



Technical Report

SnapDrive 5.2 for UNIX: Best Practices Guide

Ashwath Narayan, NetApp
August 2013 | TR-4212

TABLE OF CONTENTS

1	Introduction	4
1.1	Intended Audience	4
1.2	Prerequisites	4
2	Overview	4
2.1	Key Features	4
2.1	SnapDrive for UNIX Deployments	5
3	New Features in Version 5.2	5
4	Limitations of SnapDrive for UNIX	5
4.1	General Limitations of SnapDrive for UNIX	5
4.2	Limitations of SnapDrive for UNIX on Linux	6
4.3	Limitations of SnapDrive for UNIX on Solaris	6
4.4	Limitations of SnapDrive for UNIX on AIX	6
5	Installation and Basic Configuration	6
5.1	Checklist for Storage System	6
5.2	Checklist for Host Before SnapDrive Installation	8
5.3	snapdrive.conf File	10
5.4	Using the SnapDrive Daemon	10
5.5	Setting Up SnapDrive to Communicate with a NetApp Storage System	10
5.6	Volume Managers and File Systems	11
6	Security Features	13
6.1	Best Practices	13
6.2	Configuring SnapDrive for RBAC	14
7	Storage Provisioning	15
8	Host Partitioning Support in RHEL and SUSE Linux Enterprise Servers	16
9	Snapshot Copy Management	17
9.1	Best Practices for Snapshot Copy Management	17
9.2	Consistent Snapshot Copies	17
9.3	Snapshot Copy Space Management	18
9.4	Snap Reserve	19
10	Restoring a Snapshot Copy Using Volume-Based SnapRestore	19
10.1	Using Volume-Based SnapRestore	20
11	Cloning	20

11.1 Benefits of FlexClone	21
11.2 Best Practice for Cloning.....	22
12 SnapDrive for UNIX in a MultiStore Environment	23
13 Platform and Protocols	23
13.1 Virtualization Environment	23
13.2 ALUA Support.....	24
13.3 RDM LUN Support.....	24
14 SnapDrive for UNIX FAQ.....	28
14.1 SnapDrive for UNIX Installation FAQ.....	28
14.2 Storage Provisioning FAQ.....	30
14.3 Snapshot Copy Management FAQ.....	32
14.4 SnapDrive for UNIX in a Storage Foundation for Oracle RAC (SFRAC) Environment (Storage Creation) FAQ	33
14.5 General Queries.....	33
15 Summary	34
References.....	34
Version History	34

LIST OF TABLES

Table 1) Checklist for storage system.	7
Table 2) Manual preinstallation steps.....	9
Table 3) SnapDrive daemon operations.....	10
Table 4) Predefined roles on the Operations Manager server.....	15
Table 5) Host partition information.....	16
Table 6) FlexClone configuration values.....	21
Table 7) Default locations for snapdrive.conf file.....	30
Table 8) Target information for host operating systems.....	30
Table 9) Rescan iSCSI targets.....	31
Table 10) Steps to resolve SnapDrive storage create/connect command errors.....	31

1 Introduction

This document is a best practices guide for NetApp® storage systems using the SnapDrive® for UNIX® (SDU) solution. It also provides recommendations on various configuration options available with the solution in addition to providing guidelines for when and where to use the options.

1.1 Intended Audience

This guide provides guidelines for deploying and using SnapDrive for UNIX with NetApp storage appliances and supporting software.

This guide is intended for storage and server administrators and experts managing storage provisioning and Snapshot™ copies and NetApp storage systems using the SnapDrive for UNIX solution.

1.2 Prerequisites

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

- [SnapDrive for UNIX Installation and Administration Guide](#)
- [Clustered Data ONTAP 8.2 Documentation](#)
- [Host Utilities Installation and Setup Guide](#)

You should have a good understanding of UNIX file system administration as well as FCP and iSCSI concepts. The recommendations in this document are guidelines to assist storage and server administrators while configuring SnapDrive for UNIX.

2 Overview

SnapDrive for UNIX is an enterprise-class storage and data management solution that simplifies storage management and increases the availability and reliability of application data. The key functionality includes error-free application storage provisioning, consistent data Snapshot copies, and rapid application recovery. It also provides the ability to easily manage data residing on NetApp 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.

SnapDrive for UNIX is supported on the following platforms.

- Linux®
 - Red Hat Enterprise Linux
 - SUSE Linux Enterprise Server
 - Oracle® Enterprise Linux/Oracle Linux
- Solaris
- IBM AIX

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

2.1 Key Features

Following are the key features of SnapDrive for UNIX:

- **Advanced storage virtualization.** Virtualizes NetApp storage systems and integrates with native disk and volume management.
- **Automated mapping and management of new storage resources.** Provisions new storage resources efficiently, without application or host downtime.

- **Dynamic storage allocation.** Reallocates storage systems quickly and easily in response to shifts in application or server demand.
- **Host file system consistent Snapshot copies (consistency groups).** Implements a mechanism by which Snapshot copies that are created are file system consistent.
- **Backup and restore.** Provides a quicker way to create and restore storage from backups using NetApp Snapshot technology.
- **Role-based access control (RBAC) permissions**

For more details about SnapDrive for UNIX features, refer to the [SnapDrive for UNIX Installation and Administration Guide](#).

2.1 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 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 using Snapshot technology.

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

3 New Features in Version 5.2

Following are the new features in version 5.2:

- **The consistency group Snapshot (CGS)** capability of Data ONTAP 8.2 is leveraged to provide support for SAN protocols, with both clustered Data ONTAP and Data ONTAP operating in 7-Mode.
- **Support of clustered Data ONTAP** for both NFS and SAN.
- **Enhanced AutoSupport™ reporting** to generate reports in the following scenarios:
 - Performing a backup operation
 - Configuring a new storage system
 - Starting a new daemon
 - Operation failure
- **Support for NFS operations** on a volume containing nested junction paths that are in clustered Data ONTAP.
- **Support for Cisco® converged network adapters (CNAs).**
- **Clustered Data ONTAP support.** SnapDrive for UNIX clustered Data ONTAP support for both NFS and SAN protocols is available only with Data ONTAP 8.2 RC1 and later releases.

4 Limitations of SnapDrive for UNIX

4.1 General Limitations of SnapDrive for UNIX

- SnapDrive for UNIX does not support snap operations on a Network File System (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 any mount points.
- Snap operations might be inconsistent if you perform SnapRestore® on a mount point where a different entity is mounted other than what was created by the Snapshot copy.
- SnapDrive modifies the mount point permissions from nonroot user to root user for a qtree after volume-based SnapRestore (VBSR) operations.
- From 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 are removed.
- The clone split operation is not supported on a filespec, which is a combination of either 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.

4.2 Limitations of SnapDrive for UNIX on Linux

- In clustered Data ONTAP, to configure data a Storage Virtual Machine (SVM) (formerly called Vserver) with SnapDrive for UNIX, make sure that the IP address of the data SVM's management logical interface (LIF) is mapped to the data SVM name either in Domain Name System (DNS) or in the `/etc/hosts` file. 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.
- 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, export policy is defined only at the volume level. Therefore, the `.snapshot` directory cannot be exported to the secondary host, which does not have export permission on the parent volume.

4.3 Limitations of SnapDrive for UNIX on Solaris

- LUNs greater than 1TB with the Storage Management Initiative (SMI) label are not supported.
- SnapDrive for UNIX does not support the extensible firmware interface (EFI) and SMI LUNs in the same disk group.

4.4 Limitations of SnapDrive for UNIX on AIX

- The SnapDrive daemon takes a long time to restart if one or more Fibre Channel (FC) port links are down or not available. However, the SnapDrive daemon starts irrespective of the state of the FC ports.
- SnapDrive for UNIX, by default, creates disk groups with nonconcurrent settings, and you need to change or provide storage provisioning operations manually.

5 Installation and Basic Configuration

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 following checklists to avoid potential errors or delays during or after the installation.

5.1 Checklist for Storage System

The checklist for storage systems is described in Table 1.

Table 1) Checklist for storage system.

Step	Action
1.	<p>Verify the licenses on the storage system:</p> <ul style="list-style-type: none"> • FCP, iSCSI, or NFS license, depending on your configuration • FlexClone® • SnapRestore software • MultiStore® software (vFiler® unit environment only) <p>On Data ONTAP operating in 7-Mode, use the <code>license</code> command on the command line interface (CLI) to verify installed licenses.</p> <p>On clustered Data ONTAP 8.2, 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 SSH because it is more secure. For detailed information to help you complete these tasks, refer to the Data ONTAP Storage Management Guide.</p>
3.	<p>For an NFS environment:</p> <ol style="list-style-type: none"> 1. Make sure storage system directories are exported correctly to the host. 2. 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. <p>Refer to the Data ONTAP File Access and Protocols Management Guide for how to manage NFS protocol.</p>
4.	<p>For an FC or iSCSI environment, verify whether FCP or iSCSI is enabled on the storage system:</p> <ol style="list-style-type: none"> 1. On Data ONTAP operating in 7-Mode, run <code>fcv status</code> or <code>iscsi status</code>. 2. If the status is disabled, start the service running <code>fcv start</code> or <code>iscsi start</code>. 3. On clustered Data ONTAP 8.2, run <code>Vserver fcp</code> or <code>Vserver iscsi show</code>. 4. If no configuration exists, create the configuration for Vserver using this command: <pre>Vserver fcp create -Vserver <Vserver> or Vserver iscsi create -Vserver <Vserver></pre>
5.	<p>Note down the storage system target address using these commands on the storage system CLI:</p> <ol style="list-style-type: none"> 1. On Data ONTAP operating in 7-Mode, run: <pre>fcv nodename or iscsi nodename</pre> <ol style="list-style-type: none"> 2. On clustered Data ONTAP 8.2, run: <pre>vserver fcp show -vserver <vserver> or vserver iscsi show -vserver <vserver></pre>
6.	<p>On Data ONTAP operating in 7-Mode, make sure that the Fibre Channel port on the NetApp storage system is configured as the target using the <code>fcadmin config</code> command.</p> <p>For more details, refer to the Data ONTAP Storage Management Guide.</p>

Best Practices

By default, the iSCSI service is enabled on all 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 by using the following command:

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

Examples:

```
iscsi interface disable e0b
```

Best Practices

On clustered Data ONTAP, a particular data logical interface can be disabled for an iSCSI service 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 logical interface. First disconnect the active sessions from the host; otherwise, the storage system will display a warning message when the command is issued.

5.2 Checklist for Host Before SnapDrive Installation

Preinstallation Checks Using the SnapDrive Configuration Checker

Before using or installing SnapDrive for UNIX, make sure that the host supports all the dependent components, such as native OS version, multipath solution, cluster solution, and so on, by referring to product support matrixes and product documentation. As SnapDrive for UNIX expands the supported platforms, the combination of supported configurations grows exponentially, and it becomes a major task for administrators to validate the environment. Administrators do not know which components to check and how to get their versions released, and these instances often end up as support cases.

The configuration checker is aimed at addressing the manual reference task by allowing the user to check if the host has all the supported components of SnapDrive for UNIX to run. The configuration checker is bundled along with SnapDrive for UNIX software; however, the configuration checker is independent software and can be downloaded from [ToolChest](#) on the [NetApp Support](#) site. It can be installed and executed irrespective of whether SnapDrive for UNIX is already installed on the host.

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 `sdconfcheck` 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 adding a new configuration (for example, MPIO for high availability) is planned, then the configuration checker helps to validate if 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](#) and confirm that SnapDrive for UNIX supports your environment.

For specific information on requirements, such as patches that are required for the operating system, see the [Compatibility and Configuration Guide for NetApp FCP and iSCSI](#).

Table 2 describes the manual preinstallation steps.

Table 2) Manual preinstallation steps.

Step	Action
1.	<p>Identify if the host operating system version is supported by SnapDrive for UNIX:</p> <ul style="list-style-type: none"> Linux: <code>cat /etc/issue</code>
2.	<ul style="list-style-type: none"> For FCP environments: <ul style="list-style-type: none"> a. Check that you are using the SnapDrive for UNIX supported host utility version from the Interoperability Matrix Tool. b. Verify if NetApp host utilities are present on the host by running <code>sanlun</code> version. c. If host utilities are not installed, download and install the correct version from the NetApp Support site and refer to the Host Utilities Installation and Setup Guide to install (or upgrade) and configure host utilities. d. Identify the HBA adapters on the host and verify that the ports are enabled by running <code>sanlun fcp show adapter -v</code>. For Veritas™ environments: <ul style="list-style-type: none"> a. Check that you are using Veritas Storage Foundation for multipathing, installing, and configuring the Symantec™ Array Support Library (ASL) for NetApp storage systems. b. To determine which version of the Veritas Storage Foundation and ASL to use, check the Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products. c. For more information about obtaining and installing ASL, see the Host Utilities Installation and Setup Guide. For iSCSI environments: <ul style="list-style-type: none"> a. Check that you are using the SnapDrive for UNIX supported host utility version from the Interoperability Matrix Tool. b. Verify that the NetApp Host Utilities are present on the host by running <code>sanlun</code> version. c. If Host Utilities are not installed, download and install the correct version from the NetApp Support site and see the Host Utilities Installation and Setup Guide to install (or upgrade) and configure Host Utilities.

Best Practices

- Download the latest package of the Configuration Checker tool from [ToolChest](#) on the [NetApp Support](#) site.
- Execute Configuration Checker postinstallation and preinstallation of the SnapDrive for UNIX tool to validate and confirm the presence of dependent components.
- NetApp recommends that you have a minimum of 1GB to 2GB spare space for SnapDrive for UNIX operation in the root directory (/) to avoid the root volume running out of space. SnapDrive for UNIX maintains audit, trace, and recovery log files whose default locations are specified in the `snapdrive.conf` file. The location of 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, and hence no default size is allotted for this file.
- After the preceding checklist is verified, follow the steps in the [SnapDrive for UNIX Installation and Administration Guide](#) to install SnapDrive for UNIX on different operating systems.
- NetApp recommends that you use the default path for the SnapDrive for UNIX installation and add the SnapDrive for UNIX directory to your `$PATH` environment variable.

5.3 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 using a text editor. The modified values will take effect after the SnapDrive daemon is restarted.

To change the entries in `snapdrive.conf`, follow these 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 current values of the configurable variables.

Examples:

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

- Valid format:

```
autosupport-enabled="off" #Enable autosupport (requires autosupport-filer be set)
contact-http-port=80 #HTTP port to contact to access the filer
```

- Invalid format:

```
autosupport-enabled="off"
autosupport-enabled=off #Enable autosupport (requires autosupport-filer be set)
autosupport-enabled="off" Enable autosupport (requires autosupport-filer be set)
contact-http-port=80a #HTTP port to contact to access the filer
contact-http-port=80
```

5.4 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, performing 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 using APIs. While using the SnapDrive for UNIX daemon, all the SnapDrive commands will work as a unique process in the background.

Common Daemon Operations

The SnapDrive daemon must be running for any SnapDrive for UNIX command to work.

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

5.5 Setting Up SnapDrive to Communicate with a NetApp Storage System

Follow these steps for SnapDrive to communicate with NetApp storage systems:

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 (DNS) or is in `/etc/hosts` file.

3. Execute `snapdrive config` and verify that the 7-Mode storage system is already configured. If it is not configured, execute 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 `vsadmin` user is unlocked and administration privilege is delegated to `vsadmin` user. For more information, refer to the [Clustered Data ONTAP 8.2 System Administration Guide](#).

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 of `<userid>`.
5. Edit the `snapdrive.conf` file and verify that the default settings conform to your environment.
 - a. Transport-type setting.

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

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

- b. HTTPS setting:

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

Note: Starting from 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.

5.6 Volume Managers and File Systems

Volume Manager 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 only supports creating, renaming, restoring, and deleting Snapshot copies for an existing storage device with the JFS file system.

Multipath

Following are the supported multipathing types 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.

Note: SnapDrive for UNIX does not support user-friendly names and aliasing in Linux multipathing environments.

Multisubnet Configurations

To manage network traffic, SnapDrive for UNIX allows numerous subnet configurations to split up management and data traffic. SnapDrive for UNIX sends all management communication 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; hence, you need to configure the management interface and data interface with the same IP. In mixed SAN and NFS environments, interfaces are mapped so that one management interface and data interface pair must be the same as the ones in pure SAN environments.

View all the data interfaces configured for a management interface by executing 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, 192.168.10.10:

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

In a mixed SAN and NFS environment:

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

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

6 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 need to 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 on how to set up and view access control permissions, refer to the [SnapDrive for UNIX Installation and Administration Guide](#).

6.1 Best Practices

Use HTTPS instead of HTTP for host to storage system communication to provide added security. To use HTTPS, follow these steps:

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

NetApp recommends that you practice using Operations Manager for RBAC, as described in section 6.2, “Configuring SnapDrive for ” for Data ONTAP 7-Mode.

Note: Operations Manager for RBAC is not supported in clustered Data ONTAP 8.2.

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

1. Create a 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 will be `/vol/vol0/sdprbac/sd<hostname>.prbac`, where the host name is the name of the SnapDrive host found by the `hostname` command.
3. Specify only one level in the `sd<hostname>.prbac` file among the list of permissions listed as follows:
 - a. **NONE:** The host has no access to the storage system.
 - b. **SNAP CREATE:** The host can create Snapshot copies.
 - c. **SNAP USE:** The host can delete and rename Snapshot copies.
 - d. **SNAP ALL:** The host can create, restore, delete, and rename Snapshot copies.
 - e. **STORAGE CREATE DELETE:** The host can create, resize, and delete storage.
 - f. **STORAGE USE:** The host can connect and disconnect storage.
 - g. **STORAGE ALL:** The host can create, delete, connect, and disconnect storage.
 - h. **ALL ACCESS:** The host has access to all the preceding SnapDrive for UNIX operations.

- The permission-level string must be on the first line of the permission file.
- Set the value of the `all-access-if-rbac-unspecified` option to `off` in the `snapdrive.conf` file for security.
- Restart the SnapDrive daemon.
- 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.
- Verify access permission through the `snapdrive config access show <NetApp Controller Name>` command.

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

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

For more information on access privileges, refer to the [SnapDrive for UNIX Installation and Administration Guide](#).

6.2 Configuring SnapDrive for RBAC

RBAC is implemented 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 provide support for the nonroot local user and NIS users by using the RBAC infrastructure of Operations Manager. SnapDrive for UNIX does not require a root password of the storage system; it communicates with the storage system using `sd-<hostname> user`.

The following steps need to be performed to set up RBAC.

1. The Operations Manager administrator creates a user, `sd-admin user`, with a capability core access check over global group (`global DFM.Core.AccessCheck`). After the Operations Manager administrator configures the `sd-admin user`, the administrator has to manually send the credential information to the SnapDrive for UNIX administrator.
2. The Operations Manager administrator can also grant (`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 needs to 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 needs to turn on RBAC functionality by setting the variable `rbac-method=dfm` in `snapdrive.conf` file and restarting the SnapDrive for UNIX daemon.

RBAC with Operations Manager is not supported for clustered Data ONTAP. The Operations Manager administrator needs to grant privileges to the invoker of SnapDrive to execute SnapDrive commands. For more information on the mapping of privileges compared with commands, see the [SnapDrive for UNIX Installation and Administration Guide](#).

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

- If you are a NIS user running the `snapdrive` command, then SnapDrive for UNIX uses the format `<nisdomain>\<username>`. For example, `netapp.com\marc`.
- If you are a local user of a UNIX host such as `lnx197-141`, then SnapDrive for UNIX uses the format `<hostname>\<username>`. For example, `lnx197-141\john`.

- If you are an administrator (root) of a UNIX host, then 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 4 lists the predefined roles on the Operations Manager server.

Table 4) Predefined roles on the 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.

7 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 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 using the host commands, create a file system, and mount it. SnapDrive for UNIX achieves all these tasks with one command, reducing the time needed and the errors made during this process. While SnapDrive for UNIX can create storage using a minimum of options, NetApp recommends that you understand the default values and use them appropriately.

SnapDrive for UNIX lets you access LUNs, file systems, disk groups, and logical volumes on a host other than the one used to create it. This might be useful in scenarios in which the storage has to 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 executing the `snapdrive storage connect/disconnect` command.

Examples:

The following command removes LUN mapping from `server1` and exports the disk groups or file systems that LUNs contain.

- Server1:

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

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

- Server2:

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

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

SnapDrive for UNIX also allows you to increase the size of your disk 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 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`, then the junction path should be `/vol1`.

8 Host Partitioning Support in RHEL and SUSE Linux Enterprise Servers

Table 5 provides information about the partitioned devices on Linux.

Table 5) Host partition information.

Operating System	Single Partition	Multiple Partition	Nonpartition Devices	File System or Raw Devices
Red Hat Enterprise Linux 4x or 5x or Oracle Enterprise Linux 4x or 5x	Yes	No	No	ext3 ¹
Red Hat Enterprise Linux 6x	Yes	No	No	ext3 or ext4 ¹
SUSE Linux Enterprise Server 11	Yes	No	No	ext3 ¹
SUSE Linux Enterprise Server 10	No	No	Yes	Ext3 ³
Red Hat Enterprise Linux 4x or Oracle Enterprise Linux 4x	Yes	No	No	ASM with ASMLib ²
Red Hat Enterprise Linux 5x or Oracle Enterprise Linux 5x	Yes	No	Yes	ASM with ASMLib ²
SUSE Linux Enterprise Server 10	No	No	Yes	ASM with ASMLib ³

Note: SUSE Linux Enterprise Server 11 does not work on nonpartitioned devices with SnapManager or SnapDrive. SUSE Linux Enterprise Server 11 with partition devices works only on ext3.

¹ For a non-MPIO environment, enter the following command:

```
sfdisk -uS -f -L -q /dev/ device_name
```

For an MPIO environment, enter the following commands:

```
sfdisk -uS -f -L -q /dev/ device_name
kpartx -a -p p /dev/mapper/ device_name
```

² For a non-MPIO environment, enter the following command:

```
fdisk /dev/device_name
```

For an MPIO environment, enter the following commands:

```
fdisk /dev/mapper/device_name
kpartx -a -p p /dev/mapper/device_name
```

³ Not applicable.

9 Snapshot Copy Management

SnapDrive for UNIX integrates with NetApp Snapshot technology to make stored data reliable to host applications. The ability to create and manage Snapshot copies from the host makes SnapDrive for UNIX attractive to users.

Snapshot copies record the state of the blocks in your file system at a given point of 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 on the commands used to perform these tasks, see the [SnapDrive for UNIX Installation and Administration Guide](#).

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

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

9.1 Best Practices for Snapshot Copy Management

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 using the following commands on the storage system:

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

9.2 Consistent Snapshot Copies

Making Snapshot copies in a SAN environment differs from 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 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, then 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 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, where 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 sequence of actions:

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 a file system `/mnt/fs_multi_vol` resides over LUNs in `storage1:/vol/vol1` and `storage2:/vol/vol1` and Data ONTAP 7.2 or higher 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 the volumes.

Note: `snapdrive snap create -fs /mnt/fs_multi_vol -snapname snap1`. Crash-consistent Snapshot copies are not supported in versions earlier than clustered Data ONTAP 8.2.

Best Practices for Consistent Snapshot Copies

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

SnapDrive simplifies the use of consistency groups. When the file specs or 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 instance, when one of the controllers does not support CG or multiple applications share LUNs from a single volume where the volume I/O fencing might affect other applications, a user can specify the `-nofilerfence` option requesting 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 disk groups using a single command when the file systems or disk groups are independent of each other. NetApp recommends that you use the `-unrelated` option when NFS entities are present with other file specifications in the Snapshot copy. This is because 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 having 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 residing on one storage system and the file system on LUNs residing on a different storage system:

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

9.3 Snapshot Copy Space Management

Snapshot copy 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 at the same time provide up to 255 online point-in-time images.

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

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

Best Practices for Snapshot Copy Space Management

- 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.

9.4 Snap Reserve

Data ONTAP reserves a default of 20% of volume for configuring LUNs or for NFS files to use. 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 WAFL® (Write Anywhere File Layout) technology, snap reserve does not reserve specific physical blocks; rather, it is a logical space-accounting mechanism.

For more information on snap reserve, refer to the [Data ONTAP Block Access Management Guide](#).

Best Practices for Snap Reserve

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

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

If that volume had a 20% snap reserve policy, the volume would require 100GB for the LUN itself, 100GB for the LUN reservation, and an additional 20% for the snap reserve for a total of 220GB.

For more information about understanding fractional reserve, refer to the [Data ONTAP Block Access Management Guide for iSCSI and FCP](#).

10 Restoring a Snapshot Copy Using Volume-Based SnapRestore

For version 3.0 and earlier, SnapDrive for UNIX uses single-file SnapRestore (SFSR) implemented in Data ONTAP. While a restore is in progress for a file using 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 from version 4.0, SnapDrive for UNIX can leverage volume-based SnapRestore technology. This requires less CPU and fewer storage resources and restores instantaneously all 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 Snapshot copy that is being used for the restore operation are deleted. All the new files and new

LUNs created in this volume are also deleted. Any relationship to secondary storage systems will be broken if the restored Snapshot copy is older than the baseline Snapshot copy in the relationship.

10.1 Using Volume-Based SnapRestore

Use the following command to restore a volume 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 snapvg1lvoll -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 snapvg1 -force -reserve -vbsr execute
```

- To set space guarantee as none:

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

Best Practices

- It is a best practice to execute the `-vbsr preview` command before using the `-vbsr execute` command. If the preview option is used, SnapDrive will perform a series of checks on the volume being restored and present a file-by-file analysis of the restore operation before it occurs. The preview analysis will help you decide if you would like to proceed with the volume-based SnapRestore operation.
- Use dedicated volumes for data that will need to be restored quickly using the volume-based SnapRestore method.
- If there are Snapshot copies that were created after the Snapshot copy being used to restore them, then it is a best practice to replicate those Snapshot copies to a secondary storage system using SnapVault® technology and then perform the volume-based SnapRestore operation.

11 Cloning

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

- When there is an available update for the application 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 using 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 from 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, same as normal flexible volumes.
<code>-clone optimal</code>	<code>san-clone-method=optimal</code>	SnapDrive will attempt to create a FlexClone volume. If it is unable to, it will revert 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 in order to create a FlexClone volume:

- The storage system Data ONTAP version is 7.0 or later and FlexClone is licensed.
- The host filespec residing volume is a flexible volume.
- A host filespec can be a file system, host volume, disk group, or 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 will:

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

To use FlexClone 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 will be a SIS clone.

11.1 Benefits of FlexClone

The benefits of FlexClone are:

- Simplified data management and reduced risk.
- Flexibility and greater utilization. You can use FlexClone to create numerous copies of data for additional users without giving them access to the original data.
- 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 volume group, logical volume, and file system:

```
snapdrive snap connect -fs /tmp/test1 /tmp/test2 -destlv lvoll_dup/vg1_dup -snapname btc-ppc-42:/vol/src_snapmirror_vol:snapvg1
```

- The `-autoexpand` option eliminates the need to specify each logical volume or file system to connect the entire disk 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-ppc-42:/vol/src_snapmirror_vol:snapfs2 -
autoexpand -autorename
```

11.2 Best Practice for Cloning

Because you can make several clones from a Snapshot copy, NetApp recommends that you name 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 if they need to be destroyed.

By default, SnapDrive for UNIX assumes that no reservation is required for cloned LUNs and sets the space reservation 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 from SnapDrive for UNIX 3.0, the LUN clone split operation is supported with the `snapdrive snap connect` command. This 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 from SnapDrive for UNIX 4.2, clone split operations can be done using new clone split commands.

Consider the following 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, this can be verified 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. The clone split operation runs in the foreground in synchronous mode. 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 using a Snapshot copy:

```
snapdrive clone split estimate -hostvol vg1/lvol1 -snapname btc-ppc-
42:/vol/src_snapmirror_vol:snapvg1_lvoll -volclone
```

- The following example estimates the storage to split a LUN clone 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
```

- To split 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 off during the Snapshot copy disconnect operation, SnapDrive for UNIX will 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.

12 SnapDrive for UNIX in a 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 FlexVol[®] volumes when done 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 using the iSCSI protocol.
- SnapDrive for UNIX does not support FCP on MultiStore units.
- For SnapDrive for UNIX usage with vFiler, 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 on the following virtualization environments:

- SnapDrive for UNIX is supported on Solaris containers (global/local/native zones) for NFS configuration only.
- SnapDrive for UNIX is supported on IBM AIX LPAR/DLPAR.
- SnapDrive for UNIX is supported on the SLES guest operating system on a Hyper-V[™] VM in Windows[®] for iSCSI and NFS protocols only.
- SnapDrive for UNIX is supported on the Red Hat Enterprise Linux (RHEL) guest operating system on a KVM hypervisor for iSCSI and NFS protocols only.
- SnapDrive for UNIX supports RDM LUNs for ESX[®] inside the guest VM for FC protocol, storage resizing, and all existing operations of cloning, Snapshot copies, and restore. There are no changes with respect to SnapDrive commands. Refer to the RDM LUN support section 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 understands the performance difference between the paths if multiple paths to the LUN exist.

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.

What are the storage system and host configurations that are needed to enable ALUA from SnapDrive for UNIX?

Answer: SnapDrive for UNIX sets ALUA for the igroup and the storage system if the host uses the following multipath solutions:

- AIX native MPIO with the NetApp PCM installed
- Solaris MPXIO
- Linux native MPIO

Storage system requirements are:

- The storage system should have Data ONTAP version 7.2 or later.
- SnapDrive for UNIX requires the storage system to be configured in single image cluster (SSI) mode.

What are the protocols on which SnapDrive for UNIX enables ALUA?

Answer: SnapDrive for UNIX supports enabling ALUA only for FCP configurations. SnapDrive does not support enabling ALUA for an iSCSI configuration.

Is there a SnapDrive configuration variable to enable ALUA?

Answer: From SnapDrive 5.0 for UNIX and later, the configuration variable `enable-alua` is available to control ALUA.

13.3 RDM LUN Support

Raw device mapping (RDM) allows a special file in a 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 virtual machine 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:

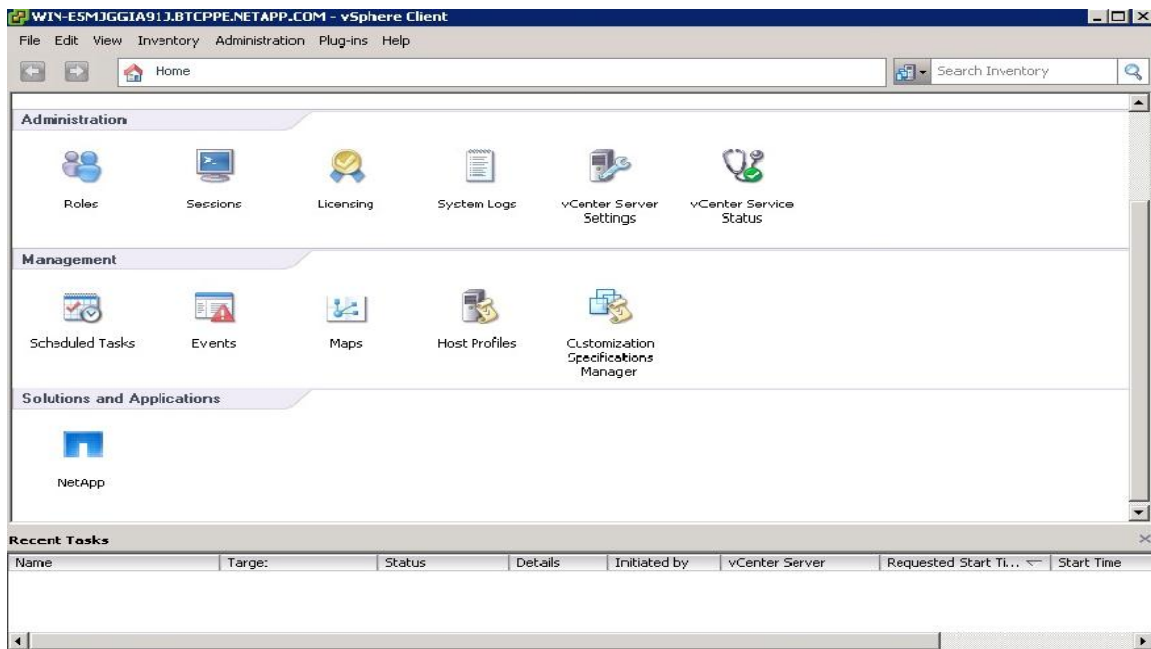
- SnapDrive for UNIX allows creating and provisioning LUNs using the FC protocol on a guest OS in a VMware® virtual environment.
- RDM LUNs for SVMs in clustered Data ONTAP are not supported.
- RDM LUNs provide all functionalities in a guest OS that are feasible in a physical OS using the FC protocol.
- The NetApp Virtual Storage Console (VSC) needs to be installed and configured with vCenter™ running in a Microsoft® Windows machine.
- SnapManager for virtual interface API calls are used to create/delete igroup mapping of the LUN and export/deport of the LUN to the guest OS.
- For more details on host OS, guest OS, 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; hence, 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.

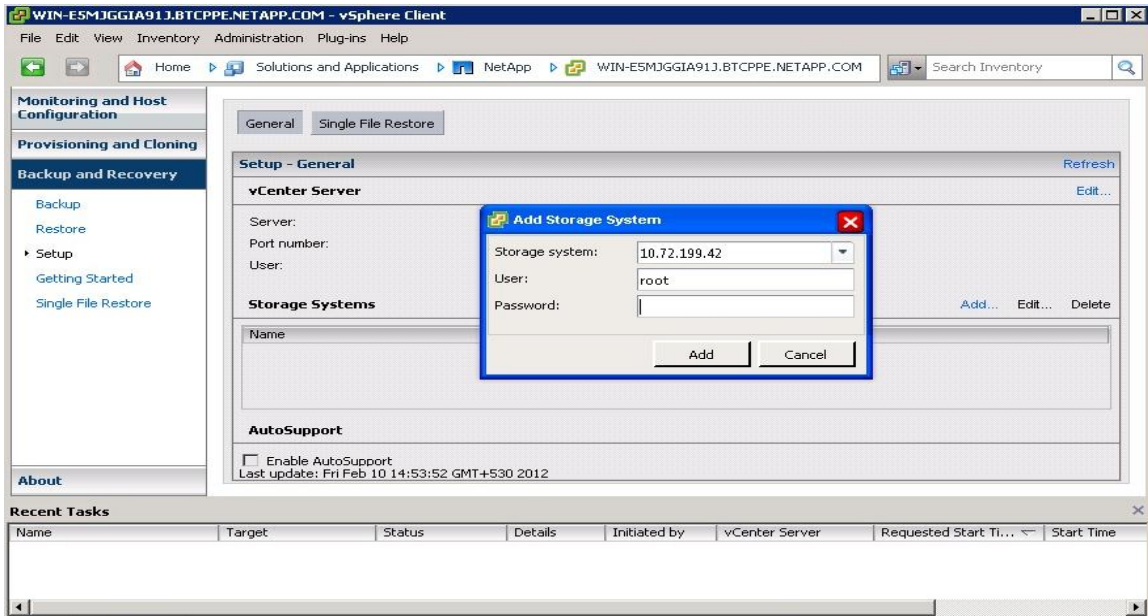
Configuring SnapDrive for UNIX for RDM LUNs

Follow these steps to configure SnapDrive for UNIX for RDM LUNs.


1. After installing the NetApp VSC plug-in on the vCenter, log into 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 vSphere NetApp plug-in. Configure the storage system on 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. Other configuration variables in `snapdrive.conf` are:
- **contact-viadmin-port.** Specifies the port on which SMVI service is running. Default is 8043.
 - **use-https-to-viadmin.** Specifies the protocol (HTTP or HTTPS) `https` for communicating with SMVI. Default is `on` (HTTPS for communication).
 - **virtualization-operation-timeout-secs.** Specifies the number of seconds to wait before timing out a call to SMVI. Default is 600 seconds.

- **vif-password-file.** Specifies the path for the virtual interface password file. Default is /opt/NetApp/snapdrive/.vifpw.

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

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

Examples:

```
# 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 where RDM LUNs are to be created.

RDM LUN Creation

The following actions are performed while creating RDM LUNs:

- SDU creates the LUN on the storage.
- SDU creates igroup mappings on storage, creates RDM mappings, and exports the LUN to the guest.
- SDU discovers the LUN and creates the host-side file specs.

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

RDM LUN Deletion

The following actions are performed while deleting RDM LUNs:

```
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 file system /mnt1
- fs /mnt1 ... deleted
deporting lun(s) from Guest OS ... done
- LUN btc-ppe-42:/vol/esxvoll/lun1 ... deleted
```

Limitations

- RDM LUNs for SVMs in clustered Data ONTAP are not supported.
- When SDU commands are executed with igroup, SDU will ignore the user-defined igroup with a warning message.

To migrate a VM using VMware vMotion[®] across the server, perform the following steps:

- a. Before migration, disconnect the exported filespec/LUN from the guest OS.
- b. Perform VM migration using VMware vMotion.
- c. After VM migration, connect the filespec/LUN to the guest OS.
- 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.
- Sometimes, you might realize that the RDM LUN was not created and behaves as an ordinary LUN. When this happens, verify that the default transport variable is set to `FCP` in the `snapdrive config` file. Execute 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 located in the ESX server that are generated due to SnapDrive operations.

14 SnapDrive for UNIX FAQ

14.1 SnapDrive for UNIX Installation FAQ

Why do I see errors during the installation process?

Answer: Verify that you have root privileges. In addition, check the [Compatibility Matrix](#) to confirm that you are using supported platforms and processor architecture.

Do I need to uninstall my existing version of SnapDrive for UNIX before upgrading?

Answer: No. The upgrade process overwrites the existing binaries. But you should stop the SnapDrive daemon before upgrading.

Does the upgrade process affect any of my existing LUNs?

Answer: No.

After upgrading, can I restore a Snapshot copy that was created using an older version?

Answer: Yes. Snapshot copy management functionality is compatible with older versions, but the LUN cloning feature in Linux, which is supported beginning with SnapDrive for UNIX 2.2, does not support cloning of LVM disk groups to the same host based on Snapshot copies created by versions earlier than 2.2.

When I upgrade SnapDrive for UNIX, is my existing `snapdrive.conf` file overwritten?

Answer: No. When you install a new version, SnapDrive for UNIX checks to see if you already have a version installed. If it detects a SnapDrive for UNIX installation, it preserves the existing `snapdrive.conf` file and renames it to the version of the file that it is installing: for example, `snapdrive.conf.5.0`. The upgrade process avoids overwriting your existing `snapdrive.conf` file, so you do not lose any customizations you made to the file.

By default, SnapDrive for UNIX includes comment text about the variables in the `snapdrive.conf` file and automatically uses the default values for all variables except those you have customized. As a result,

SnapDrive for UNIX uses the default values for variables that are new, even if you have an earlier version of the `snapdrive.conf` file. If you want to change these values, you must add the variables to your existing `snapdrive.conf` file and specify the values that you want.

You can use the `snapdrive config show` command to view the existing configuration.

What can we do if SDU 5.2 fails to start because of reported FD leaks in the `vgdisplay` command on the Linux OS?

Answer: As part of NetApp AutoSupport implementation, SDU queries the entire disk group to figure out the storage. Because of this, Linux reports leaked file descriptors in the `vgdisplay` command.

To suppress this, set the `LVM_SUPPRESS_FD_WARNINGS` environment variable to 1:

export `LVM_SUPPRESS_FD_WARNINGS=1`.

Example:

The SDU log reports the following FD leak errors:

- 16:35:40 02/23/12 [f7476b90]v,9,0,ASSISTANT EXECUTION (at 988.004188): `vgdisplay` File descriptor 6 (`/var/log/sd-trace.log`) leaked on `vgdisplay` invocation. Parent PID 11013
- 16:35:40 02/23/12 [f7476b90]v,9,0,ASSISTANT EXECUTION (at 988.006523): Output:
- 16:35:40 02/23/12 [f7476b90]j,2,6,FileSpecOperation::findDiskGroupsByDevices
- Volume manager error (6) while identifying all active disk groups could not find a disk group by that name.

What do we do if the error "WARNING!!! Unable to find a SAN storage stack" appears when installing SnapDrive for UNIX?

Answer: Verify that the appropriate transport protocol, volume manager, file system, and multipathing type are installed and configured in the system. If NFS is being used, this warning message can be ignored. Verify that the appropriate transport type, volume manager, file system, and multipathing types are properly installed and configured in the system.

Does SnapDrive for UNIX support AutoSupport?

Answer: The AutoSupport feature is supported in SnapDrive for UNIX for clustered Data ONTAP 8.2 and later. SnapDrive for UNIX for Data ONTAP 7-Mode has support for the NetApp AutoSupport feature.

This feature for clustered Data ONTAP and Data ONTAP 7-Mode can be enabled by using the following option in the `snapdrive.conf` file located in the host system:

- `autosupport-enabled=on #`

This option enables AutoSupport (requires `autosupport-filer` be set).

I am planning to have a host cluster environment. Do I need to install SnapDrive for UNIX on each host manually, or is there a tool that supports installation across multiple hosts?

Answer: In a cluster environment, you have to install SnapDrive for UNIX manually on each host, because no tool or management software is available from NetApp to support installation of software across multiple hosts.

However, you can write scripts easily in addition to quickly installing SnapDrive on the cluster nodes. In the following example, assuming an IBM AIX operating system, SMIT can be used to create a script to install SnapDrive for UNIX on all of the cluster nodes at once:

1. Write the script output from the `smit -s /tmp/install_sd.ksh, -s file to /tmp/install_sd.ksh`.

2. Mount the SnapDrive for UNIX installation directory on a shared mount point.
3. Install SnapDrive for UNIX on the host using SMIT, as explained in the [SnapDrive for UNIX Installation and Administration Guide](#).
4. After you successfully install SnapDrive for UNIX, exit the SMIT tool and copy the `install_sd.ksh` file to all cluster nodes and execute it from the host CLI.

What is the default location of the `snapdrive.conf` file?

Answer: You can find the `snapdrive.conf` file in the default installation directory. The default installation directory changes across platforms. See Table 7 for more information.

Table 7) Default locations for `snapdrive.conf` file.

Host Operating System	Path
AIX and Linux	/opt/NetApp/snapdrive/snapdrive.conf
Solaris	/opt/NTAP/snapdrive/snapdrive.conf

For more information on the `snapdrive.conf` file, see the [SnapDrive for UNIX Installation and Administration Guide](#).

14.2 Storage Provisioning FAQ

Which file system should I specify when provisioning storage?

Answer: NetApp does not recommend using any specific file system when creating the storage. You should decide which file system might be best for the application. Refer to the [Compatibility Matrix](#) to understand which file systems are supported by SnapDrive for UNIX.

Before I use the `SnapDrive storage create` command to provision an iSCSI, do I have to make sure that the host already has the target storage system in its iSCSI discovery path, or will SnapDrive for UNIX do that automatically?

Answer: You have to manually enter the target file location on the host system before you provision an iSCSI. The target file varies, depending on the operating system. Table 8 lists the target information and steps involved for different host operating systems.

1. Edit the iSCSI configuration file.

Table 8) Target information for host operating systems.

Host Operating System	File Location	Syntax
AIX	/etc/iscsi/targets	<storage_system_IPAddress/storage_system_name><port number, default is 3260> <iscsi nodename of the storage system> iscsiadm add static=config <iscsi nodename of the storage system>,<storage_system_IPAddress/storage_system_name>
Linux (RHEL and SUSE)	/etc/iscsi.conf	HeaderDigest=no DataDigest=no ImmediateData=no DiscoveryAddress=<storage_system_IPAddress/storage_system_name> The HeaderDigest, Data Digest, and ImmediateData entries are required for faster discovery.
Linux (RHEL5 and SUSE 10 and later)		iscsiadm -m discovery-address -p <ipaddress>:<portno> -t st -l

Host Operating System	File Location	Syntax
Solaris	--	iscsiadm add discovery-address <storage_system_IPaddress/storage_system_name> Or Iscsiadm add static-config <iscsi nodename of the storage system>,<new storage_system_IPaddress/storage_system_name>

2. Execute the appropriate command from Table 9 to rescan the iSCSI targets.

Table 9) Rescan iSCSI targets.

Host Operating System	Command
AIX	cfgmgr
Linux (RHEL and SUSE)	/sbin/iscsi-rescan /etc/init.d/iscsi reload
Linux (RHEL5 and SUSE 10 and later)	iscsiadm -R
Solaris	cfgadm

3. Verify the host iSCSI initiator node name in the command output by using the `iscsi show initiator` command from the storage system CLI.

Are there any special considerations of which I should be aware when using SnapDrive in a cluster environment?

Answer: Yes. Make sure that the SnapDrive for UNIX versions and the `snapdrive.conf` files are the same on all of the cluster nodes. NetApp recommends that you execute the SnapDrive for UNIX commands from the master node.

Can SnapDrive resize the disk group and reflect it in the host?

Answer: Not all volume managers support this feature. Use the `-addlun` option to resize the disk group.

Can I connect to the LUN created using System Manager or storage system CLI commands?

Answer: Yes, by using the `snapdrive storage connect` command.

Can I migrate between a cluster and a nonclustered environment: that is, between shared and dedicated file systems or disk groups?

Answer: Yes, by using the `snapdrive storage connect` or `snapdrive snap connect/restore` commands.

Why does SnapDrive produce an error when I try to use the `snapdrive storage create/connect` command?

Answer: There are many reasons for this error, but for basic troubleshooting you can use the information provided in Table 10.

Table 10) Steps to resolve SnapDrive storage create/connect command errors.

If...	Then...
SnapDrive storage create fails during LUN creation	Follow the steps in section 5.4, "Using the SnapDrive Daemon," of this report to verify reachability from the host to the storage system. If HTTPS is enabled in SnapDrive, make sure the storage system is

If...	Then...
	configured to accept HTTPS requests.
LUN discovery fails	Follow the steps in sections 5.1, "Checklist for Storage System," and 5.4, "Using the SnapDrive Daemon," of this report to confirm that the FCP or iSCSI service has been started and configured correctly.

1. Verify that the host environment is supported by SnapDrive for UNIX; see the [Interoperability Matrix Tool](#).
2. Follow the steps in section 5.1, "Checklist for Storage System," of this report to make sure storage system configurations are correct.
3. Follow the steps in section 5.2, "Checklist for Host Before SnapDrive Installation," of this report to make sure the host configurations are correct.
4. Verify configuration `snapdrive.conf` and make sure it is configured for your environment's settings.

Which entity (either cluster name or SVM) should be mentioned while configuring storage operating in clustered Data ONTAP with SnapDrive for UNIX?

Answer: You need to mention the SVM name while configuring clustered Data ONTAP storage with SDU.

Command:

```
snapdrive config set <username> <servername>
```

Example:

```
snapdrive config set vsadmin toaster-vserver
```

What can I do if RBAC is enabled and SnapDrive for UNIX operations fail occasionally due to an HTTP post error from the DataFabric® Manager server?

Answer: Increase the `dfm-api-timeout` variable's value in the `snapdrive.conf` file and restart the SnapDrive for UNIX daemon.

14.3 Snapshot Copy Management FAQ

How frequently should I create Snapshot copies?

Answer: The frequency mainly depends on your environment's specific requirements.

Can I schedule the creation of Snapshot copies using SnapDrive?

Answer: No. SnapDrive does not support scheduling Snapshot copy creation. However, you can achieve the same result by manually scripting `snapdrive snap` commands and using `cron` jobs in UNIX environments.

I accidentally deleted the file system on my host. Can I restore it using the Snapshot copy I made previously using SnapDrive for UNIX?

Answer: Yes. It is possible by using the `snapdrive snap restore` command.

Can I restore the OS file system from a Snapshot copy created by the storage system?

Answer: No. SnapDrive allows you to restore file systems only from the Snapshot copies created by SnapDrive.

Should I manually create igroups in a cluster environment when I use `snapdrive storage create`?

Answer: No. igroups need not be created manually. They are created by using the `snapdrive storage create` command.

I get the following error message when connecting to the LUN that is part of a volume/disk group, but the SnapDrive `snap list` output shows the presence of the LUN: “The following LUNs were not found.” How do I resolve this error?

Answer: SnapDrive for UNIX does not allow connecting to a LUN that is created as part of a volume group. You can connect to a Snapshot copy containing a LUN when the Snapshot copy is made only for the LUN, for example, by using the `snapdrive snap connect -lun` command.

How can I create FlexClone volumes using SnapDrive for UNIX?

Answer: You have to pass `-clone <unrestricted or optimal>` through CLI or set the following variable in the `snapdrive.conf` file:

```
san-clone-method=unrestricted | optimal (By default it has lunclone.)
```

14.4 SnapDrive for UNIX in a Storage Foundation for Oracle RAC (SFRAC) Environment (Storage Creation) FAQ

Why am I not able to use the `-nolvm` option during storage creation in an SFRAC environment?

Answer: The `-nolvm` option is not valid for clusterwide or shared resources because the operation requires the CVM to manage shared disk groups.

Why does SnapDrive for UNIX take longer to complete operations in an SFRAC environment?

Answer: This can happen if you provide many `file_specs` with the SnapDrive command in a cluster environment that has many nodes. The delay is mostly due to SnapDrive and cluster requirements to communicate with all cluster nodes during any operation involving a shared disk group or file system. To avoid this delay, NetApp recommends that you provide fewer `file_specs` (shared disk groups/file systems) in one SnapDrive command.

In the VxDMP environment on Solaris, some SnapDrive operations might fail if the paths from the host to the LUNs are changed. For example, if you created LUNs by enabling the primary path and disabling the secondary path, and, subsequently, if you enable the secondary path and then disable the primary path, the SnapDrive operations on those LUNs fail because VxDMP does not detect this change. To function properly, VxDMP requires the device path information to be updated after any change is made in the port state (enabled or disabled). If you change the path to enabled or disabled, NetApp recommends that you execute `devfsadm -Cv` followed by `vxdisk scandisks` manually on the host to refresh the VxDMP path information before any SnapDrive operation is executed.

14.5 General Queries

Can I have more than one data interface for a port set?

Answer: Yes. You can add more than one data interface for a port set using the `portset add` command on the storage system.

Does `portset` support iSCSI and FC protocols?

Answer: Yes. You can create portset of type iSCSI and FC. However, the mixed mode (both iSCSI and FC) is not supported by SnapDrive.

Is it possible to map a portset to multiple igroups?

Answer: Yes. You can map a portset to multiple igroups to limit the number of paths available to the host.

Is a portset created automatically during SVM creation?

Answer: No. A portset has to be created manually depending on your requirement. Creating a portset is an optional requirement.

15 Summary

SnapDrive for UNIX is a complete storage management solution. It helps administrators to execute nearly instantaneous Snapshot copy backups and restorations of application data. It also provisions storage from the host system as required by each application.

The recommendations made in this guide are best practices for most environments. Make sure to familiarize yourself with the prerequisites provided at the beginning of this guide.

References

- SnapDrive for UNIX Installation and Administration Guide
<http://support.netapp.com/documentation/docweb/index.html?productID=61594>
- Clustered Data ONTAP 8.2 Documentation
<http://support.netapp.com/documentation/docweb/index.html?productID=61651>
- Data ONTAP Block Access Management Guide for iSCSI and FCP
<https://support.netapp.com/documentation/productsatoz/index.html>
- Host Utilities Installation and Setup Guide
<https://support.netapp.com/documentation/productlibrary/index.html?productID=61343>
- Data ONTAP Storage Management Guide
<https://support.netapp.com/documentation/productsatoz/index.html>
- Compatibility and Configuration Guide for NetApp FCP and iSCSI
<https://support.netapp.com/documentation/productsatoz/index.html>
- Data ONTAP File Access and Protocols Management Guide
<https://support.netapp.com/portal/documentation>

Version History

Version	Date	Document Version History
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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

[Go further, faster®](#)

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, DataFabric, Data ONTAP, FlexClone, FlexVol, MultiStore, SnapDrive, SnapManager, SnapRestore, Snapshot, SnapVault, vFiler, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Oracle is a registered trademark of Oracle Corporation. ESX, vMotion, VMware, and VMware vSphere are registered trademarks and vCenter is a trademark of VMware, Inc. UNIX is a registered trademark of The Open Group. Linux is a registered trademark of Linus Torvalds. Microsoft and Windows are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. Symantec and Veritas are trademarks of Symantec Corporation. SAP is a registered trademark of SAP AG. Cisco is a registered trademark of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4212-0813



www.netapp.com