



What SnapManager for Oracle is

SnapManager Oracle

NetApp
November 04, 2025

This PDF was generated from https://docs.netapp.com/us-en/snapmanager-oracle/windows/concept_what_snapmanager_for_oraclesnapmanager_for_sap_does.html on November 04, 2025. Always check docs.netapp.com for the latest.

Table of Contents

What SnapManager for Oracle is	1
What SnapManager for Oracle does	1
Integration with other NetApp applications and technologies	4
Advantages of using SnapManager	4
Create backups using Snapshot copies	5
Why you should prune archive log files	5
Archive log consolidation	6
Full or partial restoration of databases	6
Verify backup status	6
Database backup clones	6
Track details and produce reports	7
What the SnapManager for Oracle architecture is	7
SnapManager host	8
SnapManager graphical user and command-line interfaces	8
SnapManager repository	8
SnapDrive on SnapManager server	9
What repositories are	9
What profiles are	10
What SnapManager operation states are	11
Recoverable and unrecoverable events	12
How SnapManager maintains security	12
Accessing and printing online Help	13

What SnapManager for Oracle is

SnapManager provides the tools required to perform policy-driven data management, schedule and create regular database backups, restore data from these backups in the event of data loss or disaster, and create database clones. You can create backups on primary storage and create protected backups on secondary storage by using post-processing scripts.

SnapManager leverages NetApp technologies when integrating with the latest database releases. SnapManager is integrated with the following NetApp applications and technologies:

- SnapDrive automates storage provisioning tasks and simplifies the process of creating error-free, host-consistent Snapshot copies of the storage.
- Snapshot (a feature of Data ONTAP) creates point-in-time copies of the database.
- SnapVault (a licensed feature of Data ONTAP) leverages disk-based backups for reliable, low-overhead backup and recovery of databases.
- SnapMirror (a licensed feature of Data ONTAP) replicates database data across a global network at high speeds in a simple, reliable, and cost-effective manner.
- SnapRestore (a licensed feature of Data ONTAP) recovers an entire database in seconds, regardless of the capacity or the number of files.
- FlexClone (a licensed feature of Data ONTAP) helps to create fast, space-efficient clones of databases from the Snapshot backups.

SnapManager operates across SAN (FC and iSCSI) protocols.

SnapManager also integrates with native Oracle technology, such as Oracle Recovery Manager (RMAN).

What SnapManager for Oracle does

SnapManager for Oracle simplifies and automates database backup, recovery, and cloning by leveraging the Snapshot copies, SnapRestore, and FlexClone technologies.

SnapManager provides the following benefits to database administrators (DBAs):

- Working with Database profiles
 - You can organize and retain host and database information in profiles.

When you initiate a backup based on a profile, you can reuse the information rather than having to reenter it for every backup. SnapManager also enables you to monitor operations quickly by using profiles.

- In the profile, you can define the Snapshot copies naming patterns and enter custom (prefix or suffix) text, so that all the Snapshot copies can use the same naming convention that meets business policies.
- You do not need to know the storage system name because database files are automatically mapped to the associated storage.
- When you create a new profile, you can specify the option to separate the archive log backup from the data file backup.

You can also update the existing profile to separate the archive log backup from the data file backup.

- Performing the database backup operation
 - Backup of full and partial databases
 - You can create a full or partial backup quickly in a space-efficient way, which allows you to perform backups more frequently.
 - The full database backup contains all the data files, control files, and archive log files in a single backup.
 - The partial database backup contains specified data files or tablespaces, all the control files, and all the archive log files.
 - You can protect backups to secondary storage by using post-processing scripts.
 - You can schedule backups on an hourly, weekly, daily, monthly, or unlimited basis.
 - Separate back up of data files and archive log files
 - SnapManager (3.2 or later) enables you to back up the data files and archive log files separately. To perform this operation, you must specify the option to separate the archive log files while creating or updating the profile.
 - You can specify the count and duration for which the data file backups will be retained, in the retention policy.
 - You can specify the duration for the archive log file backups to be retained in archive log retention duration.
 - SnapManager (3.2 or later) also consolidates the archive log backups to a minimum number of backups by freeing the archive log backups with duplicate archive log files and retaining only the archive log backups with unique archive log files. However, this consolidation can be optionally disabled.
- Managing the archive log files
 - SnapManager (3.2 or later) enables you to prune the archive log files from the archive log destinations.

The space occupied by the pruned archive log files is freed when the archive log backups containing these archive log files are purged.

 - SnapManager ensures that the archive log files are backed up before pruning them from the archive log destinations.
 - The archive log files, which are not backed up are not pruned.
 - SnapManager ensures that the archive log files are shipped to the Data Guard Standby database while pruning archive log files from a Data Guard Primary database.
 - SnapManager ensures that the archive log files are captured by Oracle's Streams Capture process, if any.
 - Recommendation
 - To manage archive log destination space effectively, you must create the archive log backups, and prune the archive log files along with it.
 - SnapManager consolidates the archive log backups to contain minimum number of backups by freeing the archive log backups with duplicate archive log files and retaining only the archive log backups with unique archive log files.

However, this consolidation can be optionally disabled. The archive log backups, which contain duplicate archive log files are freed and a single backup with unique archive logs is retained.

- Performing the database restore operation

- You can perform file-based restore operations.

You can also preview restore operations and obtain a file-by-file analysis of restore operations before the operation is performed.

- You can reduce the mean time to restore a database by using SnapRestore.
 - SnapManager (3.2 or later) enables you to recover the database automatically by using the archive log files from the backup even if the archive log files are not available in the archive log destination.

SnapManager (3.2 or later) also provides a way to recover the database by using the archive log files from the external location to a certain extent.

- Performing database cloning for testing and development

- You can create a clone of a database so that the database can be set up outside the production environment.

For example, you can clone in the development and test environments for testing upgrades to vital systems.

- You can clone a database on a primary storage system.

- SnapManager (3.2 or later) enables you to clone the data file backups with the archive log files available in the backup.

- You can clone the data file backups only when the archive log backup is taken along with it.
 - You can also clone the data file backups if the archive log files are available in the archive log backups made separately to a certain extent.
 - You can also clone the data file backups of a standalone database to a certain extent with archive log files from any external location accessible by Oracle.
 - If the backups are available from an external location, you can specify the external location during cloning for recovering the cloned database to a consistent state.

- Cloning of the archive log-only backups is not supported.

- General

- Integrate with existing Oracle tools, such as Recovery Manager (RMAN).

SnapManager provides the following benefits to storage administrators:

- Supports different SAN protocols.
- Enables you to optimize backups based on the type of backup (full or partial) that works best in your environment.
- Creates space-efficient database backups.
- Creates space-efficient clones.

SnapManager also works with the following Oracle features:

- SnapManager can catalog its backups with Oracle's RMAN.

If using RMAN, a DBA can make use of SnapManager backups and preserve the value of all RMAN functions, such as block-level restore. SnapManager lets RMAN use the Snapshot copies when it performs recovery or restore. For example, you can use RMAN to restore a table within a tablespace and to perform full database and tablespace restores and recoveries from Snapshot copies made by SnapManager. The RMAN recovery catalog should not be in the database that is being backed up.

Integration with other NetApp applications and technologies

SnapManager for Oracle is a stand-alone product that integrates the features from other NetApp products to enable fast backups that require only a small amount of space.

SnapManager integrates with the following NetApp applications and technologies:

Applications and technologies	Description
SnapDrive	SnapManager uses SnapDrive to create Snapshot copies of the storage. Snapshot copies ensure that backups are space-efficient and faster to create than the disk-to-disk backups.
FlexClone (a licensed feature of Data ONTAP)	SnapManager uses the FlexClone feature to create fast, space-efficient clones of backups.
Snapshot (a feature of Data ONTAP)	Snapshot technology creates point-in-time copies of the database.
SnapRestore (a licensed feature of Data ONTAP)	SnapManager reduces the mean time to recover a database by using SnapRestore. SnapRestore can recover individual files to a multi-terabyte volume so that operations can resume quickly.
SnapVault (a licensed feature of Data ONTAP)	SnapVault leverages disk-based backups for reliable, low-overhead backup and recovery of databases.
SnapMirror (a licensed feature of Data ONTAP)	SnapMirror replicates database data across a global network at high speeds in a simple, reliable, and cost-effective manner.

Advantages of using SnapManager

You can use SnapManager for Oracle to perform different tasks on the databases and manage data efficiently.

SnapManager for Oracle works with storage systems and enables you to perform the following tasks:

- Create space-efficient backups to the primary or secondary storage and schedule backups.

You can create full and partial database backups and apply retention duration policies. SnapManager (3.2 or later) enables you to create backups of only the data files and archive logs.

- SnapManager (3.2 or later) enables you to perform preprocessing or post-processing before or after the backup and restore operations.
- SnapManager (3.2 or later) enables you to protect backups by using the postprocessing scripts.
- Restore full or partial databases by using the file-based restore operation.
- Restore and recover database backups automatically.

SnapManager (3.2 or later) enables the restoration and recovery of database backups automatically. SnapManager automatically recovers the restored database by discovering, mounting, and applying the archive log files from the backups.

- Prune archive log files from the archive log destinations when creating backups for only the archive logs.
- Retain the minimum number of archive log backups automatically by retaining only the backups with unique archive log files.
- Track operation details and produce reports by host, profile, backup, or clone.
- Verify the backup status.
- Maintain the history of SnapManager operations associated with a profile.
- Create space-efficient clones of backups on the primary storage.

Create backups using Snapshot copies

SnapManager enables you to create backups on the primary (local) storage and also on the secondary (remote) storage using postprocessing scripts.

Backups created as Snapshot copies are virtual copies of the database and are stored in the same physical medium as the database. Therefore, the backup operation takes less time and requires significantly less space than full, disk-to-disk backups. SnapManager enables you to back up the following:

- All the data files, archive log files, and control files
- Selected data files or tablespaces, all the archive log files, and control files

SnapManager 3.2 or later enables you to optionally back up the following:

- All the data files and the control files
- Selected data files or tablespaces along with the control files
- Archive log files



The data files, archive log files, and control files can be located on different storage systems, storage system volumes, or logical unit numbers (LUNs). You can also use SnapManager to back up a database when there are multiple databases on the same volume or LUN.

Why you should prune archive log files

SnapManager for Oracle enables you to delete archive log files from the active file system that are already backed up.

Pruning enables SnapManager to create backups of distinct archive log files. Pruning, along with the backup retention policy, frees archive log space when backups are purged.



You cannot prune the archive log files when Flash Recovery Area (FRA) is enabled for archive log files. If you specify the archive log location in Flash Recovery Area, you must ensure that you also specify the archive log location in the archive_log_dest parameter.

Archive log consolidation

SnapManager (3.2 or later) for Oracle consolidates the archive log backups to maintain a minimum number of backups for archive log files. SnapManager for Oracle identifies and frees the backups that contain archive logs files that are subsets of other backups.

Full or partial restoration of databases

SnapManager provides the flexibility to restore full databases, specific tablespaces, files, control files, or a combination of these entities. SnapManager enables you to restore data by using a file-based restore process.

SnapManager enables database administrators (DBAs) to preview restore operations. The preview feature enables DBAs to view each restore operation on a file-by-file basis.

DBAs can specify the level to which SnapManager restores and recovers information when performing restore operations. For example, DBAs can restore and recover data to specific points in time. The restore point can be a date and time or an Oracle System Change Number (SCN).

DBAs can use SnapManager to restore the database and use another tool to recover the information. DBAs are not required to use SnapManager for both operations.

SnapManager (3.2 or later) enables you to restore and recover database backups automatically without DBA intervention. You can use SnapManager to create archive log backups, and then use those archive log backups to restore and recover the database backups. Even if the backup's archive log files are managed in an external archive log location, you can specify that external location so those archive logs can help recover the restored database.

Verify backup status

SnapManager can confirm the integrity of the backup using standard Oracle backup verification operations.

Database administrators (DBAs) can perform the verification as part of the backup operation, or at another time. DBAs can set the verify operation to occur during an off-peak time when the load on the host servers is less, or during a scheduled maintenance window.

Database backup clones

SnapManager uses the FlexClone technology to create a writable, space-efficient clone of a database backup. You can modify a clone without changing the backup source.

You might want to clone databases to enable testing or upgrades in nonproduction environments. You can clone a database residing on primary. A clone can be located on the same host or on a different host as the database.

FlexClone technology enables SnapManager to use Snapshot copies of the database to avoid creating an

entire physical, disk-to-disk copy. Snapshot copies require less creation time and take up significantly less space than physical copies.

See the Data ONTAP documentation for more information about FlexClone technology.

Related information

Data ONTAP documentation:

[\[mysupport.netapp.com/documentation/productsat0z/index.html\]\(https://mysupport.netapp.com/documentation/productsat0z/index.html\)](https://mysupport.netapp.com/documentation/productsat0z/index.html)

Track details and produce reports

SnapManager reduces the level of detail database administrators need to track the status of different operations by offering methods to monitor operations from a single interface.

After administrators specify which databases should be backed up, SnapManager automatically identifies the database files for backup. SnapManager displays information about repositories, hosts, profiles, backups, and clones. You can monitor the operations on specific hosts or databases.

What the SnapManager for Oracle architecture is


The SnapManager for Oracle architecture includes many components, such as the SnapManager for Oracle host, client, and repository. Other components include the primary and secondary storage systems and other NetApp products.

The SnapManager for Oracle architecture includes the following architectural components:

- SnapManager host
- SnapManager graphical user interface or command-line interface
- SnapManager repository
- Primary storage system
- Secondary storage systems
- SnapDrive for Windows

The following image shows the architecture of SnapManager for Oracle and related components:

Architecture



SnapManager host

The SnapManager host is a Windows server, which also runs other NetApp products.

The SnapManager host is installed with the following products:

- SnapDrive for Windows
- Host Utilities

The SnapManager host runs as a service.

SnapManager graphical user and command-line interfaces

The SnapManager client includes both a graphical user interface (GUI) and a command-line interface (CLI).

SnapManager repository

The repository stores information related to different SnapManager operations, for example, the time of backups, tablespaces and data files backed up, storage systems used, clones made, and Snapshot copies created.

The repository database cannot exist in the same database and also cannot be a part of the same database that SnapManager is backing up. This is because the repository stores the names of the database Snapshot

copies created during the backup operations. The repository must be created in a different database than the database that is being backed up. This means that you must have at least two databases: the SnapManager repository database and the target database managed by SnapManager. When you run the SnapManager services, both the databases must be up and running.



You must not perform any SnapManager operations by using the GUI or CLI when the repository database is down.

SnapDrive on SnapManager server

SnapManager uses SnapDrive for Windows to create Snapshot copies of the storage system. SnapDrive resides on the same server as SnapManager.

What repositories are

SnapManager organizes information into profiles, which are then associated with repositories. Profiles contain information about the database that is being managed, while the repository contains data about the operations that are performed on profiles.

The repository records when a backup took place, which files were backed up, and whether a clone was created from the backup. When database administrators restore a database or recover a portion of it, SnapManager queries the repository to determine what was backed up.

Because the repository stores the names of the database Snapshot copies created during backup operations, the repository database cannot exist in the same database and also cannot be a part of the same database that SnapManager is backing up. You must have at least two databases (the SnapManager repository database and the target database being managed by SnapManager) up and running when you execute SnapManager operations.

If you try to open the graphical user interface (GUI) when the repository database is down, the following error message is logged in the `sm_gui.log` file: [WARN]: SMO-01106: Error occurred while querying the repository: No more data to read from socket. Also, SnapManager operations fail when the repository database is down. For more information about the different error messages, see *Troubleshooting known issues*.

You can use any valid host name, service name, or user name to perform operations. For a repository to support SnapManager operations, the repository user name and service name must consist of only the following characters: alphabetic characters (A-Z), digits (0-9), minus sign (-), underscore (_), and period (.).

The repository port can be any valid port number and the repository host name can be any valid host name. The host name must consist of alphabetic characters (A-Z), digits (0-9), minus sign (-), and period (.), but not an underscore (_).

The repository must be created in an Oracle database. The database that SnapManager uses should be set up in accordance with Oracle procedures for database configuration.

A single repository can contain information about multiple profiles; however, each database is normally associated with only one profile. You can have multiple repositories, with each repository containing multiple profiles.


What profiles are

SnapManager uses profiles to store the information necessary to perform operations on a given database. A profile contains the information about the database including its credentials, backups, and clones. By creating a profile, you do not have to specify database details each time you perform an operation on that database.

A profile can reference only one database. The same database can be referenced by more than one profile. Backups created using one profile cannot be accessed from a different profile, even if both the profiles reference the same database.

Profile information is stored in a repository. The repository contains both the profile information for the database and information about the Snapshot copies that serve as the database backup. The actual Snapshot copies are stored on the storage system. The Snapshot copy names are stored in the repository containing the profile for that database. When you perform an operation on a database, you must select the profile from the repository.

The following figure illustrates how repositories can hold multiple profiles, but also that each profile can define only one database:



In the preceding example, Repository2 is on database DB1 and Repository1 is on the database DB4.

Each profile contains the credentials for the database associated with the profile. The credentials enable SnapManager to connect to and work with the database. The stored credentials include the user name and password pairs for accessing the host, the repository, the database, and the required connection information if you are using Oracle Recovery Manager (RMAN).

You cannot access a backup that was created using one profile from a different profile, even if both the profiles are associated with the same database. SnapManager places a lock on the database to prevent two incompatible operations from being performed simultaneously.

Profile for creating full and partial backups

You can create profiles to take full backups or partial backups.

The profiles that you specify to create the full and partial backups contain both the data files and archive log files. SnapManager does not allow such profiles to separate the archive log backups from the data file backups. The full and partial backups are retained based on the existing backup retention policies. You can schedule full and partial backups based on the time and frequency that suits you.

Profiles for creating data files-only backups and archive log-only backups

SnapManager (3.2 or later) allows you to create profiles that take backups of the archive log files separately from the data files. After you use the profile to separate the backup types, you can create either data files-only backups or archive log-only backups of the database. You can also create a backup containing both the data files and archive log files together.

The retention policy applies to all the database backups when the archive log backups are not separated. After you separate the archive log backups, SnapManager allows you to specify different retention durations .

Retention policy

SnapManager determines whether a backup should be retained by considering both the retention count (for example, 15 backups) and the retention duration (for example, 10 days of daily backups). A backup expires when its age exceeds the retention duration set for its retention class and the number of backups exceeds the retention count. For example, if the backup count is 15 (meaning that SnapManager has taken 15 successful backups) and the duration requirement is set for 10 days of daily backups, the five oldest, successful, and eligible backups expire.

Archive log retention duration

After the archive log backups are separated, they are retained based on the archive log retention duration. Archive log backups taken with data file backups are always retained along with those data file backups, regardless of the archive log retention duration.

Related information

[Managing profiles for efficient backups](#)

What SnapManager operation states are

SnapManager operations (backup, restore, and clone) can be in different states, with each state indicating the progress of the operation.

Operation state	Description
Succeeded	The operation completed successfully.
Running	The operation has started, but is not finished. For instance, a backup, which takes two minutes, is scheduled to occur at 11:00 a.m.. When you view the Schedule tab at 11:01 a.m., the operation appears as Running.
No operation found	The schedule has not run or the last run backup was deleted.
Failed	The operation failed. SnapManager has automatically executed the abort process and cleaned the operation.

Recoverable and unrecoverable events

A recoverable SnapManager event has the following problems:

- The database is not stored on a storage system that runs Data ONTAP.
- SnapDrive for Windows is not installed or cannot access the storage system.
- SnapManager fails to create a Snapshot copy or provision storage if the volume is out of space, the maximum number of Snapshot copies has been reached, or an unanticipated exception occurs.

When a recoverable event occurs, SnapManager performs an abort process and attempts to return the host, database, and storage system to the starting state. If the abort process fails, SnapManager treats the incident as an unrecoverable event.

An unrecoverable (out-of-band) event occurs when any of the following happens:

- A system issue occurs, such as when a host fails.
- The SnapManager process is stopped.
- An in-band abort operation fails when the storage system fails, the logical unit number (LUN) or storage volume is offline, or the network fails.

When an unrecoverable event occurs, SnapManager performs an abort process immediately. The host, database, and storage system might not have returned to the initial states. If this is the case, you must perform a cleanup after the SnapManager operation fails by deleting the orphaned Snapshot copy and removing the SnapManager lock file.

If you want to delete the SnapManager lock file, navigate to \$ORACLE_HOME on the target machine and delete the sm_lock_TargetDBName file. After deleting the file, you must restart the SnapManager for Oracle server.

How SnapManager maintains security

You can perform SnapManager operations only if you have the appropriate credentials. Security in SnapManager is governed by user authentication.

SnapManager maintains security by requesting user authentication through password prompts or by setting user credentials. An effective user is authenticated and authorized with the SnapManager server.

SnapManager credentials and user authentication differ significantly from SnapManager 3.0:

- In SnapManager versions earlier than 3.0, you would set an arbitrary server password when you install SnapManager. Anyone who wants to use the SnapManager server would need the SnapManager server password. The SnapManager server password would need to be added to the user credentials by using the smo credential set -host command.
- In SnapManager (3.0 and later), the SnapManager server password has been replaced by individual user operating system (OS) authentication. If you are not running the client from the same server as the host, the SnapManager server performs the authentication by using your OS user names and passwords. If you do not want to be prompted for your OS passwords, you can save the data to your SnapManager user credentials cache by using the smo credential set -host command.



The smo credential set -host command remembers your credentials when the host.credentials.persist property in the smo.config file is set to true.

Example

User1 and User2 share a profile called Prof2. User2 cannot perform a backup of Database1 in Host1 without permission to access Host1. User1 cannot clone a database to Host3 without permission to access Host3.

The following table describes different permissions assigned to the users:

Permission type	User1	User2
Host Password	Host1, Host2	Host2, Host3
Repository Password	Repo1	Repo1
Profile Password	Prof1, Prof2	Prof2

In the case where User1 and User2 do not have any shared profiles, assume User1 has permissions for the hosts named Host1 and Host2, and User2 has permissions for the host named Host2. User2 cannot run even the nonprofile commands such as dump and system verify on Host1.

Accessing and printing online Help

The online Help provides instructions for the tasks that you can perform using the SnapManager graphical user interface. The online Help also provides descriptions of fields on the windows and wizards.

1. Perform one of the following actions:
 - In the main window, click **Help > Help Contents**.
 - In any window or wizard, click **Help** to display help specific to that window.
2. Use the **Table of Contents** in the left pane to navigate through the topics.
3. Click the Printer icon at the top of the help window to print individual topics.

Copyright information

Copyright © 2025 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.