



Limitations when working with SnapManager

SnapManager Oracle

NetApp
November 04, 2025

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

Table of Contents

Limitations when working with SnapManager	1
SnapManager limitations for clustered Data ONTAP	6
Limitations related to Oracle Database	6
Deprecated versions of Oracle database	7
Volume management restrictions	7

Limitations when working with SnapManager

You must be aware of the scenarios and limitations that might affect your environment.

Limitations related to database layouts and platforms

- SnapManager supports control files on a file system or in an ASM disk group and does not support control files on raw devices.
- SnapManager operates in a Microsoft clustering (MSCS) environment but does not recognize the state of the MSCS configuration (active or passive) and does not transfer active management of a repository to a standby server in an MSCS cluster.
- In Red Hat Enterprise Linux (RHEL) and Oracle Enterprise Linux 4.7, 5.0, 5.1, 5.2, and 5.3, the ext3 file system is not supported when deploying Oracle over raw devices by using dynamic multipathing (DMP) in a multipath network I/O (MPIO) environment.

This issue is noticed in SnapManager only when using SnapDrive 4.1 for UNIX or earlier versions.

- SnapManager on RHEL does not support partitioning of disks using the **parted** utility.

This is an issue with the RHEL **parted** utility.

- In a RAC configuration, when a profile name is updated from RAC node A, the schedule file for the profile is updated only for RAC node A.

The schedule file for the same profile on RAC node B is not updated and contains the earlier schedule information. When a scheduled backup is triggered from node B, the scheduled backup operation fails because node B contains the earlier schedule file. However, the scheduled backup operation succeeds from node A, on which the profile is renamed. You can restart the SnapManager server so that you receive the latest schedule file for the profile on node B.

- The repository database might exist on a host that can be accessed by using more than one IP address.

If the repository is accessed by using more than one IP address, then the schedule file is created for each of the IP addresses. If the schedule backup is created for a profile (for example, profile A) under one of the IP addresses (for example, IP1), then the schedule file for only that IP address gets updated. If profile A is accessed from another IP address (for example, IP2), the scheduled backup is not listed because the schedule file of IP2 does not have an entry for the schedule that was created under IP1.

You can wait for the schedule to be triggered from that IP address and the schedule file to be updated, or you can restart the server.

Limitations related to SnapManager configuration

- SnapManager can be configured to catalog database backups with RMAN.
If an RMAN recovery catalog is used, the recovery catalog must be in a different database than the database that is backed up.
- SnapDrive for UNIX supports more than one type of file system and volume manager on certain platforms.

The file system and volume manager used for database files must be specified in the SnapDrive configuration file as the default file system and volume manager.

- SnapManager supports databases on MultiStore storage systems with the following requirements:
 - You must configure SnapDrive to set passwords for MultiStore storage systems.
 - SnapDrive cannot create a Snapshot copy of a LUN or file residing in a qtree in a MultiStore storage system if the underlying volume is not in the same MultiStore storage system.
- SnapManager does not support accessing two SnapManager servers running on different ports from a single client (both from the CLI or GUI).

The port numbers should be the same on the target and remote hosts.

- All LUNs within a volume should reside at the volume level or inside qtrees, but not both.

This is because if the data is residing on the qtrees and you mount the volume, then the data inside the qtrees is not protected.

- SnapManager operations fail and you cannot access the GUI when the repository database is down.

You must verify that the repository database is running when you perform any SnapManager operations.

- SnapManager does not support Live Partition Mobility (LPM) and Live Application Mobility (LAM).
- SnapManager does not support Oracle Wallet Manager and Transparent Data Encryption (TDE).
- SnapManager does not support MetroCluster configurations in raw device mapping (RDM) environments because MetroCluster configurations are yet to be supported by Virtual Storage Console (VSC).

Limitations related to profile management

- If you update the profile to separate the archive log backups, then you cannot perform a rollback operation on the host.
- If you enable a profile from the GUI to create archive log backups, and later try to update the profile by using the Multi Profile Update window or Profile Update window, then you cannot modify that profile to create a full backup.
- If you update multiple profiles in the Multi Profile Update window and some profiles have the **Backup Archivelogs separately** option enabled and other profiles have the option disabled, then the **Backup Archivelogs separately** option is disabled.
- If you update multiple profiles and some profiles have the **Backup Archivelogs separately** option enabled and other profiles have the option disabled, then the **Backup Archivelogs separately** option in the Multi Profile Update window is disabled.
- If you rename the profile, then you cannot roll back the host.

Limitations related to rolling upgrade or rollback operations

- If you try to install an earlier version of SnapManager for a host without performing the rollback operation on the host in the repository, you might not be able to do the following:
 - View the profiles that were created in earlier or later versions of SnapManager for the host.
 - Access backups or clones that were created in earlier or later versions of SnapManager.
 - Perform rolling upgrade or rollback operations on the host.
- After you separate the profiles to create archive log backups, you cannot perform a rollback operation on the related host repository.

Limitations related to backup operations

- Backup creation might fail if you run SnapManager operations concurrently on the same host against a different ASM database.
- During recovery, if the backup is already mounted, SnapManager does not mount the backup again and uses the already mounted backup.

If the backup is mounted by a different user and you do not have access to the previously mounted backup, then the other user must provide you the permission.

All archive log files have read permission for users assigned to a group; you might not have the access permission to the archive log file, if the backup is mounted by a different user group. Users can give permission to the mounted archive log files manually, and then retry the restore or recovery operation.

- SnapManager sets the backup state as “PROTECTED”, even when one of the Snapshot copies of the database backup is transferred to the secondary storage system.
- You can use the task specification file for scheduled backup only from SnapManager 3.2 or later.
- When a backup or clone operation is executed simultaneously on the 10gR2 and 11gR2 RAC databases over ASM, then one of the backup or clone creation operations fails.

This failure is because of a known Oracle limitation.

- SnapManager integrated with Protection Manager supports the backup of multiple volumes in primary storage to a single volume in secondary storage for SnapVault and qtree SnapMirror.

Dynamic secondary volume sizing is not supported. The Provisioning Manager and Protection Manager Administration Guide For Use with DataFabric Manager Server 3.8 has for more information about this.

- SnapManager does not support vaulting of backups using the post-processing script.
- If the repository database is pointing to more than one IP address and each IP address has a different host name, then the backup scheduling operation is successful for one IP address but fails for the other IP address.
- After upgrading to SnapManager 3.4 or later, any backups scheduled with post-processing scripts using SnapManager 3.3.1 cannot be updated.

You must delete the existing schedule and create a new schedule.

Limitations related to restore operations

- When you use an indirect method for performing a restore operation and the archive log files required for recovery are available only in backups from the secondary storage system, SnapManager fails to recover the database.

This is because SnapManager cannot mount the backup of archive log files from the secondary storage system.

- When SnapManager performs a volume restore operation, the archive log backup copies that are made after the corresponding backup is restored are not purged.

When the data files and archive log file destination exist on the same volume, the data files can be restored through a volume restore operation if there are no archive log files available in the archive log file destination. In such a scenario, the archive log Snapshot copies that are created after the backup of the data files are lost.

You should not delete all of the archive log files from the archive log destination.

- In an ASM environment, if the Oracle Cluster Registry (OCR) and voting disk files coexist on a disk group that has data files, then the fast restore preview operation displays the wrong directory structure for the OCR and voting disk.

Limitations related to clone operations

- You cannot view any numerical values between 0 and 100 for the progress of the clone split operation because of the speed with which the inodes are discovered and processed by the storage system containing the flexible volume.
- SnapManager does not support receiving emails only for the successful clone split operations.
- SnapManager only supports splitting a FlexClone.
- The cloning of online database backup of the RAC database that uses external archive log file location fails because of failure in recovery.

The cloning fails because Oracle fails to find and apply the archive log files for recovery from the external archive log location. This is an Oracle limitation. For more information, see the Oracle Bug ID: 13528007. Oracle does not apply archive log from the non-default location at the [Oracle support site](#). You must have a valid Oracle metalink user name and password.

- SnapManager 3.3 or later does not support using the clone specification XML file created in the releases before SnapManager 3.2.
- If temporary tablespaces are located in a different location from the datafiles location, a clone operation creates the tablespaces in the datafiles location.

However, if temporary tablespaces are Oracle Managed Files (OMFs) that are located in a different location from the datafiles location, the clone operation does not create the tablespaces in the datafiles location. The OMFs are not managed by SnapManager.

- SnapManager fails to clone a RAC database if you select the -resetlogs option.

Limitations related to archive log files and backups

- SnapManager does not support pruning of archive log files from the flash recovery area destination.
- SnapManager does not support pruning of archive log files from the standby destination.
- The archive log backups are retained based on the retention duration and default hourly retention class.

When the archive log backup retention class is modified by using the SnapManager CLI or GUI, the modified retention class is not considered for backup because archive log backups are retained based on retention duration.

- If you delete the archive log files from the archive log destinations, the archive log backup does not include archive log files older than the missing archive log file.

If the latest archive log file is missing, then the archive log backup operation fails.

- If you delete the archive log files from the archive log destinations, the pruning of archive log files fail.
- SnapManager consolidates the archive log backups even when you delete the archive log files from the archive log destinations or when the archive log files are corrupted.

Limitations related to changing of target database host name

The following SnapManager operations are not supported when you change the target database host name:

- Changing the target database host name from the SnapManager GUI.
- Rolling back of the repository database after updating the target database host name of the profile.
- Simultaneously updating multiple profiles for a new target database host name.
- Changing the target database host name when any SnapManager operation is running.

Limitations related to the SnapManager CLI or GUI

- The SnapManager CLI commands for the profile create operation that are generated from the SnapManager GUI do not have history configuration options.

You cannot use the profile create command to configure history retention settings from the SnapManager CLI.

- SnapManager does not display the GUI in Mozilla Firefox when there is no Java Runtime Environment (JRE) available on the UNIX client.
- While updating the target database host name using the SnapManager CLI, if there are one or more open SnapManager GUI sessions, then all of the open SnapManager GUI sessions fail to respond.

Limitations related to SnapMirror and SnapVault

- The SnapVault post-processing script is not supported if you are using Data ONTAP operating in 7-Mode.
- If you are using ONTAP, you cannot perform volume-based SnapRestore (VBSR) on the backups that were created in the volumes that have SnapMirror relationships established.

This is because of an ONTAP limitation, which does not allow you to break the relationship when doing a VBSR. However, you can perform a VBSR on the last or most recently created backup only when the volumes have SnapVault relationships established.

- If you are using Data ONTAP operating in 7-Mode and want to perform a VBSR on the backups that were created in the volumes that have SnapMirror relationships established, you can set the override-vbsr-snapmirror-check option to ON in SnapDrive for UNIX.

The SnapDrive documentation has more information about this.

- In some scenarios, you cannot delete the last backup associated with the first Snapshot copy when the volume has a SnapVault relationship established.

You can delete the backup only when you break the relationship. This issue is because of an ONTAP restriction with base Snapshot copies. In a SnapMirror relationship the base Snapshot copy is created by the SnapMirror engine, and in a SnapVault relationship the base Snapshot copy is the backup created by using SnapManager. For each update, the base Snapshot copy points to the latest backup created by using SnapManager.

Limitations related to Data Guard Standby databases

- SnapManager does not support Logical Data Guard Standby databases.
- SnapManager does not support Active Data Guard Standby databases.
- SnapManager does not allow online backups of Data Guard Standby databases.
- SnapManager does not allow partial backups of Data Guard Standby databases.
- SnapManager does not allow restoring of Data Guard Standby databases.

- SnapManager does not allow pruning of archive log files for Data Guard Standby databases.
- SnapManager does not support Data Guard Broker.

Related information

Documentation on the NetApp Support Site: mysupport.netapp.com

SnapManager limitations for clustered Data ONTAP

You must know the limitations for some functionalities and SnapManager operations if you are using clustered Data ONTAP.

The following functionalities are not supported if you are using SnapManager on clustered Data ONTAP:

- Data protection capabilities if SnapManager is integrated with OnCommand Unified Manager
- A database in which one LUN belongs to a system running Data ONTAP operating in 7-Mode and the other LUN belongs to a system running clustered Data ONTAP
- SnapManager for Oracle does not support migration of a Vserver, which is not supported by clustered Data ONTAP
- SnapManager for Oracle does not support the clustered Data ONTAP 8.2.1 functionality to specify different export policies for volumes and qtrees

Limitations related to Oracle Database

Before you start working with SnapManager, you must know the limitations related to Oracle Database.

The limitations are as follows:

- SnapManager supports Oracle versions 10gR2, 11gR1, 11gR2 and 12c, but does not support Oracle 10gR1 as the repository or target database.
- SnapManager will not support the use of a SCAN IP address in place of a host name.

SCAN IP is a new feature in Oracle 11gR2.

- SnapManager does not support Oracle Cluster File System (OCFS).
- Oracle 11g in a Direct NFS (dNFS) environment allows additional mount point configurations in the orafstab file, such as multiple paths for load balancing.

SnapManager does not modify the orafstab file. You must manually add any additional properties that you want the cloned database to use, in the orafstab file.

- Support for Oracle Database 9i is deprecated from SnapManager 3.2.
- Support for Oracle Database 10gR2 (earlier than 10.2.0.5) is deprecated from SnapManager 3.3.1.



Identify the different versions of Oracle databases supported by referring to the Interoperability Matrix.

Related information

Deprecated versions of Oracle database

Oracle database 9i is not supported by SnapManager 3.2 or later, and Oracle database 10gR2 (earlier than 10.2.0.4) is not supported by SnapManager 3.3.1 or later.

If you are using Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases and want to upgrade to SnapManager 3.2 or later, you cannot create new profiles; a warning message is displayed.

If you are using Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases and want to upgrade to SnapManager 3.2 or later, you must perform one of the following:

- Upgrade Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases to either Oracle 10gR2 (10.2.0.5), 11gR1, or 11gR2 databases, and then upgrade to SnapManager 3.2 or 3.3.

If you are upgrading to Oracle 12c, then you must upgrade to SnapManager 3.3.1 or later.



Oracle database 12c is supported only from SnapManager 3.3.1.

- Manage the Oracle 9i databases using a patch version of SnapManager 3.1.

You can use SnapManager 3.2 or 3.3 if you want to manage Oracle 10gR2, 11gR1, or 11gR2 databases and use SnapManager 3.3.1 or later if you want to manage Oracle 12c databases along with the other supported databases.

Volume management restrictions

SnapManager has certain volume management restrictions that might affect your environment.

You can have multiple disk groups for a database; however, the following limitations apply to all disk groups for a given database:

- Disk groups for the database can be managed by only one volume manager.
- Raw devices backed by a logical volume manager are not supported for protection of Oracle data.

Raw device storage and Automatic Storage Management (ASM) disk groups must be provisioned directly on physical devices. In some cases, partitioning is required.

- A Linux environment without logical volume management requires a partition.

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—with 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.