



Technical Report

# NetApp SnapDrive 5.3 for UNIX Best Practices Guide

Ebin Kadavy, Jeffrey Steiner, Antonio Jose Rodrigues Neto,  
Anand Ranganathan, NetApp  
October 2015 | TR-4212

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Intended Audience	4
1.2	Prerequisites	4
<b>2</b>	<b>SnapDrive for UNIX Overview</b>	<b>4</b>
2.1	Key Features in SnapDrive 5.3	5
2.2	SnapDrive for UNIX Deployments	5
<b>3</b>	<b>Limitations of SnapDrive for UNIX</b>	<b>6</b>
3.1	General Limitations of SnapDrive for UNIX	6
3.2	Limitations of SnapDrive for UNIX on Linux	6
3.3	Limitations of SnapDrive for UNIX on Solaris	6
3.4	Limitations of SnapDrive for UNIX on AIX	7
<b>4</b>	<b>Installation and Basic Configuration</b>	<b>7</b>
4.1	Checklist for Storage System	7
4.2	Checklist for Host Before SnapDrive Installation	8
4.3	snapdrive.conf File	10
4.4	Using the SnapDrive Daemon	11
4.5	Setting Up SnapDrive to Communicate with a NetApp Storage System	11
4.6	Volume Managers and File Systems	12
4.7	Multipath	12
4.8	Multisubnet Configurations	14
4.9	Configuring Data Protection for Clustered Data ONTAP	15
<b>5</b>	<b>Security Features</b>	<b>15</b>
5.1	Best Practices	15
5.2	Configuring SnapDrive for RBAC	16
<b>6</b>	<b>Storage Provisioning</b>	<b>17</b>
<b>7</b>	<b>Snapshot Management</b>	<b>18</b>
7.1	Best Practices for Snapshot Copy Management	18
7.2	Consistent Snapshot Copies	18
7.3	Snapshot Space Management	20
7.4	Snap Reserve	20
<b>8</b>	<b>Restoring a Snapshot Copy Using Volume-Based SnapRestore</b>	<b>20</b>
8.1	Using Volume-Based SnapRestore	21

<b>9 Cloning .....</b>	<b>22</b>
9.1 Benefits of FlexClone Technology .....	22
9.2 Best Practices for Cloning.....	23
<b>10 SnapDrive for UNIX with SnapMirror Load-Sharing Mirror Relationships .....</b>	<b>24</b>
<b>11 Selective LUN Map in SnapDrive for UNIX.....</b>	<b>25</b>
<b>12 SnapDrive for UNIX in a 7-Mode MultiStore Environment.....</b>	<b>25</b>
<b>13 Platform and Protocols .....</b>	<b>25</b>
13.1 Virtualization Environment .....	25
13.2 ALUA Support.....	26
13.3 RDM LUN Support.....	26
<b>14 Summary .....</b>	<b>30</b>
<b>References.....</b>	<b>30</b>
<b>Version History .....</b>	<b>30</b>

**LIST OF TABLES**

Table 1) Storage system checklist.....	7
Table 2) Manual preinstallation checklist.....	9
Table 3) SnapDrive daemon operations.....	11
Table 4) Aliased device support for SnapDrive operations.....	13
Table 5) Predefined roles on Operations Manager server.....	17
Table 6) FlexClone configuration values.....	22

# 1 Introduction

This document is a best practices guide for NetApp® storage systems that use NetApp SnapDrive® for UNIX. It also provides recommendations on various configuration options available, including providing specific information on when and where to use the options. This document is applicable for NetApp Data ONTAP® 7-Mode and clustered Data ONTAP environments. The primary differences between Data ONTAP 7-Mode and clustered Data ONTAP are clearly identified and are related to secondary protection using SnapMirror or SnapVault.

## 1.1 Intended Audience

This document provides guidelines for deploying and using SnapDrive for UNIX (SDU) with NetApp storage appliances and supporting software.

It is intended for storage and server administrators and experts who manage storage provisioning and NetApp Snapshot® copies on NetApp storage systems that use SnapDrive for UNIX as a solution.

## 1.2 Prerequisites

NetApp recommends that you refer to the following guides before reading this technical report:

- [SnapDrive 5.3 for UNIX Installation and Setup Guide for Linux for Clustered Data ONTAP](#)
- [SnapDrive 5.3 for UNIX Installation and Setup Guide for Linux for Data ONTAP 7-Mode](#)
- [SnapDrive 5.3 for UNIX Release Notes \(AIX, Linux, and Solaris\)](#)
- [SnapDrive 5.3 for UNIX Administration Guide](#)
- [Clustered Data ONTAP 8.3 Documentation](#)
- [Linux Host Utilities 6.2 Installation and Setup Guide](#)

Readers should have a good understanding of UNIX file system administration as well as the NFS, FCP and iSCSI protocols.

# 2 SnapDrive for UNIX Overview

SnapDrive for UNIX is an enterprise-class storage and data management utility that simplifies storage management and increases the availability and reliability of application data. Its key functionality includes error-free application storage provisioning, consistent-data NetApp Snapshot copies, and rapid application recovery. It also provides the ability to easily manage data that resides on NetApp Network File System (NFS) shares or NetApp LUNs. SnapDrive for UNIX complements the native file system and volume manager and integrates seamlessly with the clustering technology supported by the host operating system (OS).

SnapDrive for UNIX is supported on the following platforms.

- Linux:
  - Red Hat Enterprise Linux
  - SUSE Linux Enterprise Server
  - Oracle Enterprise Linux and Oracle Linux
- Oracle Solaris
- IBM AIX

**Note:** For the latest information about supported OS platforms and NetApp Data ONTAP versions, refer to the [Interoperability Matrix Tool](#).

## 2.1 Key Features in SnapDrive 5.3

SnapDrive 5.3 for UNIX includes the following key features:

- **Support of alias names in Linux MPIO environment.** An alias name is a globally unique name provided by an administrator for a multipath device. Alias names override the WWID. SnapDrive supports such aliased devices in a Linux multipath environment where enterprise users can set alias names for the multipath devices after creation of the storage.
- **Role-based access control (RBAC) permissions** allow administrators to restrict access to the storage system for various SnapDrive operations. The access is usually based on the role that is assigned for the user. This feature is supported through OnCommand® Unified Manager.
- **Secondary data protection.** With clustered Data ONTAP, SnapDrive can natively replicate Snapshot copies through SnapMirror or SnapVault using protection policies defined at the Data ONTAP level. With Data ONTAP 7-Mode, use OnCommand Unified Manager 5.2.x to enable the policy-based protection mechanism.
- **64-bit binaries and RHEL 7x support.** SnapDrive 5.3 supports 64-bit OS binaries and is qualified to work on latest version of the RHEL operating system.
- **Enhanced ASUP™ logging.** The event management system is a structured logging system that is used to record and report various events. The latest version has an improved logging system to track error messages.

## Existing Features in SnapDrive

The following features are already available and included in the latest version:

- The consistency group Snapshot capability of Data ONTAP 8.3 is leveraged to provide support for SAN protocols, with both clustered Data ONTAP and Data ONTAP 7-Mode.
- Support of clustered Data ONTAP is provided for both NFS and SAN.
- Support is provided for NFS operations on a volume that contains nested junction paths that are in clustered Data ONTAP.
- Cisco converged network adapters (CNAs) are supported.
- The Mount Guard feature of AIX is supported.
- VMware paravirtual SCSI-controlled devices are supported on Linux guest OSs.
- Fixes are provided for multiple BURT. For more information, refer to the [SnapDrive 5.3 for UNIX Release Notes](#).
- The load-sharing mirror (LSM) capability is supported.
- The selective LUN mapping (SLM) feature is supported from clustered Data ONTAP 8.3 onwards.
- The IPspaces for SVM feature is supported. An IPspace defines a distinct IP address space in which storage virtual machines (SVMs, formerly called Vservers) can be created. A routing table is maintained for each SVM in an IPspace; no cross-SVM or cross-IPspace traffic routing occurs. For more information, refer to the [Clustered Data ONTAP Network Management Guide](#).
- NetApp MetroCluster™ with clustered Data ONTAP is supported. For information about how to set up MetroCluster, refer to the [Clustered Data ONTAP MetroCluster Installation and Configuration Guide](#).

## 2.2 SnapDrive for UNIX Deployments

SnapDrive for UNIX can be used either as a standalone product or as part of other NetApp solutions. For example, it can be deployed along with NetApp SnapManager® for Oracle. In both types of deployment, SnapDrive for UNIX serves as a tool to create and manage storage. It also manages storage backups and restores storage from those backups through Snapshot technology.

For a complete list of supported platforms, refer to the NetApp [Interoperability Matrix Tool](#).

## 3 Limitations of SnapDrive for UNIX

This section identifies limitations of which to be aware when using SnapDrive for UNIX.

### 3.1 General Limitations of SnapDrive for UNIX

SnapDrive for UNIX has the following limitations:

- It does not support snap operations on a NFS mount point when the volume is exported with a Kerberos authentication security type, such as krb5, krb5i, or krb5p.
- In clustered Data ONTAP, SnapDrive for UNIX does not support symbolic links for mount points.
- Snapshot operations might be inconsistent if NetApp SnapRestore® is performed on a mount point where an entity is mounted that differs from the one created by the Snapshot copy.
- SnapDrive modifies the mount-point permissions from nonroot user to root user for a qtree after volume-based SnapRestore operations.
- Beginning with SnapDrive for UNIX 5.1, only asynchronous mode is supported for clone-split operations.
- The `split-clone-async` configuration variable and `-splitmode` command-line options have been removed.
- The clone-split operation is not supported on a filespec, which is a combination of volume clone and LUN clone, LUN clone and LUN, or volume clone and LUN.
- Restoring the storage stack on the host fails if you have removed a LUN by using an external command.
- SnapDrive for UNIX does not allow you to restore a Snapshot copy of a qtree if the Snapshot copy is mounted on an NFS in the SnapMirror destination.
- SnapDrive for UNIX does not support Broadcom CNAs.

### 3.2 Limitations of SnapDrive for UNIX on Linux

SnapDrive for UNIX has the following limitations in Linux environments:

- When using a clustered Data ONTAP SVM with SnapDrive for UNIX, make sure that the IP address of SVM's management logical interface (LIF) is resolvable. You must also make sure that the data SVM name is configured in SnapDrive by using the `snapdrive config set <vsadmin> <Vserver name>` command.
- In clustered Data ONTAP, aggregates that contain data SVM volumes must be assigned to the data SVM's aggregate list to configure the data SVM and perform SnapDrive operations.
  - a. Check the list of aggregates assigned to the SVM by entering the following command:

```
vserver show -fields aggr-list
```

- b. Assign one or more aggregates to the SVM by entering the using command:

```
vserver modify -vserver vserver_name -aggr-list aggr_name
```

- In clustered Data ONTAP, SnapDrive for UNIX does not support the `snap connect` operation with the `-readonly` option from the secondary host (the host that does not have export permissions on the parent volume).
- In clustered Data ONTAP, the export policy is defined only at the volume level. Therefore, the `.snapshot` directory cannot be exported to the secondary host because the secondary host does not have export permission on the parent volume.

### 3.3 Limitations of SnapDrive for UNIX on Solaris

SnapDrive for UNIX has the following limitations in Solaris environments:

- LUNs greater than 1TB with a SMI label are not supported.
- Using EFI and SMI LUNs in the same volume group is not supported.

### 3.4 Limitations of SnapDrive for UNIX on AIX

SnapDrive for UNIX has the following limitations in IBM AIX environments:

- The SnapDrive daemon will be very slow to start if one or more Fibre Channel (FC) port links are down or not available. If this becomes a problem, the devices can be removed with the `rmdev` command.
- By default, SnapDrive for UNIX creates disk/volume groups with nonconcurrent settings, and storage provisioning operations must be changed or provided manually.

## 4 Installation and Basic Configuration

You can download SnapDrive for UNIX software from the NetApp [Software Downloads](#) page. Make sure that the required UNIX platform is selected, and then read the corresponding description and download pages. Before installing SnapDrive for UNIX, use the checklists in the following sections to avoid potential errors or delays during or after the installation.

### 4.1 Checklist for Storage System

Table 1 provides the checklist for configuring the storage system.

Table 1) Storage system checklist.

Step	Action
1.	<p>Verify the licenses on the storage system:</p> <ul style="list-style-type: none"> <li>• FCP, iSCSI, or NFS license, depending on your configuration</li> <li>• NetApp FlexClone® technology</li> <li>• SnapRestore software</li> <li>• NetApp MultiStore® software (Data ONTAP 7-Mode vFile® environments only)</li> </ul> <p>In Data ONTAP 7-Mode, use the <code>license</code> command on the command-line interface (CLI) to verify installed licenses.</p> <p>In clustered Data ONTAP 8.3, run the <code>license show</code> command on the CLI.</p>
2.	<p>Enable, configure, and test RSH or SSH access on the storage systems for administrative access. NetApp recommends using SSH because it is more secure than telnet is. For detailed information to help you complete these tasks, refer to the <a href="#">Clustered Data ONTAP 8.3 Logical Storage Management Guide</a>.</p>
3.	<p>For an NFS environment:</p> <ul style="list-style-type: none"> <li>• Make sure that storage system directories are exported correctly to the host.</li> <li>• Make sure that all the host interfaces to storage have read-write permission to access the directory or at least read-only permission for <code>snapdrive snap connect</code> with the <code>-readonly</code> option.</li> </ul> <p>For information about how to manage the NFS protocol, refer to the <a href="#">Clustered Data ONTAP 8.3 File Access and Management Guide for NFS</a>.</p>

Step	Action
4.	<p>For an FC or iSCSI environment, verify whether FCP or iSCSI is enabled on the storage system:</p> <ul style="list-style-type: none"> <li>• In Data ONTAP 7-Mode, run <code>fcv status</code> or <code>iscsi status</code>.</li> <li>• If the status is disabled, start the service by running <code>fcv start</code> or <code>iscsi start</code>.</li> <li>• In clustered Data ONTAP 8.3, run <code>Vserver fcp</code> or <code>Vserver iscsi show</code>.</li> <li>• If no configuration exists, create the configuration for SVM by running:  <code>Vserver fcp create -Vserver &lt;Vserver&gt;</code> or <code>Vserver iscsi create -Vserver &lt;Vserver&gt;</code></li> </ul>
5.	<p>Note the storage system target address by running the following commands on the storage system CLI:</p> <ul style="list-style-type: none"> <li>• In Data ONTAP 7-Mode, run:  <code>fcv nodename</code> or <code>iscsi nodename</code></li> <li>• In clustered Data ONTAP 8.3, run:  <code>vserver fcp show -vserver &lt;vserver&gt;</code> or <code>vserver iscsi show -vserver &lt;vserver&gt;</code></li> </ul>
6.	<p>Verify that the FC port on the NetApp storage system is configured as the target by running the following command.</p> <ul style="list-style-type: none"> <li>• In Data ONTAP 7-Mode, run the <code>fcadmin config</code> command.</li> <li>• For clustered Data ONTAP, refer to the <a href="#">Clustered Data ONTAP 8.3 Logical Storage Management Guide</a>.</li> </ul>

### Best Practice

By default, the iSCSI service is enabled on all of the Ethernet interfaces after you enable the license. Do not use a 10/100Mb Ethernet interface for iSCSI communication. The e0 management interface on many storage systems is a 10/100 interface. So that all iSCSI commands are processed by the appropriate interfaces, disable iSCSI processing on a particular Ethernet interface with the following command:

```
iscsi interface disable <interface_name>
```

#### Examples:

```
iscsi interface disable e0b
```

In clustered Data ONTAP, a particular data LIF can be disabled for an iSCSI service by using the following command:

```
vserver iscsi interface disable -vserver <vserver> -lif <data_lif>
```

**Note:** Do not run the preceding commands while active iSCSI sessions are connected to the Ethernet or to a LIF. First disconnect the active sessions from the host; otherwise, the storage system displays a warning message when the command is issued.

## 4.2 Checklist for Host Before SnapDrive Installation

The following sections describe how to check the hosts prior to installing SnapDrive either by using the SnapDrive configuration checker or by checking manually.

### Preinstallation Checks with SnapDrive Configuration Checker

Before using or installing SnapDrive for UNIX, make sure that the host supports all of the dependent components, such as the native OS version, multipath solution, cluster solution, and so on. To verify the

support, refer to product support matrixes and product documentation. As SnapDrive for UNIX expands its supported platforms, the combination of supported configurations grows exponentially, and it becomes a major task for administrators to validate the environment. Administrators might not know which components to check or how to get their versions released, and these instances often end up as support cases.

The SnapDrive configuration checker addresses the manual reference task by allowing the user to check whether the host has all of the supported components of SnapDrive that are required for UNIX to run. The configuration checker is bundled with SnapDrive for UNIX software; however, the configuration checker software can also be used independent of SnapDrive, and it can be downloaded from the [Utility Toolchest](#) on the [NetApp Support site](#).

In the case of a new installation, NetApp recommends running the configuration checker before installing SnapDrive for UNIX to confirm that all necessary components and correct versions are available. The configuration checker validates the setup configuration against the NetApp [Interoperability Matrix Tool](#):

```
Host#/opt/NetApp/snapdrive/bin/sdconfcheck check
```

The administrator can also use an existing `snapdrive.conf` file from other hosts and validate that the new host is ready with the current working configuration. This is helpful in a disaster recovery (DR) scenario in which the host on the DR site can be verified with the primary site configuration file and it is known that both are identical:

```
DR-Host#/opt/NetApp/snapdrive/bin/sdconfcheck check -conf /snapdrive.conf
```

If SnapDrive for UNIX is already installed, and the addition of a new configuration (for example, MPIO for high availability) is planned, the configuration checker helps to validate whether SnapDrive for UNIX will support the new environment:

```
Host#/opt/NetApp/snapdrive/bin/sdconfcheck check -mptype nativempio
```

## Manual Preinstallation Steps

Before installing SnapDrive for UNIX, it is a best practice to check for the prerequisites. Refer to the [Interoperability Matrix Tool](#) to confirm that SnapDrive for UNIX supports your environment.

For specific information about requirements, such as patches that are required for the OS, refer to the OS vendor documents.

Table 2 provides the manual preinstallation checklist.

**Table 2) Manual preinstallation checklist.**

Step	Action
1.	Identify whether the host OS version is supported by SnapDrive for UNIX: <ul style="list-style-type: none"> <li>Linux: <code>cat /etc/issue</code></li> </ul>
2.	<p>For FCP environments:</p> <ul style="list-style-type: none"> <li>Confirm that you are using the SnapDrive for UNIX supported Host Utility version from the <a href="#">Interoperability Matrix Tool</a>.</li> <li>Verify that NetApp Host Utilities are present on the host by running the <code>sanlun</code> version.</li> <li>If Host Utilities are not installed, download and install the correct version from the <a href="#">NetApp Support</a> site and refer to the <a href="#">Host Utilities Installation and Setup Guide</a> to install (or upgrade) and configure Host Utilities.</li> </ul> <p>Identify the host bus adapters on the host and verify that the ports are enabled by running <code>sanlun fcp show adapter -v</code>.</p> <p>For Veritas environments:</p> <ul style="list-style-type: none"> <li>Confirm that you are using Veritas Storage Foundation for multipathing, installing, and</li> </ul>

Step	Action
	configuring the Symantec Array Support Library (ASL) for NetApp storage systems. <ul style="list-style-type: none"> <li>• Determine the version of the Veritas Storage Foundation and ASL to use.</li> <li>• For more information about obtaining and installing ASL, refer to the <a href="#">Linux Host Utilities 6.2 Installation and Setup Guide</a>.</li> </ul> For iSCSI environments: <ul style="list-style-type: none"> <li>• Confirm that you are using the SnapDrive for UNIX supported Host Utility version from the <a href="#">Interoperability Matrix Tool</a>.</li> <li>• Verify that the NetApp Host Utilities are present on the host by running the <code>sanlun</code> command.</li> <li>• If Host Utilities are not installed, download and install the correct version from the <a href="#">NetApp Support</a> site and refer to <a href="#">Host Utilities Installation and Setup Guide</a> to install (or upgrade) and configure Host Utilities.</li> </ul>

## Best Practices

- Download the latest package of the configuration checker tool from the [Utility Toolchest](#) on the [NetApp Support](#) site.
- Execute the configuration checker both before and after installing the SnapDrive for UNIX tool to validate and confirm the presence of dependent components.
- NetApp recommends having a minimum of 1GB to 2GB of spare space for SnapDrive for UNIX operation in the root directory (`/`) to prevent the root volume from running out of space. SnapDrive for UNIX maintains audit, trace, and recovery log files whose default locations are specified in the `snapdrive.conf` file. These log files can be placed into a different directory, as specified in the `snapdrive.conf` file. SnapDrive for UNIX rotates audit and trace files when they reach maximum size, as defined in the `snapdrive.conf` file. A recovery file rotates only after an operation is completed; therefore, no default size is allotted for this file.
- After the preceding checklist has been verified, follow the steps in the [SnapDrive 5.3 for UNIX Installation and Setup Guide for Linux for Clustered Data ONTAP](#) to install SnapDrive for UNIX.
- NetApp recommends using the default path for the SnapDrive for UNIX installation and adding the SnapDrive for UNIX directory to your `$PATH` environment variable.

### 4.3 snapdrive.conf File

The `snapdrive.conf` configuration file in SnapDrive for UNIX holds the default values of all configurable variables. The entries in the `snapdrive.conf` file can be manually modified by using a text editor. The modified values take effect after the SnapDrive daemon has been restarted.

To change the entries in `snapdrive.conf`, complete the following steps:

1. Copy the line that is commented (`#`) out to another line.
2. Modify the new copied line to remove only the first `#` and to set the new value for the variable.
3. Save the file and exit.
4. Restart the SnapDrive daemon.
5. Run the `snapdrive config show` command to check the current values of the configurable variables.

#### Example:

To set the `autosupport-enabled` variable to `off` and `contact-http-port` to `80` in the `snapdrive.conf` file, use this valid format:

```
autosupport-enabled="off" #Enable autosupport (requires autosupport-filer be set)
```

```
contact-http-port=80 #HTTP port to contact to access the filer
```

## 4.4 Using the SnapDrive Daemon

Most NetApp SnapManager products, such as SnapManager for Oracle and SnapManager for SAP, leverage SnapDrive for UNIX to create application-consistent backups and perform fast restores and quick clones. SnapDrive for UNIX provides a web service with a uniform interface for all NetApp SnapManager products to integrate seamlessly with SnapDrive for UNIX by using application programming interfaces (APIs). While the SnapDrive for UNIX daemon is in use, all SnapDrive commands work as a unique process in the background.

### Common Daemon Operations

For any SnapDrive for UNIX command to work, the SnapDrive daemon must be running.

Table 3 lists some of the most common daemon operations.

Table 3) SnapDrive daemon operations.

Daemon Operation	Command	OS User
Starting or restarting the daemon	<code>snapdrived start/restart</code>	Root
Verifying the status of the daemon	<code>snapdrived status</code>	Root
Stopping the daemon	<code>snapdrived stop</code>	Root
Changing the default daemon password	<code>snapdrived passwd</code>	Root

## 4.5 Setting Up SnapDrive to Communicate with a NetApp Storage System

To enable SnapDrive to communicate with NetApp storage systems, complete the following steps:

1. Verify the SnapDrive for UNIX version:

```
snapdrive version
```

2. Check the connectivity from the host to the storage system:

```
ping <storage_system_Ipaddress / storage_system_name>
```

**Note:** Make sure that the correct storage system name is registered with the domain name server or is in the `/etc/hosts` file.

3. Run `snapdrive config` and verify that the Data ONTAP 7-Mode storage system is already configured. If it is not configured, run the following command:

```
snapdrive config set vsadmin <vserver_name>storage_system_IPaddress /storage_system_name >
```

For clustered Data ONTAP, run the command `snapdrive config set vsadmin <vserver>` before configuring the management interface. Make sure that the `vsadmin` user is unlocked and that the administration privilege is delegated to the `vsadmin` user. For more information, refer to the [Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators](#).

#### Example:

```
snapdrive config set vsadmin vs1
snapdrive config set -mgmtpath vs1 10.72.199.188
```

4. To access the storage system, when prompted, enter the password for `<userid>`.
5. Edit the `snapdrive.conf` file and verify that the default settings conform to your environment:
  - Transport-type setting:

```
default-transport="iscsi" | "fcp"
```

**Note:** If the NFS configuration alone is used with SnapDrive for UNIX, the default transport variable must be commented out.

- HTTPS setting:

```
use-https-to-filer="on" | "off"
```

**Note:** Starting with SnapDrive 3.0 for UNIX, the value for HTTP/HTTPS is `on` for Solaris, AIX, Red Hat, and SUSE platforms and `off` for the HP-UX platform.

## 4.6 Volume Managers and File Systems

SnapDrive for UNIX supports more than one volume manager and file system on the same host:

- Volume manager for FCP:
  - Solaris: `vmtype="vxvm" | "svm"`
  - AIX: `vmtype="vxvm" | "lvm"`
  - Linux: `vmtype="lvm"`
- Volume manager for iSCSI:
  - Solaris: `vmtype="svm"`
  - AIX: `vmtype="lvm"`
  - Linux: `vmtype="lvm"`
- File system for FCP:
  - Solaris: `fstype="vxfs" | "ufs"`
  - AIX: `fstype="jfs2" | "vxfs"`
  - Linux: `fstype="ext3" | "ext4"`
- File system for iSCSI:
  - Solaris: `fstype="ufs"`
  - AIX: `fstype="jfs2"`
  - Linux: `fstype="ext3" | "ext4"`

**Note:** SnapDrive for UNIX supports all commands on any AIX system with the JFS2 or VxFS file systems. However, it supports only creating, renaming, restoring, and deleting Snapshot copies for an existing storage device with the JFS file system.

## 4.7 Multipath

The following multipathing types are supported on different platforms:

- For FCP:
  - Solaris: `multipathing-type="DMP" | "mpxio" | "none"`
  - AIX: `multipathing-type="NativeMPIO" | "DMP" | "none"`
  - Linux: `multipathing-type="NativeMPIO" | "none"`
- For iSCSI:
  - Solaris: `multipathing-type="mpxio" | "none"`
  - AIX: `multipathing-type="NativeMPIO" | "none"`
  - Linux: `multipathing-type="NativeMPIO" | "none"`Note

**Note:** All values are case sensitive.

## Alias Names in Linux MPIO Environment

An alias name is a globally unique name that an administrator provides for a multipath device. Alias names override the WWID, and this can be specified by using the alias option in multipath section in /etc/multipath.conf. SnapDrive 5.3 for UNIX allows the use of alias names for multipath devices.

To create alias names, complete the following steps.

1. Make sure that the alias names are added in the multipath.conf file:

```
## Set alias name for the multipath devices in /etc/multipath.conf
[root@x336-206-32 linux]# cat /etc/multipath.conf
multipaths {
  multipath {
    wwid 3600a098054336650313f4662306d4a74
    alias FIRST_LUN
  }
  multipath {
    wwid 3600a098054336650313f4662306d4a75
    alias SECOND_LUN
  }
}
```

2. Restart the multipath daemon and SnapDrive to enable the changes:

```
[root@x336-206-32 linux]# /etc/init.d/multipathd restart
ok
Stopping multipathd daemon:                [ OK ]
Starting multipathd daemon:                 [ OK ]
## Restart the snapdrived daemon.
[root@x336-206-32 linux]# snapdrived restart
Successfully stopped daemon
Successfully started daemon
[root@x336-206-32 linux]#
## Check the alias names in corresponding device
```

3. Query the SnapDrive storage to check if the alias names are displayed for multipath devices:

```
[root@x336-206-32 linux]# snapdrive storage show -all
WARNING: This operation can take several minutes
based on the configuration.
dg: alias_lvm_SdDg                dgtype lvm
hostvol: /dev/mapper/alias_lvm_SdDg-alias_lvm_SdHv    state: AVAIL
fs: /dev/mapper/alias_lvm_SdDg-alias_lvm_SdHv    mount point: /mnt/alias_lvm (persistent)
fstype ext3
device filename                adapter path    size    proto    state    clone    lun path
backing snapshot
-----
-----
-----
sd_60_vs5:/vol/bril_3/lun12    -    P    1g    iscsi    online    No
sd_60_vs5:/vol/bril_1/lun21    -    P    1g    iscsi    online    No
[root@x336-206-32 linux]#
```

## Support for Aliased Devices

Table 4 lists the alias devices supported in SnapDrive 5.3 for UNIX.

Table 4) Aliased device support for SnapDrive operations.

Basic Operations	RAW LUN	LVM	NOLVM
snapdrive storage/host connect	Displays aliased devices	Displays aliased devices	Displays aliased devices
snapdrive snap restore	Displays aliased	Displays aliased	Displays aliased

Basic Operations	RAW LUN	LVM	NOLVM
(file or volume-based SnapRestore)	devices	devices	devices
snapdrive snap connect	Displays nonaliased devices	Displays nonaliased devices	Displays nonaliased devices
snapdrive snap create	Supported	Supported	Supported
snapdrive snap list/show -v	Does not show device path	Supported	Supported
snapdrive storage show -all	Supported	Supported	Does not show for "raw device" and "mount point"
snapdrive storage show -fs	Displays aliased devices	Displays aliased devices	Displays aliased devices

## 4.8 Multisubnet Configurations

To manage network traffic, SnapDrive for UNIX allows multiple subnet configurations to split up management and data traffic. SnapDrive for UNIX sends all management communication by using NetApp APIs over the management interface, and the data interface is used to exchange data between systems. The `snapdrive config set` command with the `-mgmt` option can be used to configure numerous subnets.

### Usage:

```
snapdrive config set -mgmtpath mgmt_interface data_interface [mgmt_interface data_interface ...]
```

In a pure SAN environment, the concept of a data path does not exist; therefore, you must configure the management interface and the data interface with the same IP address. In mixed SAN and NFS environments, interfaces are mapped so that one management interface and data interface pair must be the same as those in purely SAN environments.

You can view all of the data interfaces configured for a management interface by running the `snapdrive config list` command with the `-mgmtpath` option.

The following example configures 10.72.199.42 as the management interface and 10.73.68.142 as the data interface in an NFS environment:

```
snapdrive config set -mgmtpath 10.72.199.42 10.73.68.142
```

In a pure SAN environment, both management and data interfaces are set with the same IP address, 192.168.10.10:

```
snapdrive config set -mgmtpath 192.168.10.10 192.168.10.10
```

In a mixed SAN and NFS environment, the interfaces have different IP addresses:

```
snapdrive config set -mgmtpath 10.72.199.42 10.72.199.42 10.72.199.42 10.73.68.142
```

**Note:** The first two interfaces in the preceding command are for SAN configurations; therefore, they have the same value for the IP addresses. The third and the fourth interfaces are the management and data interfaces for NFS configurations.

## 4.9 Configuring Data Protection for Clustered Data ONTAP

To enable data protection to secondary storage, SnapDrive should be configured to access the secondary SVM. SnapManager for Oracle communicates to storage by using SnapDrive commands to vault or mirror the Snapshot copies.

To enable secondary protection in SnapDrive, do as follows:

1. Make sure that the storage system management LIF is resolvable.
2. Register the SVMs to the SnapDrive so that it's aware that the secondary storage has been enabled for protection.

```
snapdrive config set vsadmin <vserver>
```

## 5 Security Features

SnapDrive for UNIX provides the following basic levels of security:

- **Access control** allows you to specify the operations that a host running SnapDrive for UNIX can perform on a storage system. You must set the access control permissions individually for each host.
- **HTTPS** allows all interactions with the storage system and host through the Data ONTAP interface, including sending the passwords in a secure manner.

**Note:** For more information about how to set up and view access control permissions, refer to the [SnapDrive 5.3 for UNIX Administration Guide for Linux](#).

### 5.1 Best Practices

To provide added security, use HTTPS instead of HTTP for host-to-storage system communication. To use HTTPS, complete the following steps:

1. Enable HTTPS on the storage system side.
2. Set the configuration variables `contact-ssl-port=443` and `use-https-to-filer=on` in the `snapdrive.conf` file.

**Note:** OnCommand Unified Manager for RBAC is however supported for clustered Data ONTAP and 7-mode. Check IMT for exact versions.

In an environment in which OnCommand Unified Manager is not used for RBAC, complete the following steps:

1. Create the directory `sdprbac` (SnapDrive permissions for RBAC) under the root volume of the storage system.
2. Create a `prbac` file under the `sdprbac` directory. The file name is `/vol/vol10/sdprbac/sd<hostname>.prbac`, in which the host name is the name of the SnapDrive host identified by the `hostname` command.
3. Specify only one level in the `sd<hostname>.prbac` file from among the following list of permissions:
  - **NONE:** The host has no access to the storage system.
  - **SNAP CREATE:** The host can create Snapshot copies.
  - **SNAP USE:** The host can delete and rename Snapshot copies.
  - **SNAP ALL:** The host can create, restore, delete, and rename Snapshot copies.
  - **STORAGE CREATE DELETE:** The host can create, resize, and delete storage.
  - **STORAGE USE:** The host can connect and disconnect storage.
  - **STORAGE ALL:** The host can create, delete, connect, and disconnect storage.

- ALL ACCESS: The host has access to all the preceding SnapDrive for UNIX operations.

**Note:** The permission-level string must be on the first line of the permission file.

4. For security, set the value of the `all-access-if-rbac-unspecified` option to `off` in the `snapdrive.conf` file.
5. Restart the SnapDrive daemon.
6. If the permission file is not on the storage system, SnapDrive for UNIX checks the variable `all-access-if-rbac-unspecified` in the `snapdrive.conf` file. If this variable is set to `on`, the host has access to all SnapDrive commands.
7. Run the `snapdrive config access show <NetApp Controller Name>` command to verify access permission.

After you configure and verify the preceding tasks based on your environment and needs, you can use the SnapDrive CLI commands to perform the following tasks:

- Create storage entities such as LUNs, volume groups, logical volumes, and file systems.
- Create Snapshot copies and restoration file systems, LUNs, volume groups, and NFS files.
- Connect and disconnect from the Snapshot copy on the host.

For more information about access privileges, refer to the [SnapDrive 5.3 for UNIX Administration Guide for Linux](#).

## 5.2 Configuring SnapDrive for RBAC

The RBAC feature is supported through OnCommand Unified Manager for both Data ONTAP 7-Mode and clustered Data ONTAP systems. OnCommand Unified Manager 5.2.1 RC1 alone supports the RBAC feature in clustered Data ONTAP. RBAC is implemented by using the Operations Manager infrastructure. Limited access control was observed in versions earlier than SnapDrive for UNIX 4.0, and only the root user could perform SnapDrive for UNIX operations. SnapDrive for UNIX 4.0 and later versions support the nonroot local user and NIS users through the RBAC infrastructure of Operations Manager. SnapDrive for UNIX does not require a root password for the storage system; it communicates with the storage system by using `sd-<hostname>` user.

The following tasks provide a high-level overview about how to set up RBAC:

1. The Operations Manager administrator creates a user, `sd-admin` user with the capability of core access check over the global group (`global DFM.Core.AccessCheck`). After the Operations Manager administrator configures the `sd-admin` user, the administrator must manually send the credential information to the SnapDrive for UNIX administrator.
2. The Operations Manager administrator can also grant global write privileges (`global DFM.Database.Write`) to `sd-admin` to enable SnapDrive for UNIX to refresh storage system entities on Operations Manager.
3. The Operations Manager administrator must create the `sd-<hostname>` user on the storage system.
4. The SnapDrive for UNIX administrator receives user credentials for `sd-admin` and `sd-<hostname>` from the Operations Manager administrator.
5. The UNIX administrator must turn on RBAC functionality by setting the variable `rbac-method=dfm` in the `snapdrive.conf` file and restarting the SnapDrive for UNIX daemon.
6. RBAC with Operations Manager is not supported for clustered Data ONTAP. The Operations Manager administrator must grant privileges to the invoker of SnapDrive to execute SnapDrive commands. For more information about the mapping of privileges compared with commands, refer to the [SnapDrive 5.3 for Unix Administration Guide](#).

SnapDrive for UNIX uses the following formats to verify whether a user is authorized to perform these tasks:

- For a NIS user running the `snapdrive` command, SnapDrive for UNIX uses the format `<nisdomain>\<username>`; for example, `netapp.com\marc`.
- For a local user of a UNIX host, such as `lnx197-141`, SnapDrive for UNIX uses the format `<hostname>\<username>`; for example, `lnx197-141\john`.
- For an administrator (`root`) of a UNIX host, SnapDrive for UNIX treats the administrator as a local user and uses the format `lnx197-141\root`.

Preconfigured roles simplify the task of assigning roles to users. Table 5 lists these roles on the Operations Manager server.

Table 5) Predefined roles on Operations Manager server.

Role Name	Description
GlobalSDStorage	Manages storage with SnapDrive for UNIX
GlobalSDConfig	Manages configurations with SnapDrive for UNIX
GlobalSDSnapshot	Manages Snapshot copies with SnapDrive for UNIX
GlobalSDFullControl	Fully uses SnapDrive for UNIX

## 6 Storage Provisioning

SnapDrive for UNIX allows you to create and delete LUNs on a storage system or connect to the storage that is already created by using command-line options from the host. If you need to perform these tasks without SnapDrive for UNIX, you must log in to the storage system to create and map a LUN, identify the LUN on the host through the host commands, create a file system, and mount the file system. SnapDrive for UNIX achieves all of these tasks with one command, reducing the time needed and the possible errors made during this process. SnapDrive for UNIX can create storage by using a minimum of options, but NetApp recommends that you understand the default values and use them appropriately.

SnapDrive for UNIX lets you access LUNs, file systems, volume groups, and logical volumes on a host other than the one used for creating them. This might be useful in scenarios in which the storage must be migrated to a new server or when you want to back up a storage entity from a different server to another medium. SnapDrive for UNIX does this by running the `snapdrive storage connect/disconnect` command.

### Examples:

The `server1` command removes LUN mapping from server 1 and exports the volume groups or file systems that the LUNs contain.

```
snapdrive storage disconnect -fs /mnt/test
```

The `server2` command maps and identifies the LUN, imports the underlying volume group, activates the volume group, and mounts the file system on a new server.

```
snapdrive storage connect -fs /mnt/test -lvol test_SdDg/test_SdHv -lun btc-ppc-158:/vol/src_vol/lun4
```

By using SnapDrive for UNIX, you can even export and import volume groups or file systems without unmapping the LUNs from the server by using the `snapdrive host disconnect/connect` command.

SnapDrive for UNIX also allows you to increase the size of your volume group by using the `snapdrive storage resize` command with the `-addlun` option. Use the `snapdrive storage resize`

command with the `-growby` option to increase the size of the volume group by adding a LUN of a specified size to the volume group:

```
snapdrive storage resize -vg vg1 -growby 10g -addlun
```

To avoid space contentions, do not have LUNs on the same storage system volume as other data (for example, an NFS share).

**Note:** For clustered Data ONTAP, the junction path name should be the same as the volume name. For example, if the volume name is `vol1`, the junction path should be `/vol1`.

## 7 Snapshot Management

SnapDrive for UNIX integrates with NetApp Snapshot technology to make stored data reliable for host applications. The ability to create and manage Snapshot copies from the host is one of the main features that make SnapDrive for UNIX attractive to users.

Snapshot copies record the state of the blocks in your file system at a given point in time and provide read-only access to that image of the NetApp LUN or NetApp NFS share. SnapDrive for UNIX enables you to create, restore, and delete Snapshot copies of the file system, volume group, host volume, or LUN and to clone storage entities from Snapshot copies. For more information about the commands used to perform these tasks, refer to the [SnapDrive 5.3 for UNIX Administration Guide](#).

SnapDrive for UNIX Snapshot copies are widely used because SnapDrive for UNIX offers the following distinct advantages:

- Hosts consistent Snapshot copies (restorable copy)
- Has faster restore times
- Creates backups of larger amounts of data quickly

### 7.1 Best Practices for Snapshot Copy Management

As a best practice, you should disable automatic Snapshot copy creation on the storage system for the volume on which the LUNs are created and set the Snapshot space reservation to 0 by running the following commands on the storage system:

```
vol options <vol-name> nosnap {on | off}  
snap reserve <vol_name> 0
```

### 7.2 Consistent Snapshot Copies

The process of creating Snapshot copies in a SAN environment differs from that of a NAS environment in one fundamental way: In a SAN environment, the storage system does not control the state of the file system.

Snapshot copies of a single storage system volume that contain all of the LUNs supporting a particular host file system are consistent. If the LUNs in the host file system span are different from the storage system volumes or storage systems, the copies might not be consistent unless they are created at exactly the same time across different storage system volumes or storage systems and can be restored successfully. Starting with SnapDrive 5.2 for UNIX, consistent Snapshot copies can be created by using the clustered Data ONTAP consistency group feature, which is supported beginning with clustered Data ONTAP 8.2.

A consistency group is a grouping of a set of volumes that must be managed as a single logical entity. The functional objective of a consistency group is to provide storage-based, crash-consistent checkpoints from which an application can restart. These checkpoints are used without interaction or coordination with the source application.

A checkpoint represents a collection of Snapshot copies, one Snapshot copy per volume, for all volumes defined in a consistency group. This collection of Snapshot copies is not the same as a regular group of Snapshot copies, in which each copy is created independently of each volume. This special collection of Snapshot copies, the checkpoint, has some distinct characteristics:

- The copy of volumes occurs as an automatic operation.
- The resulting Snapshot copy preserves write ordering across all volumes for dependent writes.

The internal operation of a consistency group can be expressed by the following high-level process:

1. SnapDrive issues a start checkpoint call to all participating controllers.
2. The controllers fence write access on volumes in a consistency group.
3. The controllers prepare a Snapshot copy of all volumes.
4. SnapDrive receives fence success from all participating controllers and issues a commit checkpoint to all controllers.
5. Upon receiving a commit checkpoint from SnapDrive, controllers commit the `snapshot create` in all volumes.
6. The controllers unfence all volumes in the consistency group.

#### Example:

If the file system `/mnt/fs_multi_vol` resides over LUNs in `storage1:/vol/vol1` and `storage2:/vol/vol1`, and Data ONTAP 7.2 or later is installed on `storage1` and `storage2`, then the following command creates a consistency group for the storage system volumes `storage1:/vol/vol1` and `storage2:/vol/vol1` and creates a Snapshot copy consistent with respect to both of the volumes:

```
snapdrive snap create -fs /mnt/fs_multi_vol -snapname snap1
```

**Note:** Crash-consistent Snapshot copies are not supported in versions earlier than clustered Data ONTAP 8.2.

## Best Practices for Consistency Group Snapshot Copies

In an environment in which all participating controllers support consistency groups, SnapDrive uses a Data ONTAP consistency group Snapshot copy as the preferred (default) method of capturing multicontroller or volume Snapshot copies.

SnapDrive simplifies the use of consistency groups. When the file specs or the file system dictates a Snapshot copy that spans multiple volumes and controllers and all target controllers support consistency groups, SnapDrive automatically recognizes this requirement and creates consistency groups to enable crash-consistent Snapshot copies. No change to the SnapDrive syntax is necessary to take advantage of consistency groups.

When the Data ONTAP consistency group feature is not supported or is not preferred—for example, when one of the controllers does not support consistency groups, or multiple applications share LUNs from a single volume for which the volume I/O fencing might affect other applications—a user can specify the `-nofilerfence` option, which requests that SnapDrive not use consistency groups in creating Snapshot copies. In this case, SnapDrive uses the best-effort mode to create a consistent Snapshot copy. In a best-effort mechanism, SnapDrive for UNIX creates a Snapshot copy that spans multiple storage system volumes without freezing the target storage and then checks for read-write I/Os that occurred when the Snapshot copy was created. If any I/Os occurred during Snapshot copy creation, SnapDrive for UNIX discards the Snapshot copy and reports the failure to the user.

SnapDrive for UNIX allows making Snapshot copies of multiple file systems or volume groups by using a single command when the file systems or volume groups are independent of each other. NetApp recommends using the `-unrelated` option when NFS entities are present with other file specifications in

the Snapshot copy. The reason is that the best-effort mechanism cannot be applied on NFS entities because Data ONTAP does not provide the required statistics as it does for LUNs.

SnapDrive for UNIX creates a Snapshot copy of unrelated storage entities that have no dependent writes during Snapshot copy creation. Each of the storage entities is crash consistent individually, but they are not consistent as a group. In the following example, the volume group consists of LUNs that reside on one storage system, and the file system consists of LUNs that reside on a different storage system:

```
snapdrive snap create -fs /mnt/fs1 -vg vg1 -unrelated -snapname snapfs1_vg1
```

### 7.3 Snapshot Space Management

Snapshot backups occur within seconds, and each copy typically consumes only the amount of data that has changed since the previous copy was created. Snapshot copies consume minimal disk space while providing up to 255 online point-in-time images.

The amount of disk space consumed by an individual Snapshot is determined by the following factors:

- The rate at which the data changes within the active file systems, which can be measured in megabytes per second or megabytes per hour
- The amount of time that elapses between creation of Snapshot copies

#### Best Practices

- Disable automatic Snapshot copy creation on the storage system for the volume on which the LUNs are created.
- Periodically use the `snapdrive snap list` command and delete old Snapshot copies, which could unnecessarily occupy space.

### 7.4 Snap Reserve

Data ONTAP reserves a default of 5% of the volume for Snapshot copy space consumption. This is because Snapshot copies need space, which they consume in the snap reserve area. After the snap reserve area is filled, the Snapshot copies start to take space from the active file system. Because of NetApp WAFL® (Write Anywhere File Layout) technology, snap reserve does not reserve specific physical blocks; rather, it is a logical space-accounting mechanism.

#### Best Practices for Snap Reserve

NetApp recommends setting the snap reserve value to 0% on any volume that contains LUNs exclusively. Use of any other value does not offer any benefit and might result in wasted space through unintended interaction with the fractional reserve policy on the volume.

For example, a fully provisioned 100GB LUN with a 100% space-reservation policy requires 100GB of additional space in a volume to create a Snapshot copy and support 100% turnover of the LUN data. The result is 100GB space consumption in the LUN, 100GB fractional reserve space consumption, and a total volume size of 200GB.

For more information about fractional reserve, refer to the [Clustered Data ONTAP 8.3 Logical Storage Management Guide](#).

## 8 Restoring a Snapshot Copy Using Volume-Based SnapRestore

In version 3.0 and earlier, SnapDrive for UNIX uses single-file SnapRestore implemented in Data ONTAP. While a restore is in progress for a file through single-file SnapRestore, any operation that tries to change the file is suspended until the restore operation is complete. For LUNs, when a single-file SnapRestore operation is in progress, the LUN is available, and I/Os (both reads and writes) are allowed. Single-file

SnapRestore for normal files as well as LUNs might take a long time, depending on the size of the LUN or the file being restored. Because of the length of time, for some environments single-file SnapRestore might not be the best option.

Starting with version 4.0, SnapDrive for UNIX can leverage volume-based SnapRestore technology. This requires less CPU and fewer storage resources, and it instantaneously restores all of the LUNs or normal files in a volume from the same Snapshot copy or backup.

Volume-based Snapshot technology should be used with caution because all Snapshot copies created after the one being used for the restore operation are deleted. All new files and new LUNs created in this volume are also deleted. Any relationship to secondary storage systems is broken if the restored Snapshot copy is older than the baseline Snapshot copy in the relationship.

## 8.1 Using Volume-Based SnapRestore

Run the following command to restore a volume by using the volume-based SnapRestore method:

```
snapdrive snap restore [-lun | -dg | -vg | -hostvol | -lvol | -fs | -file]
file_spec [file_spec ...] [{-lun | -dg | -vg | -hostvol | -lvol | -fs | -file} file_spec
[file_spec ...] ...] -snapname snap_name [-force [-noprompt]] [{-reserve | -noreserve}] [-vbsr
[preview|execute]]
```

### Examples:

- Volume-based SnapRestore in preview mode:

```
snapdrive snap restore -hostvol vgl/lvol1 -snapname snapvgllvol1 -vbsr
```

- Volume-based SnapRestore without prompting any confirmation message:

```
snapdrive snap restore -fs /mnt/fsl -snapname snapfsl -force -noprompt -vbsr execute
```

- Using the `-reserve` or `-noreserve` option with the `snapdrive snap restore` command overrides the space-guarantee policy on volumes set with the `space-reservations-volume-enabled` variable in the `snapdrive.conf` file.

- The following command is an attempt to set a space guarantee for volume as `volume`:

```
snapdrive snap restore -vg vgl -snapname snapvgl -force -reserve -vbsr execute
```

- The following command is to set a space guarantee as `none`:

```
snapdrive snap restore -fs /mnt/fsl -snapname snapfsl -force -noprompt -noreserve -vbsr execute
```

### Best Practices

- Run the `-vbsr preview` command before you use the `-vbsr execute` command. If the `preview` option is used, SnapDrive performs a series of checks on the volume that is being restored and presents a file-by-file analysis of the restore operation before it occurs. The preview analysis can help you decide whether you want to proceed with the volume-based SnapRestore operation.
- Use dedicated volumes for data that must be restored quickly by implementing the volume-based SnapRestore method.
- If there are Snapshot copies that were created after the Snapshot copy being used to restore them, replicate those Snapshot copies to a secondary storage system by using NetApp SnapVault® technology and then perform the volume-based SnapRestore operation.

## 9 Cloning

SnapDrive for UNIX allows you to clone existing file systems from Snapshot copies. You might use the cloning feature in the following scenarios:

- When there is an available update for the application that is running on the storage system LUNs or NFS file system, you can clone the storage, update the software, and verify that the software runs satisfactorily before you use it in production.
- You can create a copy from a Snapshot copy backup of an existing file system on the NetApp storage system that can be mounted on the same host or on a different host to separate the upgrade and testing process. After the new application update, the cloned file system can be split from the parent volume or LUN.

Earlier versions of SnapDrive for UNIX leveraged FlexClone technology in an NFS environment and LUN clone technology in a SAN environment. Starting with version 4.0, SnapDrive for UNIX can leverage FlexClone technology even in a SAN environment based on the value of the `san-clone-method` option in the `snapdrive.conf` file. Table 6 lists the various possible values that can be specified for the `san-clone-method` option in the `snapdrive.conf` file.

Table 6) FlexClone configuration values.

CLI Usage	snapdrive.conf Variable and Value	Description
<code>-clone lunclone</code>	<code>san-clone-method=lunclone</code>	Creates a LUN clone.
<code>-clone unrestricted</code>	<code>san-clone-method=unrestricted</code>	Creates a FlexClone volume that can be used as a back end for provisioning and Snapshot operations, the same as for normal flexible volumes.
<code>-clone optimal</code>	<code>san-clone-method=optimal</code>	SnapDrive attempts to create a FlexClone volume. If it is unable to do so, it reverts to the LUN clone method.

If the `-clone <lunclone|unrestricted|optimal>` option is used with the `snapdrive snap connect` command, it overrides the `san-clone-method` value defined in the `snapdrive.conf` file.

SnapDrive for UNIX verifies the following to create a FlexClone volume:

- The storage system Data ONTAP version is 7.0 or later, and the FlexClone software is licensed.
- The host filespec-residing volume is a flexible volume.
- A host filespec can be a file system, a host volume, a volume group, or a LUN.
- The host filespec-residing volume is not the root volume.
- Enough space is available on the aggregate.

If any of these checks fails, SDU takes one of the following actions:

- Errors out for an unrestricted FlexClone volume
- Falls back to the LUN clone method for a restricted FlexClone volume (if the clone method is optimal)

To use FlexClone technology in a SAN environment, a FlexClone license is required. No separate license is required for creating LUN clones.

For clustered Data ONTAP, the LUN clone created is a single-instance storage (SIS) clone.

### 9.1 Benefits of FlexClone Technology

FlexClone technology provides the following benefits:

- Simplified data management and reduced risk

- Flexibility and greater utilization. You can use FlexClone technology to create numerous copies of data for additional users without giving them access to the original data.
- Faster cloning. It is faster than a LUN clone because a LUN clone requires allocating new inodes.

### Examples:

- The following command connects to a volume group with two LUNs and one file system. It specifies a destination name for the volume group, the logical volume, and the file system:

```
snapdrive snap connect -fs /tmp/test1 /tmp/test2 -destlv lvoll_dup/vgl_dup -snapname btc-ppe-42:/vol/src_snapmirror_vol:snapvg1
```

- The `-autoexpand` option eliminates the need to specify each logical volume or file system to connect the entire volume group. The `-autorename` option along with `-autoexpand` renames the entities when the default name is in use:

```
snapdrive snap connect -fs /tmp/fs2 -snapname btc-ppe-42:/vol/src_snapmirror_vol:snapfs2 -autoexpand -autorename
```

## 9.2 Best Practices for Cloning

Because you can make several clones from a Snapshot copy, NetApp recommends naming the Snapshot copy in a way that indicates its usage. One simple naming convention is to prefix the characters `cl_` to the cloned file system name; however, SnapDrive for UNIX does not have a convention for naming Snapshot copies. You can also use the `prefix-clone-name` option in the `snapdrive.conf` file to have SnapDrive automatically prefix a string to the FlexClone copy.

Data ONTAP locks any Snapshot copies used to back up cloned volumes and LUNs until the clone is either split or destroyed. Any disk blocks associated with such a copy remain locked and cannot be reused until the copy is deleted. NetApp recommends that you regularly review your existing Snapshot copies and clones to determine whether they need to be destroyed.

By default, SnapDrive for UNIX assumes that no space reservation is required for cloned LUNs and sets it to 0%. If you want to set the space reservation, you can use the `snapdrive snap connect` command with the `-reserve` option to enable the storage reservation. You must consider capacity planning when you use clones.

Starting with SnapDrive for UNIX 3.0, the LUN clone-split operation is supported with the `snapdrive snap connect` command. This operation can be achieved either by having `enable-split-clone=on` in the `snapdrive.conf` file or by using the `-split` option in the `snapdrive snap connect` command.

Starting with SnapDrive for UNIX 4.2, clone-split operations can be performed through the new clone-split commands.

Consider the following guidelines when performing a clone split:

- Before executing the LUN clone-split operation, make sure the volume has enough space to accommodate the cloned LUN; otherwise, resize the volume. Starting with SnapDrive for UNIX 4.2, the space can be verified by using the `snapdrive clone split estimate` command.
- Plan the clone-split operation during a low I/O time period.
- Use LUN clone split only if you are expecting the cloned LUN to be read-write and are expecting it to affect the original LUN.
- NetApp recommends using the `-split` option instead of `enable-split-clone=on` so that you can control the behavior per command. Starting with SnapDrive for UNIX 4.2, the `snapdrive clone split start` command can be used for splitting the clones.

SnapDrive for UNIX allows you to start a clone-split operation in asynchronous or synchronous mode. In asynchronous mode, the clone-split operation runs in the background, and a job ID is displayed. In

synchronous mode, the clone-split operation runs in the foreground. The mode can be set either by the `split-clone-async=on|off` in `snapdrive.conf` file or by using the `-splitmode` option as `async|sync` in the `snapdrive clone split` command.

#### Examples:

- The following example estimates storage space for a volume clone by using a Snapshot copy:

```
snapdrive clone split estimate -hostvol vg1/lvol1 -snapname btc-ppe-42:/vol/src_snapmirror_vol:snapvg1_lvol1 -volclone
```

- The following example estimates the storage to split a LUN clone by using a Snapshot copy with the `-fs` option:

```
snapdrive clone split estimate -fs /mnt/fs1 -snapname btc-ppe-42:/vol/src_snapmirror_vol:snapfs1 -lunclone
```

- The following example splits a volume clone in asynchronous mode:

```
snapdrive clone split start -fs /tmp/sfp1 /tmp/sfp2 -splitmode async
```

**Note:** If you set the `enable-split-clone` configuration variable value to `on` or `sync` during the Snapshot copy-connect operation and to `off` during the Snapshot copy-disconnect operation, SnapDrive for UNIX does not delete the original volume or LUN that is present in the Snapshot copy.

**Note:** LUN clone split is not supported with MultiStore.

**Note:** LUN clone split is not possible for clustered Data ONTAP because the created clones are SIS clones.

## 10 SnapDrive for UNIX with SnapMirror Load-Sharing Mirror Relationships

SnapMirror load-sharing mirrors increase performance and availability for NAS clients by distributing an SVM namespace root volume to other nodes in the same cluster and distributing data volumes to other nodes in the cluster to improve performance for large read-only workloads.

By default, all client requests for access to a volume in a load-sharing mirror set are granted read-only access. Read-write access is granted by accessing a special administrative mount point, which is the path that servers requiring read-write access into the load-sharing mirror set must mount. All other clients have read-only access. After changes are made to the source volume, the changes must be replicated to the rest of the volumes in the load-sharing mirror set.

There are no changes in operations on the source volume when SnapDrive for UNIX is used to perform them. If there is a failover of load-sharing mirror relationships, the user can create a new backup and then perform restore or clone operations on the promoted target volume.

For more information, refer to the [Clustered Data ONTAP 8.3 Logical Storage Management Guide](#).

**Note:** If you are using clustered Data ONTAP 8.2, 8.2.1, or 8.2.2, the load-sharing mirror update fails if the cluster administrator is not configured. You must configure the cluster administrator by using the `snapdrive config set -cserver` command.

## Best Practices

Create a load-sharing mirror of a NAS SVM namespace root volume on every node in the cluster so that the root of the namespace is available regardless of node outages or node failovers.

When a client requests access to a volume configured with a load-sharing mirror set, Data ONTAP directs all client connections to the load-sharing mirror destination volumes only; therefore, a destination volume should be created on the same node on which the source volume resides, allowing the namespace to provide a direct data access path to data volumes on that node.

## 11 Selective LUN Map in SnapDrive for UNIX

Beginning with clustered Data ONTAP 8.3, the selective LUN map (SLM) feature is enabled by default on all new LUN maps. When you create a new LUN map, the LUN is accessible only through paths found on the node that owns that LUN and its HA partner.

LUNs are accessible on all of the LIFs of an SVM. You should assign LIFs to the SVMs on each cluster node in your network. As the number of nodes in the cluster increases, the number of potential paths also multiplies, which can result in an excessive number of paths to a LUN, multiple igroups per host, and disruptive mobility events. SLM solves these problems by restricting LUN accessibility to the node that owns the LUN and the HA partner node. It also creates a single igroup per host and supports nondisruptive LUN mobility operations that do not require port set manipulation or LUN remapping.

For information about configuration parameters, refer to the section “Using Selective LUN Map in SnapDrive for UNIX” in the [SnapDrive for UNIX Administration Guide](#).

**Note:** SLM does not automatically apply to LUN maps created earlier than clustered Data ONTAP 8.3. If you are accessing the LUN through the node that owns the LUN, the path is called active optimized. However, if you access that LUN through the HA partner node, the path is called active nonoptimized.

## 12 SnapDrive for UNIX in a 7-Mode MultiStore Environment

SnapDrive for UNIX does not distinguish between a physical storage system and a vFiler unit. Therefore, there are no changes in operations on NetApp FlexVol<sup>®</sup> volumes when the operations are performed by using SnapDrive for UNIX on a vFiler unit.

Consider the following points when working in a MultiStore unit:

- Verify that a MultiStore license is available on the storage system.
- SnapDrive for UNIX can manage LUNs on MultiStore units by using the iSCSI protocol.
- SnapDrive for UNIX does not support FCP on MultiStore units.
- For SnapDrive for UNIX usage with vFiler units, refer to [KB1011960](#).
- SnapDrive operations on a volume are allowed only when the entire volume is owned by a vFiler unit.
- The Data ONTAP configuration option `vfiler.vol_clone_zapi_allow` must be set to `on` to connect to a copy of a volume or LUN in a vFiler unit.

## 13 Platform and Protocols

### 13.1 Virtualization Environment

SnapDrive for UNIX is supported in the following virtualization environments:

- Solaris containers (global, local, and native zones) for the NFS configuration only
- IBM AIX LPAR and DLPAR
- The SLES guest OS on a Microsoft Hyper-V virtual machine (VM) in Microsoft Windows for the iSCSI and NFS protocols only
- The Red Hat Enterprise Linux (RHEL) guest OS on a KVM hypervisor for the iSCSI and NFS protocols only

SnapDrive for UNIX supports raw device mapping (RDM) LUNs for VMware ESX inside the guest VM for the FC protocol; storage resizing; and all existing operations of cloning, Snapshot copies, and restore.

**Note:** There are no changes with respect to SnapDrive commands. Refer to section 13.3, “RDM LUN Support,” for instructions on how to set up SnapDrive for UNIX for use with RDM LUNs within a VM.

RDM LUN support is available for the following ESX guest platforms only:

- Red Hat Enterprise Linux
- Solaris

For a complete list of supported versions, platforms, and configurations, refer to the [Interoperability Matrix Tool](#).

## 13.2 ALUA Support

Asymmetric Logical Unit Access (ALUA) defines the protocol for how multipath I/O should be managed between hosts and storage devices. This standard acknowledges the performance difference between the paths if there are multiple paths to the LUN.

SnapDrive for UNIX enables ALUA for the igroup if the Data ONTAP version supports the ALUA command set and if the multipath solution in the host understands it.

## 13.3 RDM LUN Support

RDM allows a special file in a Virtual Machine File System (VMFS) volume on an ESX host to act as a raw device inside a VM. With RDM, the ESX server accesses the LUN, which in turn passes access directly to a VM for use with its native file system. SnapDrive for UNIX supports RDM LUNs on Linux and Solaris.

While working on RDM LUNs, consider the following guidelines:

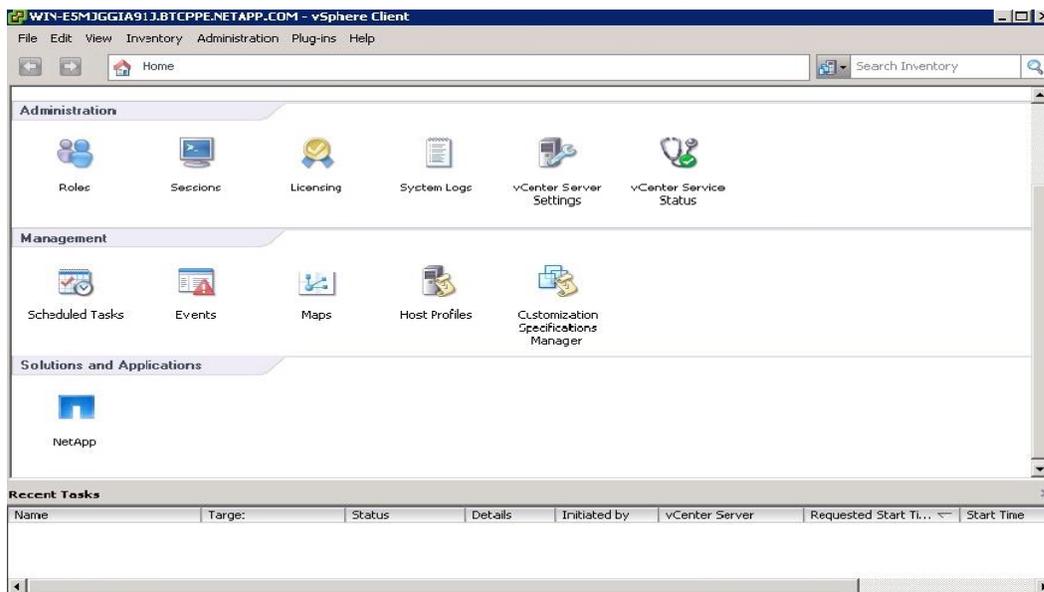
- SnapDrive for UNIX allows you to create and provision LUNs by using the FC protocol on a guest OS in a VMware virtual environment.
- RDM LUNs for SVMs in clustered Data ONTAP are not supported.
- RDM LUNs provide all functionalities in a guest OS that are feasible in a physical OS by using the FC protocol.
- The NetApp Virtual Storage Console (VSC) must be installed and configured with VMware vCenter running in a Microsoft Windows machine.
- SnapManager for virtual interface API calls are used to create and delete igroup mapping of the LUN and export and deport of the LUN to the guest OS.
- For more details about host OS, guest OS, VMware vCenter, and VSC supported versions, refer to the [Interoperability Matrix Tool](#).
- Verify that the ESX server has FC adapters installed.
- Verify that the zoning is configured properly between the ESX server and the storage system.
- Storage system credentials should be configured on the backup and recovery capability of VSC.
- FCP Host Utilities are not required in a guest OS for RDM LUN support.

- VMware limitations are:
  - Each guest can be configured with 4 SCSI controllers, each of which can be mapped to 16 devices. However, one device is reserved per controller; therefore, 60 (16 x 4 – 4) LUNs can be mapped to the guest.
  - Each ESX server can have a maximum of 256 LUNs mapped to it.
  - The maximum supported LUN size is 2TB minus 512B.

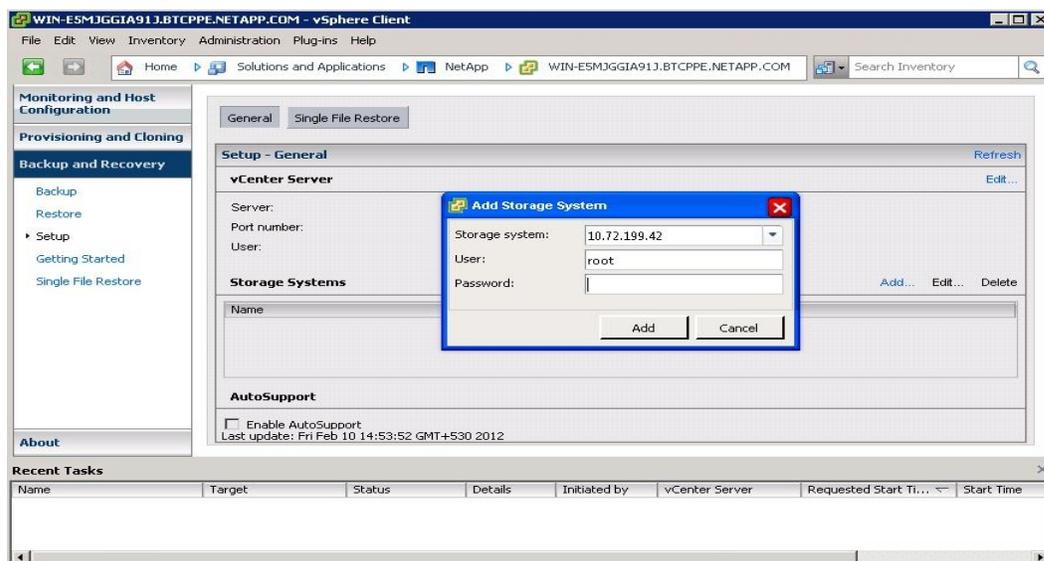
## Configure SnapDrive for UNIX for RDM LUNs

To configure SnapDrive for UNIX for RDM LUNs, complete the following steps:

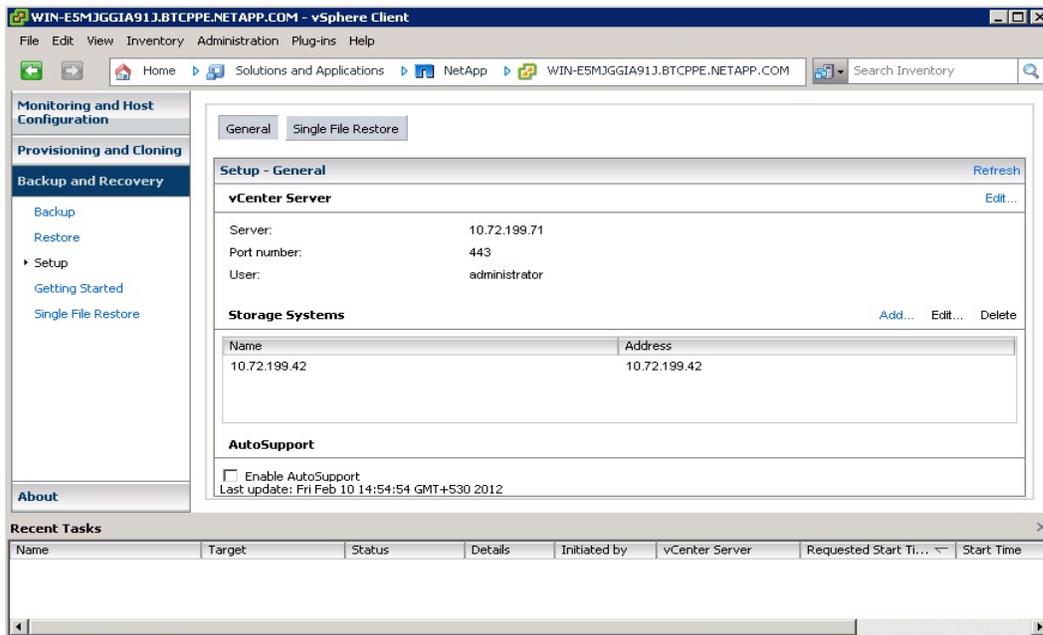
1. After installing the NetApp VSC plug-in on the vCenter instance, log in to the vCenter server through VMware vSphere. The NetApp icon appears in the vSphere Client.



2. Click the NetApp icon to display the home view for the VMware vSphere NetApp plug-in. Configure the storage system by using the Backup and Recovery option.



3. The configured storage system with backup and recovery capability appears.



4. Set `default-transport` as `fc` in the `snapdrive.conf` file and restart the SnapDrive for UNIX daemon.

**Note:** The following configuration variables are also included in `snapdrive.conf`:

- `contact-viadmin-port` specifies the port on which the SnapManager for Virtual Infrastructure (SMVI) service is running. The default is 8043.
- `use-https-to-viadmin` specifies the protocol (HTTP or HTTPS) for communicating with SMVI. The default is `on` (HTTPS for communication).
- `virtualization-operation-timeout-secs` specifies the number of seconds to wait before timing out a call to SMVI. The default is 600 seconds.
- `vif-password-file` specifies the path for the virtual interface password file. The default is `/opt/NetApp/snapdrive/.vifpw`.

5. Use the following command to specify the VSC server with backup and recovery capabilities:

```
snapdrive config set -viadmin <vi-admin-user> <vi_server>
```

**Example:**

```
# snapdrive config set -viadmin admin 10.72.199.71
```

- Password for admin: <SMVI installed server-level user password of 10.72.199.71>
  - Retype password:

**Note:** User credentials are not validated by the SMVI server.

- Verify by using the `snapdrive config` command. The appliance type for the specified SMVI server is `Virtual Interface`.

**Example:**

```
Snapdrive config list
username appliance name appliance type
-----
root 10.72.199.42 StorageSystem
```

```
admin 10.72.199.71 VirtualInterface
```

These steps should be performed on each guest OS on which RDM LUNs are to be created.

## RDM LUN Creation

SnapDrive for UNIX performs the following actions in the creation of RDM LUNs:

- Creates the LUN on the storage
- Creates igroup mappings on the storage, creates RDM mappings, and exports the LUN to the guest
- Discovers the LUN and creates the host-side file specifications.

### Example:

```
snapdrive storage create -lun 10.72.199.42:/vol/esxvol1/lun1 -lunsize 20m -fs /mnt1 -nolvm LUN
btc-ppe-42:/vol/esxvol1/lun1 ... created
exporting new lun(s) to Guest OS ... done
discovering new lun(s) ... Done
LUN to device file mappings:
- btc-ppe-42:/vol/esxvol1/lun1 => /dev/sdh
filesystem /mnt1 created
```

## RDM LUN Deletion

SnapDrive for UNIX performs the following actions in the deletion of RDM LUNs:

- Deletes the host-side file specifications.
- Deports the LUN and removes the RDM mapping
- Deletes the LUN on the storage

### Example:

```
SDU deletes the host-side file specs.
SDU deports the LUN, removes the RDM mapping.
SDU deletes the LUN on the storage.
snapdrive storage delete -fs /mnt1
delete filesystem /mnt1
- fs /mnt1 ... deleted
deporting lun(s) from Guest OS ... done
- LUN btc-ppe-42:/vol/esxvol1/lun1 ... deleted
```

## Limitations

SnapDrive for UNIX has the following limitations in relation to RDM LUNs:

- RDM LUNs for SVMs in clustered Data ONTAP are supported.
- When SDU commands are executed with an igroup, SDU ignores the user-defined igroup with a warning message. For example, when the transport protocol is FCP, the igroup that is specified in the CLI command is ignored by SnapDrive, and that igroup is automatically created by the virtual interface.
- Multipathing is not supported in the guest OS.
- SMVI does not support virtual compatibility mode RDM.
- SMVI does not perform any authentication checks that are configured in SDU.

## Best Practices

- Creating an RDM LUN in a VMware environment enables the guest OS to perform fast I/O operations. Use SnapDrive for UNIX 4.2 or later to perform Snapshot management and storage provisioning operations.
- If you discover that the RDM LUN was not created and it behaves as an ordinary LUN, verify that the default transport variable is set to `FCP` in the `snapdrive config` file. Run the `snapdrive config list` command to validate the presence of the virtual interface.
- When you execute storage provisioning operations for RDM LUN mapping, the entries are listed in the ESX server. To avoid overheads, NetApp recommends removing the stale entries in the ESX server that are generated by SnapDrive operations.

## 14 Summary

SnapDrive for UNIX is a storage management tool. It helps storage administrators create backups based on Snapshot copies and helps restore Snapshot copies that contain application data. SDU can also be used to provision storage from the host system for any application.

The recommendations made in this guide are best practices for most deployments.

## References

- SnapDrive for UNIX Installation and Administration Guide  
[https://library.netapp.com/ecm/ecm\\_get\\_file/ECMP12476257](https://library.netapp.com/ecm/ecm_get_file/ECMP12476257)
- [SnapDrive 5.3 for UNIX Installation and Setup Guide for Linux for Clustered Data ONTAP](#)
- [SnapDrive 5.3 for UNIX Installation and Setup Guide for Linux for 7-Mode](#)
- Clustered Data ONTAP 8.3 Documentation  
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=61898&language=en-US>
- Host Utilities 6.2 Installation and Setup Guide  
[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP1217221](https://library.netapp.com/ecm/ecm_download_file/ECMP1217221)
- Clustered Data ONTAP 8.3 Logical Storage Management Guide  
<https://library.netapp.com/ecmdocs/ECMP1610211/html/frameset.html>
- Clustered Data ONTAP 8.3 SAN Administration Guide  
<https://library.netapp.com/ecmdocs/ECMP1636035/html/frameset.html>
- Clustered Data ONTAP 8.3 File Access Management Guide for NFS  
<https://library.netapp.com/ecmdocs/ECMP1610208/html/frameset.html>

## Version History

Version	Date	Document Version History
Version 1.3	July 2015	Updated with clustered Data ONTAP features and multipath alias support.
Version 1.2	November 2014	Updated to include 5.3 release components.
Version 1.1	April 2014	Updated to include 5.2.1 release components.
Version 1.0	August 2013	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### Copyright Information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

TR-4212-1015