

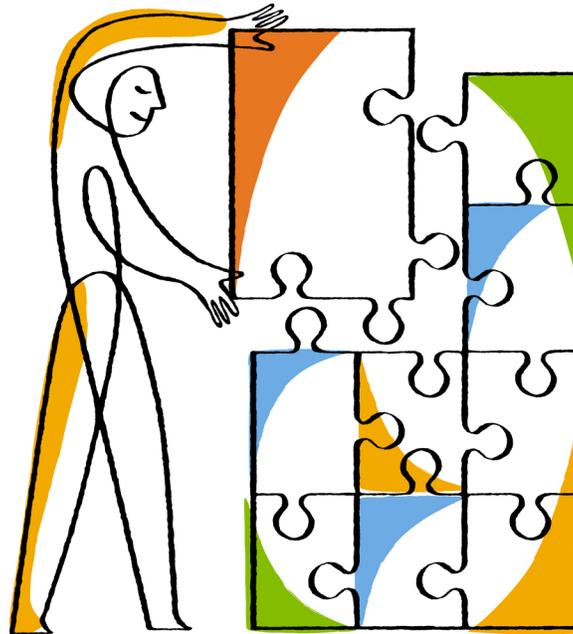


NetApp®

SnapManager® 3.4 for Oracle®

Installation and Setup Guide for UNIX®

For Clustered Data ONTAP®



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-10240_A0
November 2015

Contents

Product overview	4
SnapManager architecture	4
Deployment workflow	7
Preparing for deployment	8
SnapManager licensing	8
Supported configurations	9
Supported storage types	9
UNIX host requirements	10
Configuring databases	11
Configuring the target database	11
Creating an Oracle database user for the target database	12
Creating the repository database instance	12
Creating an Oracle user for the repository database instance	13
Verifying the Oracle listener configuration	13
Installing SnapManager	14
Setting up SnapManager	16
Starting the SnapManager server	16
Accessing the SnapManager user interface	16
Configuring the SnapManager repository	17
Preparing storage systems for SnapMirror and SnapVault	
replication	19
Understanding the differences between SnapMirror and SnapVault	19
Preparing storage systems for SnapMirror replication	19
Preparing storage systems for SnapVault replication	21
Backing up and verifying your databases	24
SnapManager backup overview	24
Defining a backup strategy	24
Creating a profile for your database	26
Backing up your database	28
Verifying database backups	30
Scheduling recurring backups	30
Where to go next	33
Copyright information	34
Trademark information	35
How to send comments about documentation and receive update	
notifications	36
Index	37

Product overview

SnapManager for Oracle automates and simplifies the complex, manual, and time-consuming processes associated with the backup, recovery, and cloning of Oracle databases. You can use SnapManager with Data ONTAP SnapMirror technology to create copies of backups on another volume and with Data ONTAP SnapVault technology to archive backups efficiently to disk.

SnapManager integrates with native Oracle technologies such as Oracle Real Application Clusters (Oracle RAC), Automatic Storage Management (ASM), and Direct NFS across FC, iSCSI, and NFS protocols. Optionally, backups created by using SnapManager can be cataloged with Oracle Recovery Manager (RMAN) to preserve the backup information; these backups can be used later in block-level restore or tablespace point-in-time recovery operations.

SnapManager highlights

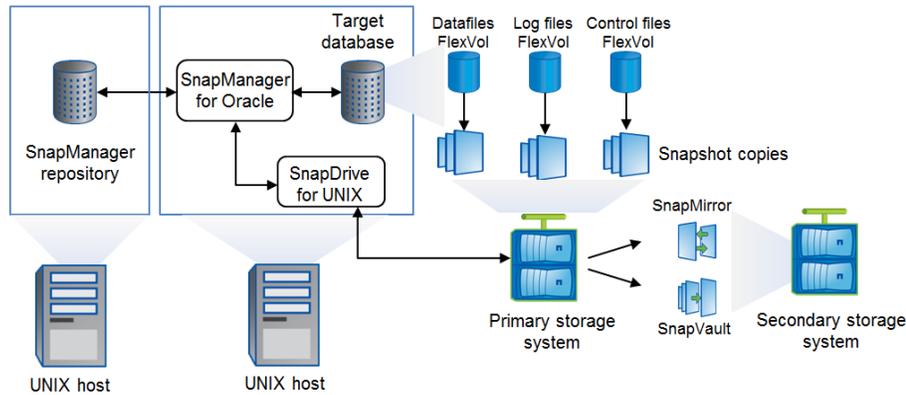
SnapManager features seamless integration with Oracle databases on the UNIX host and with NetApp Snapshot, SnapRestore, and FlexClone technologies on the back end. It offers an easy-to-use user interface (UI) as well as command-line interface (CLI) for administrative functions.

SnapManager enables you to perform the following database operations and manage data efficiently:

- Creating space-efficient backups on primary or secondary storage
SnapManager enables you to back up the data files and archive log files separately.
- Scheduling backups
- Restoring full or partial databases by using a file-based or volume-based restore operation
- Recovering databases by discovering, mounting, and applying archive log files from backups
- Pruning archive log files from archive log destinations when creating backups of only the archive logs
- Retaining a minimum number of archive log backups automatically by retaining only the backups that contain unique archive log files
- Tracking operation details and generating reports
- Verifying backups to ensure that backups are in a valid block format and that none of the backed-up files are corrupted
- Maintaining a history of operations performed on the database profile
A profile contains information about the database to be managed by SnapManager.
- Protecting backups on the secondary storage systems.
- Creating space-efficient clones of backups on primary or secondary storage
SnapManager enables you to split a clone of a database.

SnapManager architecture

SnapManager for Oracle includes components that work together to provide a comprehensive and powerful backup, restore, recovery, and cloning solution for Oracle databases.



SnapDrive for UNIX

SnapManager requires SnapDrive to establish connection with the storage system. You must install SnapDrive for UNIX on every target database host before installing SnapManager.

SnapManager for Oracle

You must install SnapManager for Oracle on every target database host.

You can either use the command-line interface (CLI) or UI from the database host where SnapManager for Oracle is installed. You can also use the SnapManager UI remotely by using a web browser from any system running on an operating system supported by SnapManager.

Note: The supported JRE versions are 1.5, 1.6, and 1.7.

Target database

The target database is an Oracle database that you want to manage using SnapManager by performing backup, restore, recovery, and clone operations.

The target database can be a standalone, Real Application Clusters (RAC), or reside on Oracle Automatic Storage Management (ASM) volumes. For details about the supported Oracle database versions, configurations, operating systems, and protocols, see the NetApp Interoperability Matrix Tool.

SnapManager repository

The SnapManager repository resides in an Oracle database and stores metadata about profiles, backups, restore, recovery, and clone. A single repository can contain information about operations performed on multiple database profiles.

The SnapManager repository cannot reside in the target database. The SnapManager repository database and the target database must be online before performing SnapManager operations.

Primary storage system

SnapManager backs up the target databases on the primary NetApp storage system.

Secondary storage system

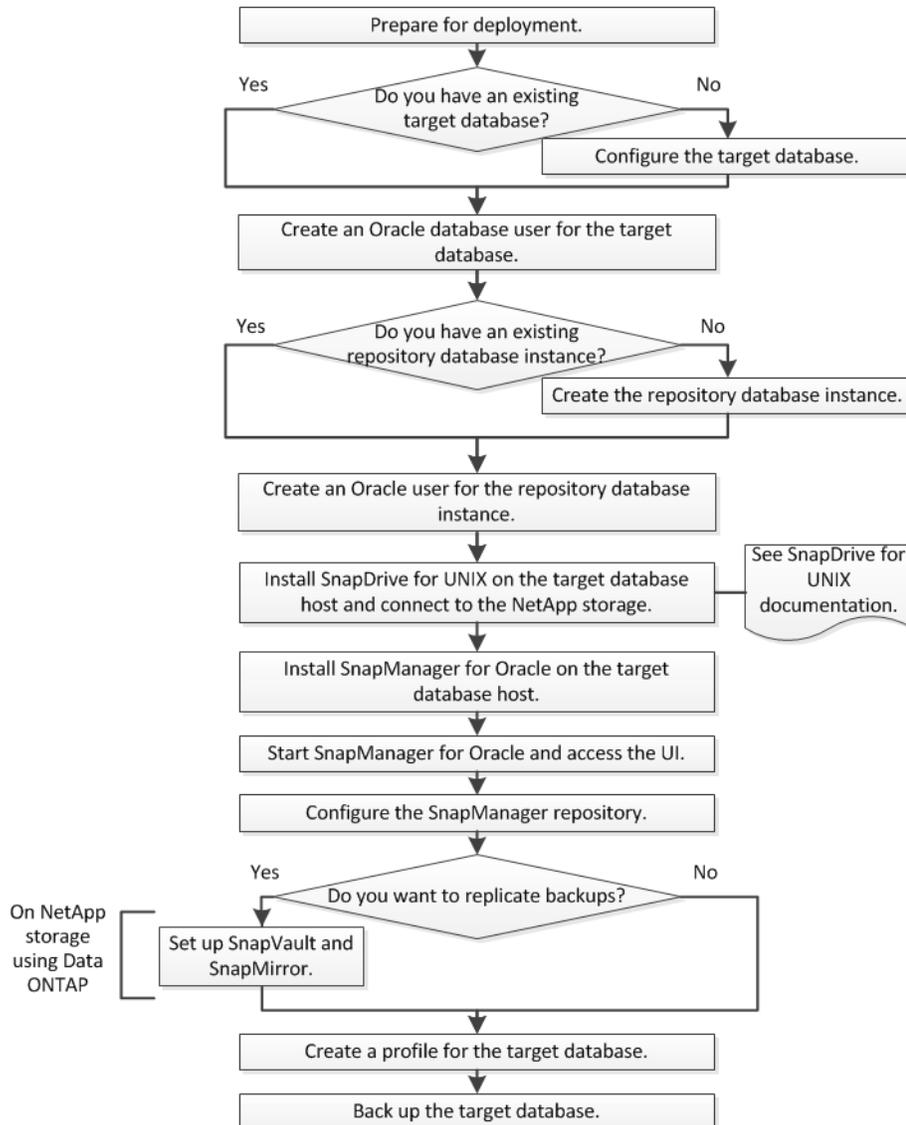
When you enable data protection on a database profile, the backups created on the primary storage system by SnapManager are replicated to a secondary NetApp storage system using SnapVault and SnapMirror technologies.

Related information

[NetApp Interoperability Matrix Tool](#)

Deployment workflow

Before you can create backups with SnapManager, you need to first install SnapDrive for UNIX and then install SnapManager for Oracle.



Preparing for deployment

Before you deploy SnapManager, you must ensure your storage system and UNIX hosts meet the minimum resource requirements.

Steps

1. Verify that you have the required licenses.
2. Verify the supported configurations.
3. Verify the supported storage types.
4. Verify that your UNIX hosts meet SnapManager requirements.

Related references

[SnapManager licensing](#) on page 8

[Supported configurations](#) on page 9

[Supported storage types](#) on page 9

[UNIX host requirements](#) on page 10

SnapManager licensing

A SnapManager license and several storage system licenses are required to enable SnapManager operations. The SnapManager license is available in two licensing models: *per-server licensing*, in which the SnapManager license resides on each database host; and *per-storage system licensing*, in which the SnapManager license resides on the storage system.

The SnapManager license requirements are as follows:

License	Description	Where required
SnapManager per-server	A host-side license for a specific database host. Licenses are required only for database hosts on which SnapManager is installed. No SnapManager license is required for the storage system.	On the SnapManager host. A SnapManager license is not required on primary and secondary storage systems when using per-server licensing.
SnapManager per-storage system	A storage-side license that supports any number of database hosts. Required only if you are not using a per-server license on the database host.	On primary and secondary storage systems.
SnapRestore	A required license that enables SnapManager to restore databases.	On primary and secondary storage systems. Required on SnapMirror destination systems to perform remote verification. Required on SnapVault destination systems to perform remote verification and to restore from a backup.

License	Description	Where required
FlexClone	An optional license for cloning databases.	On primary and secondary storage systems. Required on SnapVault destination systems when creating clones from a backup.
SnapMirror	An optional license for mirroring backups to a destination storage system.	On primary and secondary storage systems.
SnapVault	An optional license for archiving backups to a destination storage system.	On primary and secondary storage systems.
Protocols	NFS, iSCSI, or FC license is required depending on the protocol used.	On primary and secondary storage systems. Required on SnapMirror destination systems to serve data if a source volume is unavailable.

Supported configurations

The hosts on which you are installing SnapManager must meet the specified software, browser, database, and operating system requirements. You must verify support for your configuration before you install or upgrade SnapManager.

For information about supported configurations, see the Interoperability Matrix tool.

Related information

[NetApp Interoperability Matrix Tool](#)

Supported storage types

SnapManager supports a wide range of storage types on both physical and virtual machines. You must verify support for your storage type before you install or upgrade SnapManager.

Machine	Storage type
Physical server	<ul style="list-style-type: none"> NFS-connected volumes FC-connected LUNs iSCSI-connected LUNs
VMware ESX	<ul style="list-style-type: none"> NFS volumes connected directly to the guest system RDM LUNs on the guest operating system

UNIX host requirements

You must install SnapManager for Oracle on every host where an Oracle database you want to backup is hosted. You must ensure that your hosts meet the minimum requirements for SnapManager configuration.

- You must install SnapDrive on the database host before you install SnapManager.
- You can install SnapManager either on physical or virtual machines.
- You must install the same version of SnapManager on all hosts that share the same repository.
- You must install Oracle patch 13366202 if you are using Oracle databases 11.2.0.2 or 11.2.0.3. If you are using DNFS, you must also install the patches listed in the My Oracle Support (MOS) report 1495104.1 for maximum performance and stability.

Configuring databases

You must configure at least two Oracle databases: a target database that you want to back up using SnapManager; and a repository database to store the target database metadata. The target database and the SnapManager repository database must be configured and online before performing SnapManager operations.

Steps

1. [Configuring the target database](#) on page 11
The target database is an Oracle database that can be configured either as standalone, Real Application Clusters (RAC), Automatic Storage Management (ASM), or any other supported combinations.
2. [Creating an Oracle database user for the target database](#) on page 12
An Oracle database user is required to log in to the database and perform SnapManager operations. You must create this user with the `sysdba` privilege if a user with the `sysdba` privilege does not exist for the target database.
3. [Creating the repository database instance](#) on page 12
The repository database instance is an Oracle database in which you create the SnapManager repository. The repository database instance must be a stand-alone database and cannot be the target database.
4. [Creating an Oracle user for the repository database instance](#) on page 13
An Oracle user is required to log in to and access the repository database instance. You must create this user with `connect` and `resource` privileges.
5. [Verifying the Oracle listener configuration](#) on page 13
The listener is a process that listens for client connection requests. It receives incoming client connection requests and manages the traffic of these requests to the database. Before connecting to a target database or repository database instance, you can use the `STATUS` command to verify the listener configuration.

Configuring the target database

The target database is an Oracle database that can be configured either as standalone, Real Application Clusters (RAC), Automatic Storage Management (ASM), or any other supported combinations.

Step

1. Configure the target database by referring *TR-3633*.

Related information

[NetApp Technical Report 3633: Best Practices for Oracle Databases on NetApp Storage](#)

Creating an Oracle database user for the target database

An Oracle database user is required to log in to the database and perform SnapManager operations. You must create this user with the *sysdba* privilege if a user with the *sysdba* privilege does not exist for the target database.

About this task

SnapManager can use any Oracle user with the *sysdba* privilege that exists for the target database. For example, SnapManager can use the default *sys* user. However, even if the user exists, you can create a new user for the target database and assign the *sysdba* privilege.

You can also use the OS authentication method wherein the operating system (OS) allows the Oracle database to use the credentials that are maintained by the OS to authenticate users to log in to the database and perform SnapManager operations. If you are authenticated by the OS, you can connect to the Oracle database without specifying a user name or password.

Steps

1. Log in to SQL *Plus:


```
sqlplus '/ as sysdba'
```
2. Create a new user with an administrator password:


```
create user user_name identified by admin_password;
```

user_name is the name of the user you are creating and *admin_password* is the password that you want to assign to the user.
3. Assign the sysdba privilege to the new Oracle user:


```
grant sysdba to user_name;
```

Creating the repository database instance

The repository database instance is an Oracle database in which you create the SnapManager repository. The repository database instance must be a stand-alone database and cannot be the target database.

Before you begin

You must have an Oracle database and a user account to access the database.

Steps

1. Log in to SQL *Plus:


```
sqlplus '/ as sysdba'
```
2. Create a new tablespace for the SnapManager repository:


```
create tablespace tablespace_name datafile '/u01/app/oracle/oradata/
datafile/tablespace_name.dbf' size 100M autoextend on;
```

tablespace_name is the name of the tablespace.
3. Verify the block size of the tablespace:


```
select tablespace_name, block_size from dba_tablespaces;
```

SnapManager requires a minimum 4-K block size for the tablespace.

Related information

[NetApp Technical Report 3761: SnapManager for Oracle: Best Practices](#)

Creating an Oracle user for the repository database instance

An Oracle user is required to log in to and access the repository database instance. You must create this user with *connect* and *resource* privileges.

Steps

1. Log in to SQL *Plus:

```
sqlplus '/ as sysdba'
```

2. Create a new user and assign an administrator password to that user:

```
create user user_name identified by admin_password default tablespace
tablespace_name quota unlimited on tablespace_name;
```

- *user_name* is the name of the user you are creating for the repository database.
- *admin_password* is the password you want to assign to the user.
- *tablespace_name* is the name of the tablespace created for the repository database.

3. Assign *connect* and *resource* privileges to the new Oracle user:

```
grant connect, resource to user_name;
```

Verifying the Oracle listener configuration

The listener is a process that listens for client connection requests. It receives incoming client connection requests and manages the traffic of these requests to the database. Before connecting to a target database or repository database instance, you can use the *STATUS* command to verify the listener configuration.

About this task

The *STATUS* command displays basic status information about a specific listener, including a summary of listener configuration settings, listening protocol addresses, and a summary of services registered with that listener.

Step

1. Enter the following command at the command prompt:

```
lsnrctl STATUS
```

The default value assigned to the listener port is 1521.

Installing SnapManager

You must install SnapManager on each host where the database you want to backup is running.

Before you begin

You must have installed SnapDrive for UNIX on the database host and established a connection to the storage system.

For information about how to install SnapDrive and establish connection to storage system, see SnapDrive for UNIX documentation.

About this task

You must install one SnapManager instance per database host. If you are using a Real Application Cluster (RAC) database and want to back up the RAC database, you must install SnapManager on all the hosts of the RAC database.

Steps

1. Download the SnapManager for Oracle install package for UNIX from the NetApp Support Site and copy it to the host system.

[NetApp Downloads: Software](#)

2. Log in to the database host as the root user.
3. From the command prompt, navigate to the directory where you copied the install package.
4. Make the install package executable:

```
chmod 755 install_package.bin
```
5. Install SnapManager:

```
./ install_package.bin
```
6. Press `Enter` to continue.
7. Perform the following steps:

- a. Press `Enter` to accept the default value for the operating system user.

The default value for the user is *oracle*.

- b. Press `Enter` to accept the default value for operating system group.

The default value for the group is *dba*.

- c. Press `Enter` to accept the default value for the startup type.

The configuration summary is displayed.

8. Review the configuration summary and press `Enter` to continue.

SnapManager is installed at `/opt/NTAPsmo` for Solaris and `/opt/NetApp/` for all other UNIX hosts.

Related tasks

[Setting up SnapManager](#) on page 16

Related information

[NetApp Documentation: SnapDrive for UNIX](#)

Setting up SnapManager

You can start SnapManager and access it by using either the user interface (UI) or the command-line interface (CLI). After accessing SnapManager, you must create the SnapManager repository before performing any SnapManager operations.

Steps

1. [Starting the SnapManager server](#) on page 16
You must start the SnapManager server from the target database host.
2. [Accessing the SnapManager user interface](#) on page 16
You can access the SnapManager user interface (UI) remotely by using a web browser from any system running on an operating system supported by SnapManager. You can also access the SnapManager UI from the target database host by running the `smogui` command.
3. [Configuring the SnapManager repository](#) on page 17
You must configure the SnapManager repository in the repository database instance. The repository database stores metadata for databases managed by SnapManager.

Starting the SnapManager server

You must start the SnapManager server from the target database host.

Step

1. Log in to the target database host and start the SnapManager server:

```
smo_server start
```

The following message is displayed: `SnapManager Server started on secure port port_number with PID PID_number.`

Note: The default port is `27214`.

After you finish

You can verify that SnapManager is running correctly:

```
smo system verify
```

The following message is displayed: `Operation Id operation_ID_number succeeded.`

Accessing the SnapManager user interface

You can access the SnapManager user interface (UI) remotely by using a web browser from any system running on an operating system supported by SnapManager. You can also access the SnapManager UI from the target database host by running the `smogui` command.

Before you begin

- You must ensure that SnapManager is running.
- You must ensure that the supported operating system and Java are installed on the system where you want to access the SnapManager UI.

For information about the supported operating system and Java, see the Interoperability Matrix tool.

Steps

1. In the web browser window, enter the following:

`https://server_name.domain.com:port_number`

- `server_name` is the name of the target database host where SnapManager is installed. You can also enter the IP address of the target database host.
- `port_number` is the port on which SnapManager is running. The default value is 27214.

2. Click the **Launch SnapManager for Oracle** link.

The SnapManager for Oracle UI is displayed.

Configuring the SnapManager repository

You must configure the SnapManager repository in the repository database instance. The repository database stores metadata for databases managed by SnapManager.

Before you begin

- You must have created the repository database instance.
- You must have created the Oracle user for the repository database instance with required privileges.
- You must have included the repository database instance details in the `tnsnames.ora` file.

About this task

You can configure the SnapManager repository either from the SnapManager user interface (UI) or command-line interface (CLI). These steps show how to create a repository using the SnapManager UI. You can also use the CLI if you prefer.

For information about how to create the repository by using CLI, see the *SnapManager for Oracle Administration Guide for UNIX*.

Steps

1. In the left pane of the SnapManager UI, right-click **Repositories**.
2. Select **Create New Repository** and click **Next**.
3. In the **Repository Database Configuration Information** window, enter the following information:

In this field...	Do this...
User Name	Enter the name of the user you created for the repository database instance.
Password	Enter the password.
Host	Enter the IP address of the host where the repository database instance is created.

In this field...	Do this...
Port	Enter the port used to connect to the repository database instance. The default port is 1521.
Service Name	Enter the name that SnapManager uses to connect to the repository database instance. Depending on the details included in the <code>tnsnames.ora</code> file, this might be the short service name or the fully qualified service name.

4. In the **Perform Repository Add Operation** window, review the configuration summary and click **Add**.

If the operation fails, click the **Operation Details** tab to view what caused the operation to fail. The error details are also captured in the operation log located at `/var/log/smo`.

5. Click **Finish**.

The repository is listed in the left pane under the **Repositories** tree. If you do not see the repository, right-click **Repositories** and click **Refresh**.

Related information

[SnapManager 3.4 for Oracle Administration Guide for UNIX](#)

Preparing storage systems for SnapMirror and SnapVault replication

You can use SnapManager with Data ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with Data ONTAP SnapVault technology to archive backups efficiently to disk. Before you can perform these tasks in SnapManager, you must configure a *data-protection relationship* between the source and destination volumes and *initialize* the relationship.

Note: You cannot configure both a SnapMirror relationship and a SnapVault relationship on the same clustered Data ONTAP source volume. You must configure these relationships on different source volumes.

Related concepts

[Understanding the differences between SnapMirror and SnapVault](#) on page 19

Related tasks

[Preparing storage systems for SnapMirror replication](#) on page 19

[Preparing storage systems for SnapVault replication](#) on page 21

Understanding the differences between SnapMirror and SnapVault

SnapMirror is disaster recovery technology, designed for failover from primary storage to secondary storage at a geographically remote site. SnapVault is archiving technology designed for disk-to-disk Snapshot copy replication for standards compliance and other governance-related purposes.

These objectives account for the different balance each technology strikes between the goals of backup currency and backup retention:

- SnapMirror stores *only* the Snapshot copies that reside in primary storage, because, in the event of a disaster, you need to be able to fail over to the most recent version of primary data you know to be good. Your organization, for example, might mirror hourly copies of production data over a ten-day span. As the failover use case implies, the equipment on the secondary system needs to be equivalent or nearly equivalent to the equipment on the primary system to serve data efficiently from mirrored storage.
- SnapVault, in contrast, stores Snapshot copies *whether or not* they currently reside in primary storage because, in the event of an audit, access to historical data is likely to be as important as access to current data. You might want to keep monthly Snapshot copies of your data over a 20-year span (to comply with government accounting regulations for your business, for example). Since there is no requirement to serve data from secondary storage, you can use slower, less expensive disks on the vault system.

Of course, the different weights SnapMirror and SnapVault give to backup currency and backup retention ultimately derive from the 255-Snapshot copy limit for each volume. Where SnapMirror retains the most recent copies, SnapVault retains the copies taken over the longest period of time.

Preparing storage systems for SnapMirror replication

Before you can use SnapManager's integrated SnapMirror technology to mirror Snapshot copies, you need to configure a *data-protection relationship* between the source and destination volumes, then

initialize the relationship. Upon initialization, SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks that it references to the destination volume. It also transfers any other, less recent Snapshot copies on the source volume to the destination volume.

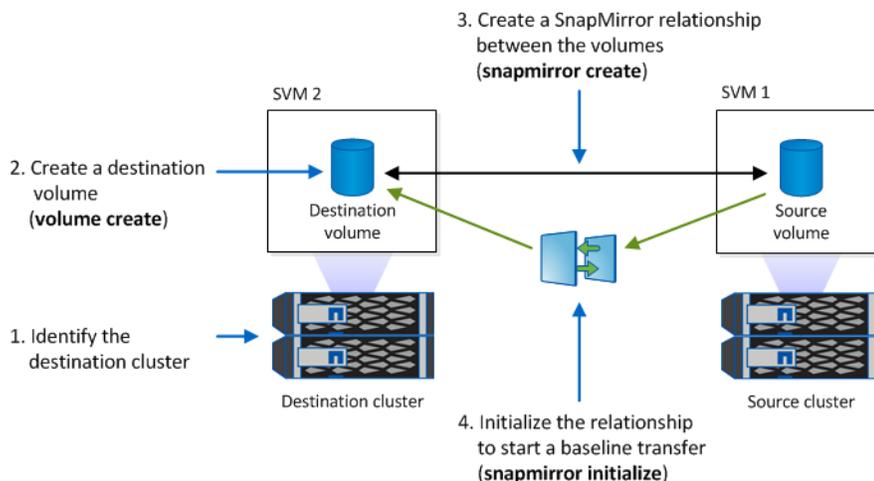
Before you begin

- You must create the source and destination volumes in peered clusters with peered Storage Virtual Machines (SVMs). For more information, see the *Clustered Data ONTAP Cluster Peering Express Guide*.
- You must be a cluster administrator.
- For Snapshot copy verification on the destination volume, the source and destination Storage Virtual Machines (SVMs) must have a management LIF as well as a data LIF. The management LIF must have the same DNS name as the SVM. Set the management LIF role to **data**, the protocol to **none**, and the firewall policy to **mgmt**.

About this task

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create a SnapMirror relationship. The following procedure assumes you are using the CLI. For information about how to create SnapMirror relationship by using OnCommand System Manager, see the *Clustered Data ONTAP Volume Disaster Recovery Preparation Express Guide*.

The following illustration shows the procedure for initializing a SnapMirror relationship:



Steps

1. Identify the destination cluster.
2. On the destination cluster, use the `volume create` command with the `-type DP` option to create a SnapMirror destination volume that is either the same or greater in size than the source volume.

Important: The language setting of the destination volume must match the language setting of the source volume.

Example

The following command creates a 2-GB destination volume named `dstvolB` in SVM2 on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate node01_aggr -type DP
-size 2GB
```

3. On the destination SVM, use the `snapmirror create` command with the `-type DP` parameter to create a SnapMirror relationship.

The `DP` type defines the relationship as a SnapMirror relationship.

Example

The following command creates a SnapMirror relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2. By default, the command assigns the default SnapMirror policy `DPDefault`:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB
-type DP
```

Note: Do not define a mirror schedule for the SnapMirror relationship. SnapManager does that for you when you create a backup schedule.

If you do not want to use the default SnapMirror policy, you can invoke the `snapmirror policy create` command to define a SnapMirror policy.

4. Use the `snapmirror initialize` command to initialize the relationship.

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume. It also transfers any other Snapshot copies on the source volume to the destination volume.

Example

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Related information

[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)

[Clustered Data ONTAP 8.3 Volume Disaster Recovery Preparation Express Guide](#)

Preparing storage systems for SnapVault replication

Before you can use SnapManager's integrated SnapVault technology to archive Snapshot copies to disk, you need to configure a *data-protection relationship* between the source and destination volumes, then initialize the relationship. On initialization, SnapVault makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

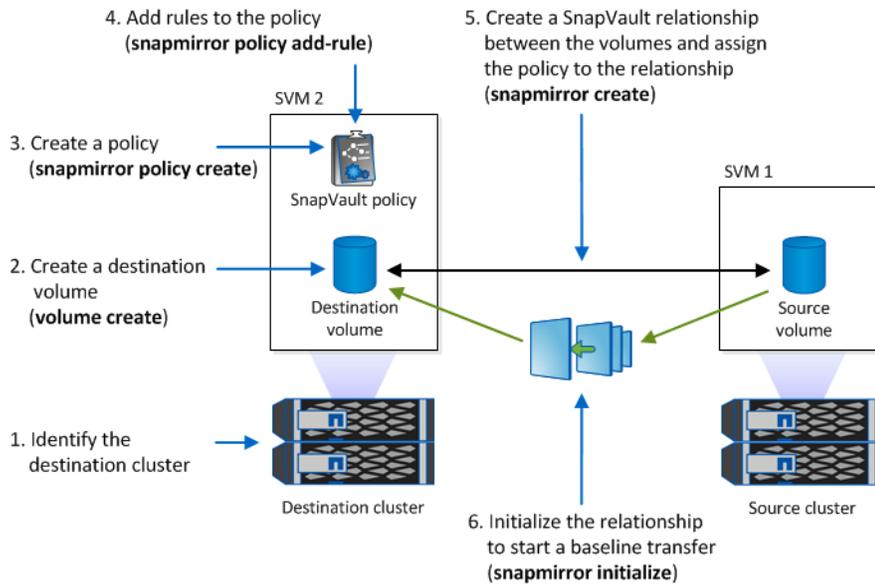
Before you begin

- You must create the source and destination volumes in peered clusters with peered Storage Virtual Machines (SVMs). For more information, see the *Clustered Data ONTAP Cluster Peering Express Guide*.
- You must be a cluster administrator.

About this task

You can use the Data ONTAP command-line interface (CLI) or OnCommand System Manager to create SnapVault relationships. The following procedure assumes you are using the CLI. For information about how to create SnapVault relationship by using OnCommand System Manager, see the *Clustered Data ONTAP Volume Backup Using SnapVault Express Guide*.

The following illustration shows the procedure for initializing a SnapVault relationship:



Steps

1. Identify the destination cluster.
2. On the destination cluster, use the `volume create` command with the `-type DP` option to create a SnapVault destination volume that is the same or greater in size than the source volume.

Important: The language setting of the destination volume must match the language setting of the source volume.

Example

The following command creates a 2-GB destination volume named `dstvolB` in SVM2 on the aggregate `node01_aggr`:

```
cluster2::> volume create -vserver SVM2 -volume dstvolB -aggregate node01_aggr -type DP
-size 2GB
```

3. On the destination SVM, use the `snapmirror policy create` command to create a SnapVault policy.

Example

The following command creates the SVM-wide policy `SVM1-vault`:

```
SVM2::> snapmirror policy create -vserver SVM2 -policy SVM1-vault
```

Note: Do not define a cron schedule or Snapshot copy policy for the SnapVault relationship. SnapManager does that for you when you create a backup schedule.

4. Use the `snapmirror policy add-rule` command to add a rule to the policy that defines the following Snapshot copy labels and the retention policy for each label:
 - Daily
 - Weekly
 - Monthly
 - Hourly
 - Unlimited

Important: The labels are case-sensitive.

These are fixed labels that SnapManager uses. You select one of these options when you archive a backup. You must execute this command once for each of the rules you are adding.

Example

The following command adds a rule to the SVM1-vault policy that defines the “daily” label and specifies that thirty Snapshot copies matching the label should be kept in the vault:

```
SVM2::> snapmirror policy add-rule -vserver SVM2 -policy SVM1-vault
-snapmirror-label Daily -keep 30
```

5. Use the `snapmirror create` command with the `-type XDP` parameter and the `-policy` parameter to create a SnapVault relationship and assign a vault policy.

The `XDP` type defines the relationship as a SnapVault relationship.

Example

The following command creates a SnapVault relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2. It assigns the policy named `SVM1-vault`:

```
SVM2::> snapmirror create -source-path SVM1:srcvolA -destination-path SVM2:dstvolB
-type XDP -policy SVM1-vault
```

6. Use the `snapmirror initialize` command to initialize the relationship.

The initialization process performs a *baseline transfer* to the destination volume. SnapMirror makes a Snapshot copy of the source volume, then transfers the copy and all the data blocks it references to the destination volume.

Example

The following command initializes the relationship between the source volume `srcvolA` on SVM1 and the destination volume `dstvolB` on SVM2:

```
SVM2::> snapmirror initialize -destination-path SVM2:dstvolB
```

Related information

[Clustered Data ONTAP 8.3 Cluster Peering Express Guide](#)

[Clustered Data ONTAP 8.3 Volume Backup Using SnapVault Express Guide](#)

Backing up and verifying your databases

After installing SnapManager, you can create a basic backup of your database and verify that backup to ensure it does not contain any corrupt files.

Related concepts

[SnapManager backup overview](#) on page 24

Related tasks

[Creating a profile for your database](#) on page 26

[Backing up your database](#) on page 28

[Verifying database backups](#) on page 30

[Scheduling recurring backups](#) on page 30

Related references

[Defining a backup strategy](#) on page 24

SnapManager backup overview

SnapManager uses NetApp Snapshot technology to create backups of databases. You can use the DBVERIFY utility, or you can use SnapManager to verify the integrity of the backups.

SnapManager backs up a database by creating Snapshot copies of the volumes containing data files, control files, and archive log files. Together, these Snapshot copies comprise a backup set that SnapManager can use to restore a database.

Defining a backup strategy

Defining a backup strategy before creating your backups ensures that you have backups to successfully restore your databases. SnapManager provides flexible granular backup schedule to meet your Service Level Agreement (SLA).

Note: For SnapManager best practices, see *TR 3761*.

What mode of SnapManager backup do you need?

SnapManager supports two modes of backups:

Backup mode	Description
Online backup	Creates a backup of the database when the database is in online state. This backup mode is also called a hot backup.
Offline backup	Creates a backup of the database when the database is either in a mounted or shutdown state. This backup mode is also called a cold backup.

What type of SnapManager backup do you need?

SnapManager supports three types of backups:

Backup type	Description
Full backup	Creates a backup of the entire database, which includes all the datafiles, control files, and archive log files.
Partial backup	Creates a backup of selected datafiles, control files, tablespaces, and archive log files
Archive log-only backup	Creates a backup of only the archive log files. You must select Backup Archivelogs Separately while creating the profile.

What type of database profile do you need?

SnapManager creates backups based on whether the database profile separates the archive log backups from the data file backups.

Profile type	Description
A single database profile for combined backup of data files and archive logs	Allows you to create: <ul style="list-style-type: none"> • Full backup containing all the data files, archive log files, and control files • Partial backup containing selected data files, tablespaces, archive log files, and control files
Separate database profiles for archive log backups and data file backups	Allows you to create: <ul style="list-style-type: none"> • Combined backup with different labels for data file backup and archive log backup • Data-files-only backup of all the data files along with the control files • Partial data-files-only backup of selected data files or tablespaces along with the control files • Archive-logs-only backup

What naming conventions should be used for Snapshot copies?

Snapshot copies created by backups can follow a custom naming convention. Custom text or built-in variables such as the profile name, the database name, and the database SID provided by SnapManager can be used to create the naming convention. You can create the naming convention while creating the policy.

Note: You must include the *smid* variable in the naming format. The *smid* variable creates a unique Snapshot identifier.

The Snapshot copy naming convention can be changed during or after the creation of a profile. The updated pattern applies only to Snapshot copies that have not yet been created; existing Snapshot copies retain the previous pattern.

How long do you want to retain backup copies on the primary storage system and the secondary storage system?

A backup retention policy specifies the number of successful backups to retain. You can specify the retention policy while creating the policy.

You can select hourly, daily, weekly, monthly, or unlimited as the retention class. For each retention class, you can specify the retention count and retention duration, either together or individually.

- Retention count determines the minimum number of backups of a particular retention class that should be retained.

For example, if backup schedule is *daily* and retention count is *10*, then 10 daily backups are retained.

Note: The maximum number of Snapshot copies that Data ONTAP allows you can retain is 255. After it reaches the maximum limit, by default the creation of new Snapshot copies fail. However, you can configure the rotation policy in Data ONTAP to delete older Snapshot copies.

- Retention duration determines the minimum number of days for which the backup should be retained.

For example, if backup schedule is *daily* and retention duration is *10*, then 10 days of daily backups are retained.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

Note: For long-term retention of backup copies, you should use SnapVault.

Do you want to verify backup copies using the source volume or a destination volume?

If you use SnapMirror or SnapVault, you can verify backup copies using the Snapshot copy on the SnapMirror or SnapVault destination volume rather than the Snapshot copy on the primary storage system. Using a destination volume for verification reduces the load on the primary storage system.

Related information

[NetApp Technical Report 3761: SnapManager for Oracle: Best Practices](#)

Creating a profile for your database

You must create a profile for your database to perform any operation on that database. The profile contains information about the database and can reference only one database; however, a database can be referenced by multiple profiles. A backup that is created using one profile cannot be accessed from a different profile, even if both profiles are associated with the same database.

Before you begin

You must ensure that target database details are included in the `/etc/oratab` file.

About this task

These steps show how to create a profile for your database using the SnapManager UI. You can also use the CLI if you prefer.

For information about how to create profiles using the CLI, see the *SnapManager for Oracle Administration Guide for UNIX*.

Steps

- From the Repositories tree, right-click the repository or the host and select **Create Profile**.
- On the **Profile Configuration Information** page, enter the custom name and password for the profile.
- On the **Database Configuration Information** page, enter the following information:

In this field...	Do this...
Database Name	Enter the name of the database you want to backup.
Database SID	Enter the secure ID (SID) of the database. The database name and the database SID can be the same.
Host	Enter the IP address of the host where the target database resides. You can also specify the host name if the host name is specified in the Domain Name System (DNS).
Host Account	Enter the Oracle user name of the target database. The default value for the user is oracle.
Host Group	Enter the Oracle user group name. The default value is dba.

4. On the **Database Connection Information** page, select one of the following:

Choose this...	If you want to...
Use O/S Authentication	Use the credentials maintained by the operating system to authenticate users who access the database.
Use Database Authentication	Allow Oracle to authenticate an administrative user using password file authentication. Enter the appropriate database connection information. <ul style="list-style-type: none"> In the SYSDBA Privileged User Name field, enter the name of the database administrator with administrative privileges. In the Password field, enter the password of the database administrator. In the Port field, enter the port number used to connect to the host where the database resides. The default value is 1521.
Use ASM Instance Authentication	Allow Automatic Storage Management (ASM) database instance to authenticate an administrative user. Enter the appropriate ASM instance authentication information. <ul style="list-style-type: none"> In the SYSDBA/SYSASM Privileged User Name field, enter the user name of the ASM instance administrator with administrative privileges. In the Password field, enter password of the administrator.

Note: You can select the ASM authentication mode only if you have an ASM instance on the database host.

5. On the **RMAN Configuration Information** page, select one of the following:

Choose this...	If...
Do not use RMAN	You are not using RMAN to manage backup and restore operations.
Use RMAN via the control file	You are managing the RMAN repository using control files.

Choose this...	If...
Use RMAN via Recovery Catalog	You are managing the RMAN repository using recovery catalog database. Enter the user name who has access to recovery catalog database, password, and the Oracle net service name of the database that manages the Transparent Network Substrate (TNS) connection.

6. On the **Snapshot Naming Information** page, select the variables to specify a naming format for the Snapshot copy.

You must include the *smid* variable in the naming format. The *smid* variable creates a unique Snapshot identifier.

7. On the **Policy Settings** page, perform the following:

- a. Enter the retention count and duration for each retention class.

- b. From the **Protection Policy** drop-down list, select the protection policy.

You must select either *SnapManager_cDOT_Mirror* or *SnapManager_cDOT_Vault* policies depending on whether SnapMirror or SnapVault relationship is established.

- c. If you want to back up archive logs separately, select the **Backup Archivelogs Separately** checkbox, specify the retention, and select the protection policy.

You can select a policy which is different from the policy associated for datafiles. For example, if you have selected *SnapManager_cDOT_Mirror* for datafiles, you can select *SnapManager_cDOT_Vault* for archive logs.

8. On the **Configure Notification Settings** page, specify the email notification settings.
9. On the **History Configuration Information** page, select one of the options to maintain the history of SnapManager operations.
10. On the **Perform Profile Create Operation** page, verify the information and click **Create**.
11. Click **Finish** to close the wizard.

If the operation fails, click **Operation Details** to view what caused the operation to fail.

Related information

[SnapManager 3.4 for Oracle Administration Guide for UNIX](#)

Backing up your database

After creating a profile, you must back up your database. You can schedule recurring backups after the initial backup and verification.

About this task

These steps show how to create a backup of your database using the SnapManager user interface. You can also use the command-line interface (CLI) if you prefer.

For information about how to create backups using CLI, see the *SnapManager for Oracle Administration Guide for UNIX*.

Steps

1. From the Repositories tree, right-click the profile containing the database you want to back up, and select **Backup**.

2. In **Label**, enter a custom name for the backup.

You must not include spaces or special characters in the name. If you do not specify a name, SnapManager automatically creates a backup label.

From SnapManager 3.4, you can modify the backup label created automatically by SnapManager. You can edit the `override.default.backup.pattern` and `new.default.backup.pattern` configuration variables to create your own default backup label pattern.

3. Optional: In **SnapVault Label**, you must enter the SnapMirror label that you specified in the rules of the SnapMirror policy while setting up the SnapVault relationship.

Note: The **SnapVault Label** field appears only if you have selected *SnapManager_cDOT_Vault* as the protection policy while creating the profile.

4. Select **Allow startup or shutdown of database, if necessary** to modify the state of the database, if required.

This option ensures that if the database is not in the required state to create a backup, SnapManager automatically brings the database to the desired state to complete the operation.

5. On the **Database, Tablespaces or Datafiles to Backup** page, perform the following:

- a. Select **Backup Datafiles** to back up either the full database, selected data files, or selected tablespaces.
- b. Select **Backup Archivelogs** to back up the archive log files separately.
- c. Select **Prune Archivelogs** if you want to delete the archive log files from the active file system that is already backed up.

Note: If Flash Recovery Area (FRA) is enabled for archive log files, then SnapManager fails to prune the archive log files.

- d. Select **Protect the backup** if you want to enable backup protection.

This option is enabled only if the protection policy was selected while creating the profile.

- e. From the **Type** drop-down list, select the type of backup (offline or online) you want to create. If you select *Auto*, SnapManager creates a backup based on the current state of the database.

- f. From the **Retention Class** drop-down list, select the retention class.

- g. Select the **Verify backup using the Oracle DBVERIFY utility** check box if you want to ensure that the backed-up files are not corrupted.

6. On the **Task Enabling** page, specify whether you want to perform tasks before and after backup operations are completed.

7. On the **Perform Backup Operation** page, verify the information and click **Backup**.

8. Click **Finish** to close the wizard.

If the operation fails, click **Operation Details** to view what caused the operation to fail.

Related information

[SnapManager 3.4 for Oracle Administration Guide for UNIX](#)

Verifying database backups

You can verify the backup of your database to ensure that the backed-up files are not corrupted.

About this task

If you did not select the **Verify backup using the Oracle DBVERIFY utility** check box while creating a backup, you must perform these steps manually to verify the backup. However, if you selected the check box, SnapManager automatically verifies the backup.

Steps

1. From the **Repositories** tree, select the profile.
2. Right-click the backup that you want to verify, and select **Verify**.
3. Click **Finish**.

If the operation fails, click **Operation Details** to view what caused the operation to fail.

In the **Repository** tree, right-click the backup, and then click **Properties** to view the results of the verify operation.

After you finish

You can use backed-up files to perform restore operations. For information about how to perform restore operations using the SnapManager user interface (UI), see the *Online Help*. If you want to use the command-line interface (CLI) to perform restore operations, see the *SnapManager for Oracle Administration Guide for UNIX*.

Related information

[SnapManager 3.4 for Oracle Administration Guide for UNIX](#)

Scheduling recurring backups

You can schedule backup operations so that the backups are initiated automatically at regular intervals. SnapManager allows you to schedule backups on an hourly, daily, weekly, monthly, or one-time basis.

About this task

You can assign multiple backup schedules for a single database. However, when scheduling multiple backups for the same database, you must ensure that the backups are not scheduled at the same time.

These steps show how to create a backup schedule for your database using the SnapManager user interface (UI). You can also use the command-line interface (CLI) if you prefer. For information about how to schedule backups using the CLI, see the *SnapManager for Oracle Administration Guide for UNIX*.

Steps

1. From the Repositories tree, right-click the profile containing the database for which you want to create a backup schedule, and select **Schedule Backup**.
2. In **Label**, enter a custom name for the backup.
You must not include spaces or special characters in the name. If you do not specify a name, SnapManager automatically creates a backup label.

From SnapManager 3.4, you can modify the backup label created automatically by SnapManager. You can edit the `override.default.backup.pattern` and `new.default.backup.pattern` configuration variables to create your own default backup label pattern.

3. Optional: In **SnapVault Label**, you must enter the SnapMirror label that you specified in the rules of the SnapMirror policy while setting up the SnapVault relationship.

Note: The **SnapVault Label** field appears only if you have selected *SnapManager_cDOT_Vault* as the protection policy while creating the profile.
4. Select **Allow startup or shutdown of database, if necessary** to modify the state of the database, if required.

This option ensures that if the database is not in the required state to create a backup, SnapManager automatically brings the database to the desired state to complete the operation.
5. On the **Database, Tablespaces or Datafiles to Backup** page, perform the following:
 - a. Select **Backup Datafiles** to back up either the full database, selected data files, or selected tablespaces.
 - b. Select **Backup Archivelogs** to back up the archive log files separately.
 - c. Select **Prune Archivelogs** if you want to delete the archive log files from the active file system that is already backed up.

Note: If Flash Recovery Area (FRA) is enabled for archive log files, then SnapManager fails to prune the archive log files.
 - d. Select **Protect the backup** if you want to enable backup protection.

This option is enabled only if the protection policy was selected while creating the profile.
 - e. From the **Type** drop-down list, select the type of backup (offline or online) you want to create.

If you select *Auto*, SnapManager creates a backup based on the current state of the database.
 - f. From the **Retention Class** drop-down list, select the retention class.
 - g. Select the **Verify backup using the Oracle DBVERIFY utility** check box if you want to ensure that the backed-up files are not corrupted.
6. In the **Schedule Name** field, enter a custom name of the schedule.

You must not include spaces in the name.
7. On the **Configure Backup Schedule** page, perform the following:
 - a. From the **Perform this operation** drop-down list, select the frequency of the backup schedule.
 - b. In the **Start Date** field, specify the date when you want to initiate the backup schedule.
 - c. In the **Start Time** field, specify the time when you want to initiate the backup schedule.
 - d. Specify the interval in which backups will be created.

For example, if you have selected the frequency as hourly and specify the interval as 2, then backups will be scheduled every 2 hours.
8. On the **Task Enabling** page, specify whether you want to perform tasks before and after backup operations are completed.
9. On the **Perform Backup Schedule Operation** page, verify the information and click **Schedule**.

10. Click **Finish** to close the wizard.

If the operation fails, click **Operation Details** to view what caused the operation to fail.

Related information

[SnapManager 3.4 for Oracle Administration Guide for UNIX](#)

Where to go next

After installing SnapManager and successfully creating a backup, you can use SnapManager to perform restore, recovery, and cloning operations. In addition, you might want to find information about other SnapManager features such as scheduling, managing SnapManager operations, and maintaining a history of operations.

You can find more information about these features as well as release-specific information for SnapManager in the following documentation, all of which is available on the [NetApp Support](#).

- [SnapManager 3.4 for Oracle Administration Guide for UNIX](#)
Describes how to configure and administer SnapManager for Oracle. Topics include how to configure, back up, restore, and clone databases, perform secondary protection, plus an explanation of CLI commands and instructions on how to upgrade and uninstall the product.
- [SnapManager 3.4 for Oracle Release Notes](#)
Describes new features, fixed issues, important cautions, known issues, and limitations for SnapManager for Oracle.
- [SnapManager for Oracle Online Help](#)
Describes the step-by-step procedures for performing different SnapManager operations using the SnapManager UI.
Note: The *Online Help* is integrated with the SnapManager UI and is not available on the Support Site.
- [NetApp Technical Report 3761: SnapManager for Oracle: Best Practices](#)
Describes SnapManager for Oracle best practices.
- [NetApp Technical Report 3633: Best Practices for Oracle Databases on NetApp Storage](#)
Describes best practices to configure Oracle databases on NetApp storage system.

Related information

[NetApp Support](#)

[NetApp Documentation: Product Library A-Z](#)

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, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and 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>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- accessing
 - SnapManager UI [16](#)
- architecture
 - of SnapManager for Oracle [4](#)
- archiving Snapshot copies
 - preparing for SnapVault replication [21](#)

B

- backup mode
 - offline backups [24](#)
 - online backups [24](#)
 - retention class
 - number of backups [24](#)
 - number of days [24](#)
- backup types
 - archive log-only backups [24](#)
 - full backups [24](#)
 - partial backups [24](#)
- backups
 - creating for databases [28](#)
 - creating scheduling in SnapManager for Oracle [30](#)
 - verifying [30](#)

C

- comments
 - how to send feedback about documentation [36](#)
- components
 - of SnapManager for Oracle [4](#)
- configuration
 - workflow [7](#)
- configurations
 - verifying support for systems [9](#)
- configuring
 - repositories [17](#)
 - target databases [11](#)
- creating
 - database profiles [26](#)
 - repository database instances [12](#)
- creating Oracle users
 - for repository database instances [13](#)

D

- database backups
 - verifying [30](#)
- database profiles
 - creating [26](#)
- databases
 - backing up [28](#)
 - backup overview [24](#)
 - backup strategies, SnapManager for Oracle [24](#)
 - creating scheduled backups in SnapManager for Oracle [30](#)
 - Oracle [11](#)

- preparing for SnapMirror replication [19](#)
 - preparing for SnapVault replication [21](#)
 - repository [11](#)
 - target [11](#)
- databases, metadata
 - configuring the repository [17](#)
 - role of the repository [17](#)
- deployment
 - workflow [7](#)
- documentation
 - about cloning operations [33](#)
 - about restore and recovery operations [33](#)
 - additional information about administering SnapManager [33](#)
 - how to receive automatic notification of changes to [36](#)
 - how to send feedback about [36](#)

F

- feedback
 - how to send comments about documentation [36](#)

H

- hosts
 - requirements [10](#)

I

- IMT
 - using to verify support for system configurations [9](#)
 - verifying host configurations [10](#)
- information
 - how to send feedback about improving documentation [36](#)
- installation
 - workflow [7](#)
- installing
 - prerequisites [14](#)
 - SnapManager for Oracle [14](#)
- interfaces
 - CLI [4](#)
 - UI [4](#)
- Interoperability Matrix
 - See IMT

L

- licensing
 - for SnapManager for Oracle [8](#)
 - for storage systems [8](#)
- listeners
 - verifying the configuration of [13](#)

M

- metadata, databases

- configuring the repository [17](#)
- role of the repository [17](#)
- mirroring Snapshots
 - preparing for SnapMirror replication [19](#)

O

- Oracle listeners
 - verifying the configuration of [13](#)
- overview
 - of SnapManager for Oracle [4](#)

P

- per-server licensing
 - requirements for SnapManager for Oracle [8](#)
- per-storage system licensing
 - requirements for SnapManager for Oracle [8](#)
- prerequisites
 - for installing SnapManager for Oracle [14](#)
- profiles, database
 - creating [26](#)

R

- recurring backups
 - scheduling in SnapManager for Oracle [30](#)
- repositories
 - configuring [17](#)
- repository database instances
 - creating [12](#)
 - creating Oracle users for [13](#)
- requirements
 - for deploying SnapManager [8](#)
 - for hosts [10](#)

S

- SnapManager
 - configuring repository databases [11](#)
 - configuring target databases [11](#)
- SnapManager for Oracle
 - archiving Snapshots [21](#)
 - backup overview [24](#)
 - backup strategies [24](#)
 - highlights and benefits of [4](#)
 - installing [14](#)
 - mirroring Snapshots [19](#)
 - overview of [4](#)
- SnapManager server
 - starting [16](#)
- SnapManager UI
 - accessing [16](#)
- SnapMirror

- differences from SnapVault [19](#)
- preparing to mirror backups [19](#)
- Snapshot copies
 - naming conventions of [24](#)
- SnapVault
 - differences from SnapMirror [19](#)
 - preparing to archive backups [21](#)
- starting
 - SnapManager server [16](#)
- storage types
 - supported by SnapManager [9](#)
- suggestions
 - how to send feedback about documentation [36](#)
- support
 - storage types [9](#)
 - verifying system configurations [9](#)
- system configurations
 - verifying support for [9](#)

T

- target database hosts
 - starting SnapManager server from [16](#)
- target databases
 - configuring [11](#)
 - creating Oracle database users for [12](#)
 - verifying the listener configuration of [13](#)

- twitter
 - how to receive automatic notification of documentation changes [36](#)

U

- UI
 - accessing [16](#)
- users
 - creating for repository database instances [13](#)
 - creating for target databases [12](#)

V

- verifying
 - backups [30](#)
 - listener configurations [13](#)

W

- workflows
 - configuration [7](#)
 - deployment [7](#)
 - installation [7](#)