



Troubleshooting known issues

SnapManager Oracle

NetApp
January 29, 2021

This PDF was generated from https://docs.netapp.com/us-en/snapmanager-oracle/unix-administration/reference_troubleshooting_known_issues.html on September 13, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Troubleshooting known issues 1
 - SnapManager for Oracle fails to identify Cluster-Mode profiles 1
 - The server fails to start 1
 - Terminating a currently running SnapManager operation 2
 - Unable to delete or free the last protected backup 2
 - Unable to manage archive log file destination names if the destination names are part of other destination names 3
 - Restoring control files that are multiplexed on Automatic Storage Management (ASM) and non-ASM storage fails 3
 - SnapManager clone operation fails 4
 - Repository database size grows with time and not with the number of backups 4
 - The SnapManager GUI cannot be accessed and SnapManager operations fail when the repository database is down 5
 - Unable to create temporary files for the cloned database 5
 - Unable to migrate the protocol from NFSv3 to NFSv4 6
 - Back up of Data Guard Standby database fails 6

Troubleshooting known issues

You should be aware of some known issues that might occur when you use SnapManager, and how to work around them.

SnapManager for Oracle fails to identify Cluster-Mode profiles

If the Cluster-Mode profile name is not present in the `cmode_profiles.config` file in the SnapManager for Oracle installation directory, the following error message might trigger:

Please configure DFM server using `snapdrive config set -dfm user_name appliance_name`.

Also, while upgrading the SnapManager for Oracle, if you delete the `/opt/NetApp/smo/*` folder, then the `cmode_profiles.config` file that has the Cluster-Mode profile names also get deleted. This issue also triggers the same error message.

Workaround

Update the profile: `smo profile update-profile <profile_name>`



If SnapManager for Oracle is installed in the `/opt/NetApp/smo/` path, then the file location is `/opt/NetApp/smo/cmode_profile/cmode_profiles.config`.

The server fails to start

When starting the server, you might see an error message similar to the following:

SMO-01104: Error invoking command: SMO-17107: SnapManager Server failed to start on port 8074 because of the following errors: `java.net.BindException: Address already in use`

This might be because the SnapManager listening ports (27214 and 27215, by default) are currently in use by another application.

This error can also occur if the `smo_server` command is already running, but SnapManager does not detect the existing process.

Workaround

You can reconfigure either SnapManager or the other application to use different ports.

To reconfigure SnapManager, edit the following file: `/opt/NTAP/smo/properties/smo.config`

You assign the following values:

- `SMO Server.port=27214`
- `SMO Server.rmiRegistry.port=27215`
- `remote.registry.ocjdbc.port= 27215`

The `remote.registry.ocjdbc.port` must be the same as `Server.rmiRegistry.port`.

To start the SnapManager server, enter the following command: `smo_server start`



An error message is displayed if the server is already running.

If the server is already running, perform the following steps:

1. Stop the server by entering the following command: `smo_server stop`
2. Restart the server by entering the following command: `smo_server start`

Terminating a currently running SnapManager operation

If SnapManager server hangs and you cannot execute any operations successfully, you can terminate SnapManager and its operations.

Workaround

SnapManager works with both SnapManager and Protection Manager. You must perform the following steps to list the different processes running and stop the last process running.

1. List all SnapDrive processes that are running: `ps`

Example: `ps | grep snapdrive`

2. Stop the SnapDrive process or processes: `kill <pid>`

`pid` is the list of processes you found using the `ps` command.



Do not stop all SnapDrive processes. You might want to end only the last process that is running.

3. If one of the operations involves restoring a protected backup from secondary storage, open the Protection Manager console and perform the following:
 - a. From the System menu, select **Jobs**.
 - b. Select **Restore**.
 - c. Check for the name of the dataset that matches the one in the SnapManager profile.
 - d. Right-click and select **Cancel**.
4. List the SnapManager processes:
 - a. Log in as a root user.
 - b. List the processes by using the `ps` command.

Example: `ps | grep java`

5. End the SnapManager process: `kill <pid>`

Unable to delete or free the last protected backup

When you create the first backup for a profile on secondary storage, SnapManager sends all the information about the backup to Protection Manager. For subsequent backups related to this profile, SnapManager sends only the modified information. If you remove the last protected backup, SnapManager loses the ability to

identify the differences between backups and must to find a way to rebaseline these relationships. Therefore, attempting to delete the last protected backup results in an error message being displayed.

Workaround

You can delete the profile or only the profile backup.

To delete the profile, perform the following steps:

1. Delete the profile's backups.
2. Update the profile and disable protection in the profile.

This deletes the dataset.

3. Delete the last protected backup.
4. Delete the profile.

To delete only the backup, perform the following steps:

1. Create another backup copy of the profile.
2. Transfer that backup copy to secondary storage.
3. Delete the previous backup copy.

Unable to manage archive log file destination names if the destination names are part of other destination names

While creating an archive log backup, if the user excludes a destination that is part of other destination names, then the other destination names are also excluded.

For example, assume that there are three destinations available to be excluded: /dest, /dest1, and /dest2. While creating the archive log file backup, if you exclude /dest by using the command

```
smo backup create -profile almsamp1 -data -online -archivelogs -exclude  
-dest /dest
```

, SnapManager for Oracle excludes all the destinations starting with /dest.

Workaround

- Add a path separator after destinations are configured in v\$archive_dest. For example, change the /dest to /dest/.
- While creating a backup, include destinations instead of excluding any destination.

Restoring control files that are multiplexed on Automatic Storage Management (ASM) and non-ASM storage fails

When the control files are multiplexed on ASM and non-ASM storage, the backup operation is successful. However, when you try to restore control files from that successful backup, the restore operation fails.

SnapManager clone operation fails

When you clone a backup in SnapManager, the DataFabric Manager server might fail to discover volumes, and display the following error message:

```
SMO-13032: Cannot perform operation: Clone Create. Root cause: SMO-11007: Error cloning from snapshot:
FLOW-11019: Failure in ExecuteConnectionSteps: SD-00018: Error discovering storage for
/mnt/datafile_clone3: SD-10016: Error executing snapdrive command "/usr/sbin/snapdrive storage show -fs
/mnt/datafile_clone3": 0002-719 Warning: Could not check SD.Storage.Read access on volume
filer:/vol/SnapManager_20091122235002515_vol1 for user user-vm5\oracle on Operations Manager servers
x.x.x.x
```

Reason: Invalid resource specified. Unable to find its Id on Operations Manager server 10.x.x.x

This occurs if the storage system has large number of volumes.

Workaround

You must perform one of the following:

- From the Data Fabric Manager server, run `dfm host discover storage_system`.

You can also add the command in a shell script file and schedule a job in the DataFabric Manager server to run the script at a frequent interval.

- Increase the value of `dfm-rbac-retries` in the `Snapdrive.conf` file.

SnapDrive uses the default refresh interval value and default number of retries. The default value of `dfm-rbac-retry-sleep-secs` is 15 seconds and `dfm-rbac-retries` is 12 iterations.



The Operations Manager refresh interval depends on the number of storage systems, number of storage objects in the storage system, and the load on the DataFabric Manager server.

As a recommendation, perform the following:

- a. From the DataFabric Manager server, manually run the following command for all the secondary storage systems associated with the dataset: `dfm host discover storage_system`
- b. Double the time taken to perform the host discovery operation and assign that value to `dfm-rbac-retry-sleep-secs`.

For example, if the operation took 11 seconds, you can set the value of `dfm-rbac-retry-sleep-secs` to 22 (11*2).

Repository database size grows with time and not with the number of backups

The repository database size grows with time because SnapManager operations insert or delete data within the schema in the repository database tables, which results in high index space usage.

Workaround

You must monitor and rebuild the indexes according to the Oracle guidelines to control the space consumed by the repository schema.

The SnapManager GUI cannot be accessed and SnapManager operations fail when the repository database is down

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

The following table lists the different actions you might want to perform, and their exceptions:

Operations	Exceptions
Opening a closed repository	The following error message is logged in sm_gui.log: [WARN]: SMO-01106: Error occurred while querying the repository: Closed Connection java.sql.SQLException: Closed Connection.
Refreshing an opened repository by pressing F5	A repository exception is displayed in the GUI and also logs a NullPointerException in the sm_gui.log file.
Refreshing the host server	A NullPointerException is logged in the sumo_gui.log file.
Creating a new profile	A NullPointerException is displayed in the Profile Configuration window.
Refreshing a profile	The following SQL exception is logged in sm_gui.log: [WARN]: SMO-01106: Error occurred while querying the repository: Closed Connection.
Accessing a backup	The following error message is logged in sm_gui.log: Failed to lazily initialize a collection.
Viewing clone properties	The following error message is logged in sm_gui.log and sumo_gui.log: Failed to lazily initialize a collection.

Workaround

You must ensure that the repository database is running when you want to access the GUI or want to perform any SnapManager operations.

Unable to create temporary files for the cloned database

When temporary tablespace files of the target database are placed in mount points different from the mount point of the data files, the clone create operation is successful but SnapManager fails to create temporary files for the cloned database.

Workaround

You must perform either of the following:

- Ensure that the target database is laid out so that temporary files are placed in the same mount point as that of the data files.
- Manually create or add temporary files in the cloned database.

Unable to migrate the protocol from NFSv3 to NFSv4

You can migrate the protocol from NFSv3 to NFSv4 by enabling the `enable-migrate-nfs-version` parameter in the `snapdrive.conf` file. During the migration, SnapDrive considers only the protocol version, irrespective of the mount point options such as `rw`, `largefiles`, `nosuid`, and so on.

However, after migrating the protocol to NFSv4, when you restore the backup that was created by using NFSv3, the following occurs:

- If NFSv3 and NFSv4 are enabled at the storage level, the restore operation is successful but is mounted with the mount point options that were available during backup.
- If only NFSv4 is enabled at the storage level, the restore operation is successful and only the protocol version (NFSv4) is retained.

However, the other mount point options such as `rw`, `largefiles`, `nosuid`, and so on are not retained.

Workaround

You must manually shut down the database, unmount the database mount points, and mount with the options available prior to the restore.

Back up of Data Guard Standby database fails

If any archive log location is configured with the service name of the primary database, the back up of Data Guard Standby database fails.

Workaround

In the GUI, you must clear **Specify External Archive Log location** corresponding to the service name of the primary database.

Copyright Information

Copyright © 2021 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, 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.