eth1.arp_announce=1
```

rp_filter

INTEGER defines different modes for packet filtering, and returns to the destination interface.

0 - No source validation done against incoming packets.

1 - Strict mode as defined in RFC3704 Strict Reverse Path. Each incoming packet is tested against the FIB and if the interface is not the best reverse path, the packet check will fail. By default, failed packets are discarded.

2 - Loose mode (default) as defined in RFC3704 Loose Reverse Path. Each incoming packet's source address is also tested against the FIB and if the source address is not reachable via any interface, the packet check will fail.

```
net.ipv4.conf.eth0.rp_filter=2
net.ipv4.conf.eth1.rp_filter=2
```

Helpful Information

This section contains information about how to perform certain functions from the storage array that can be helpful in working with Linux and the array OS.

Alignment Requirements for Linux

For additional information about alignment requirements, see *KB-000010 Proper Block Alignment*. At the time this document was published, the KB was located at https://infosight.hpe.com/InfoSight/media/kb/active/sup_KB-000010_Proper_Block_Alignment.pdf.

