

Troubleshooting SnapManager

SnapManager Oracle

NetApp August 30, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapmanageroracle/windows/task_creating_operation_level_dump_files.html on August 30, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Troubleshooting SnapManager	1
Dump files	7
Troubleshooting clone issues	12
Troubleshooting graphical user interface issues	15
Troubleshooting known issues	20
Running multiple parallel operations fails in SnapManager	23
Unable to restore RAC database from one of the RAC nodes where the profile was not created $\ .$	23
Where to go for more information	24

Troubleshooting SnapManager

You can find information about some of the most common issues that might occur and how you can resolve them.

The following table describes common issues and possible solutions:

Issue-driven question	Possible solution
Are the target database and listener running?	Run the Isnrctl status command. Ensure that the database instance is registered with the listener.
Is the storage visible?	 Perform the following steps: 1. Right-click My Computer, and then select Manage. 2. Click storage > snapdrive > Hostname > Disks.
Is the SnapManager server running?	 Check the status, and then start the server by using Service Configuration. Before you can use the graphical user interface (GUI) or the command-line interface (CLI) to initiate SnapManager commands related to profiles, the server must be running. You can create or update repositories without starting the server, but to execute all other SnapManager operations, the server must be running. To start the SnapManager server, enter the following command: smo_server start.
Are all the components required to run SnapManager set up correctly?	Run the smo system verify command to verify that SnapDrive is set up correctly.
Do you have the correct version of SnapManager?	Use the smo version command to check the SnapManager version.

Have you looked at the SnapManager log files to determine if the error messages can help isolate the issue?	SnapManager records all log entries into one set of rotating log files. The log files are found at C:\program_files\NetApp\SnapManager for Oracle\logs. If you are using Windows 2008, the logs are at the following locations: • Operation logs: • C:\Program Files\NetApp\SnapManager for Oracle\var\log\smo • Client logs: • C:\Users\Administrator\AppData\Roaming\Net App\smo\3.3.0\ It might also be helpful to look at the logs in the following location: C:\Documents and Settings\hostname\Application Data\NetApp\smo\3.3.0\log Each operation log is written to its own log file of the form smo_of_date_time.log.
If you have archive logs stored on a storage system that is not running Data ONTAP, have you excluded them from consideration for backup with SnapManager?	The smo.config file enables you to exclude certain archive log files. For Windows, the file is at the following location: C:\program_files\NetApp\smo\properties\smo.config Use the format mentioned in the file to exclude the local archive logs. For additional information, see the "Setting configuration properties" topic. You can also exclude the archive log destinations while creating a backup from the SnapManager CLI. For additional information, see the "Creating database backups" topic. You can also exclude the archive log destinations while creating a backup from the SnapManager GUI.
Do you have an MS-DOS window open in the directory in which you are attempting to install or upgrade SnapManager on Windows?	You will see an error message similar to the following: Directory C:\Program Files\NetApp\SnapManager for Oracle\bin is currently in use by another program. Any window, opened by you or another user, that is currently referencing this directory must be closed before installation can proceed. Close the window and attempt the installation or upgrade again.

Were you unable to connect to the repository?	If connecting to a repository fails, run the Isnrctl status command on the repository database and check the active service names. When SnapManager connects to the repository database, it uses the service name of the database. Depending on how the listener is setup, this might be the short service name or the fully qualified service name. When SnapManager connects to a database for a backup, restore, or other operation, it uses the host name and the SID. If the repository does not initialize correctly because it is currently unreachable, you receive an error message asking whether you want to remove the repository. You can remove the repository from your current view so that you can perform operations on other repositories. Also, check whether the corresponding service is running.
Can system resolve the host name?	Check whether the specified host name is on a different subnet. If you receive an error message that SnapManager cannot resolve the host name, then add the host name in the host file. Add the host name to the file located at C:\WINDOWS\system32\drivers\etc\hosts: xxx.xxx.xxx hostname IP address
Is SnapDrive running?	To view the status of SnapDrive, go to Services, and then select the SnapDrive service.
Which storage systems are configured to be accessed with SnapDrive?	 To find the storage systems configured for SnapDrive, perform the following steps: 1. Right-click My Computer, and then select Manage. 2. Click Storage > SnapDrive. 3. Right-click the host name, and then select transport protocol settings.

How can SnapManager GUI performance be improved?	• Ensure that you have valid user credentials for the repository, profile host, and profile.
	If your credential is invalid, then clear the user credentials for the repository, profile host, and profile. Reset the same user credentials that you set before for the repository, profile host, and profile. For additional information about setting the user credentials again, see "Setting credentials after clearing credential cache".
	Close the unused profiles.
	If the number of profiles that you have opened is more, the SnapManager GUI performance slows down.
	 Check whether you enabled Open On Startup in the User Preferences window under the Admin menu, from the SnapManager GUI.
	If this is enabled, then the user configuration (user.config) file available at C:\Documents and Settings\Administrator\Application Data\NetApp\smo\3.3.0\gui\state is displayed as openOnStartup=PROFILE.
	Because Open On Startup is enabled, you must check for recently opened profiles from the SnapManager GUI, using lastOpenProfiles in the user configuration (user.config) file: lastOpenProfiles=PROFILE1,PROFILE2,PROFIL E3,
	You can delete the profile names listed and always keep a minimum number of profiles as open.
	 Before installing the new version of SnapManager on the Windows-based environment, delete the SnapManager client-side entries available at the following location:
	C:\Documents and Settings\Administrator\Application Data\NetApp
SnapManager GUI takes more time to refresh when there are multiple SnapManager operations started and running simultaneously in the background. When you right-click the backup (that is already deleted but still gets displayed in the SnapManager GUI), the backup options for that backup are not enabled in the Backup or Clone window.	You need to wait until the SnapManager GUI gets refreshed, and then check for the backup status.

What would you do when the Oracle database is not set in English?	 SnapManager operations might fail if the language for an Oracle database is not set to English.Set the language of the Oracle database to English: 1. Verify that the NLS_LANG environment variable is not set: echo%NLS_LANG% 2. Add the following line to the wrapper.conf file located at C:\SnapManager_install_directory\service: set.NLS_LANG=AMERICAN_AMERICA.WE8MS WIN1252 3. Restart the SnapManager server: smo_server restart If the system environment variable is set to NLS_LANG, you must edit the script to not overwrite NLS_LANG. 	
What would you do when the backup scheduling operation fails if the repository database points to more than one IP and each IP has a different host name?	 script to not overwrite NLS_LANG. Stop the SnapManager server. Delete the schedule files in the repository directory from the hosts where you want to trigge the backup schedule. The schedule file names can be in the following formats: repository#repo_username#repository_datal se_name#repository_host#repo_port repository-repo_usernamerepository_database_name-repository_host-repo_port Note: You must ensure that you delete the schedule file in th format that matches the repository details. Restart the SnapManager server. Open other profiles under the same repository from the SnapManager GUI to ensure that you c not miss any schedule information of those profiles. 	

What would you do when the SnapManager operation fails with credential file lock error?	SnapManager locks the credential file before updating, and unlocks it after updating.When multiple operations run simultaneously, one of the operations might lock the credential file to update it. If another operation tries to access the locked credential file at the same time, the operation fails with the file lock error.	
	Configure file depend operations	the following parameters in the smo.config ding on the frequency of simultaneous :
	 fileLoc 	k.retryInterval = 100 milliseconds
	 fileLoc 	k.timeout = 5000 milliseconds
	i	The values assigned to the parameters must be in milliseconds.
What would you do when the backup verify operation's intermediate status shows failed in the Monitor tab even though the backup verify operation is still running?	The error message is logged in the sm_gui.log file. You must look in the log file to determine the new values for the operation.heartbeatInterval and operation.heartbeatThreshold parameters which will resolve this issue.	
	1. Add th file:	e following parameters in the smo.config
	∘ op	eration.heartbeatInterval = 5000
	∘ op de 500	eration.heartbeatThreshold = 5000 The fault value assigned by SnapManager is 00.
	2. Assign	the new values to these parameters.
	()	The values assigned to the parameters must be in milliseconds.
	3. Restar operat	t the SnapManager server and perform the ion again.

What to do when you encounter a heap-space issue?	When you encounter a heap-space issue during SnapManager for Oracle operations, you must perform the following steps:	
	 Navigate to the SnapManager for Oracle installation directory. 	
	 Open the launchjava file from the installationdirectory\bin\launchjava path. 	
	 Increase the value of the java -Xmx160m Java heap-space parameter. 	
	For example, you can increase the default value of 160m to 200m.	
	(i) If you have increased the value of the Java heap-space parameter in the earlier versions of SnapManager for Oracle, you should retain that value.	
What would you do when the SnapManager services do not start in a Windows environment and the following error message is displayed: Windows could not start Snap Manager on Local Computer. For more information, review the System Event log. If this is a non-Microsoft service, contact service vendor, and refer to service-specific error code 1?	Configure the following parameters in the wrapper.conf file located at Installation_directory\service.	
	• The wrapper startup timeout parameter defines the maximum permissible time between the wrapper starting the Java Virtual Machine (JVM) and response from the JVM that the application has started.	
	The default value is set to 90 seconds. However, you can change a value greater than 0. If you specify an invalid value, the default is used instead.	
	• The wrapper.ping.timeout parameter defines the maximum permissible time between the wrapper pinging the JVM and the response from the JVM. The default value is set to 90 seconds.	
	However, you can change to a value greater than 0. If you specify an invalid value, the default is used instead.	

Dump files

The dump files are compressed log files containing information about SnapManager and its environment. The different types of log files created are operation, profile, and system dump file.

You can use the dump command or the **Create Diagnostics** tab in the graphical user interface (GUI) to collect information about an operation, a profile, or the environment. A system dump does not require a profile; however, the profile and operation dumps require profiles.

SnapManager includes the following diagnostic information in the dump file:

- The steps performed
- · The length of time for each step to complete
- The outcome of each step
- Error, if any, that occurred during the operation



SnapManager log files or dump files enable read and write permissions only for the root users and the other users who belong to root user group.

SnapManager also includes the following information in the file:

- · Operating system version and architecture
- · Environment variables
- Java version
- SnapManager version and architecture
- SnapManager preferences
- SnapManager messages
- log4j properties
- SnapDrive version and architecture
- SnapDrive log files
- Oracle version
- · Oracle OPatch local inventory details
- · Repository database Oracle version
- Target database type (stand alone)
- Target database role (primary, physical standby, or logical standby)
- Target database Oracle Recovery Manager (RMAN) setup (no RMAN integration, RMAN with control files, or RMAN with catalog file)
- Target database Oracle version
- · System identifier (SID) of the target database
- · RMAN database name and TNS connection name
- · Repository database service name
- · Database instances installed on the host
- · Profile descriptor
- · Shared memory maximum
- Swap space information
- · Memory information

- Multipath environment
- Host utilities version
- Microsoft Internet Small Computer System Interface (iSCSI) software initiator version for Windows
- · Output of the system verify command

The dump file also lists the SnapManager limitations on Windows.

SnapManager dump files also contain the SnapDrive data collector file and the Oracle alert log file. You can collect the Oracle alert log file by using the smo operation dump and smo profile dump commands.



System dump does not contain Oracle alert logs; however, the profile and operation dumps contain the alert logs.

Even if the SnapManager host server is not running, you can access the dump information by using the command-line interface (CLI) or the GUI.

If you encounter a problem that you cannot resolve, you can send these files to NetApp Global Services.

Creating operation-level dump files

You can use the smo operation dump command with the name or ID of the failed operation to get log information about a particular operation. You can specify different log levels to gather information about a specific operation, profile, host, or environment.

1. Enter the following command:smo operation dump -idguid



The smo operation dump command provides a super set of the information provided by the smo profile dump command, which in turn provides a super set of the information provided by the smo system dump command.

Dump file location:

```
Path:\<user-home>\Application
Data\NetApp\smo\3.3.0\smo dump 8abc01c814649ebd0114649ec69d0001.jar
```

Creating profile-level dump files

You can find log information about a particular profile by using the smo profile dump command with the name of the profile.

1. Enter the following command: smo profile dump -profile profile_name

Dump file location:

```
Path:\<user-home>\Application
Data\NetApp\smo\3.3.0\smo dump 8abc01c814649ebd0114649ec69d0001.jar
```



If you encounter an error while creating a profile, use the smosystem dump command. After you have successfully created a profile, use the smooperation dump and smoprofile dump commands.

Creating system-level dump files

You can use the smo system dump command to get log information about the SnapManager host and environment. You can specify different log levels to collect information about a specific operation, profile, or host and environment.

1. Enter the following command: smo system dump

Resulting dump

```
Path:\<user-home>\Application
Data\NetApp\smo\3.3.0\smo_dump_server_host.jar
```

How to locate dump files

The dump file is located at the client system for easy access. These files are helpful if you need to troubleshoot a problem related to profile, system, or any operation.

The dump file is located in the user's home directory on the client system.

• If you are using the graphical user interface (GUI), the dump file is located at:

```
user_home\Application Data\NetApp\smo\3.3.0\smo_dump dump_file_type_name
server_host.jar
```

• If you are using the command-line interface (CLI), the dump file is located at:

```
user_home\.netapp\smo\3.3.0\smo_dump_dump_file_type_name server_host.jar
```

The dump file contains the output of the dump command. The name of the file depends on the information supplied. The following table shows the types of dump operations and the resulting file names:

Type of dump operation	Resulting file name
Operation dump command with operation ID	smo_dump_operation-id.jar

Operation dump command with no operation ID	<pre>smo operation dump -profile VH1-verbose The following output is displayed: smo operation dump -profile VH1 -verbose [INFO] SMO-13048: Dump Operation Status: SUCCESS [INFO] SMO-13049: Elapsed Time: 0:00:01.404 Dump file operated</pre>	
	Path: user_home\Application Data\ontap\smo\3.3.0\smo_dump_VH1_ kaw.rtp.foo.com.jar	
System dump command	smo_dump_host-name.jar	
Profile dump command	smo_dump_profile-name_host-name.jar	

How to collect dump files

You can include -dump in the SnapManager command to collect the dump files after a successful or failed SnapManager operation.

You can collect dump files for the following SnapManager operations:

- Creating profiles
- · Updating profiles
- Creating backups
- Verifying backups
- Deleting backups
- Freeing backups
- Mounting backups
- Unmounting backups
- Restoring backups
- Creating clones
- Deleting clones



When you create a profile, you can collect dump files only if the operation is successful. If you encounter an error while creating a profile, you must use the smosystem dump command. For successful profiles, you can use the smooperation dump and smoprofile dump commands to collect the dump files.

Example

smo backup create -profile targetdb1_prof1 -auto -full -online
-dump

Collecting additional log information for easier debugging

If you need additional logs to debug a failed SnapManager operation, you must set an external environment variable server.log.level. This variable overrides the default log level and dumps all the log messages in the log file. For example, you can change the log level to DEBUG, which logs additional messages and can assist in debugging issues.

The SnapManager logs can be found at the following locations:

SnapManager_install_directory\log

To override the default log level, you must perform the following steps:

- 1. Create a platform.override text file in the SnapManager installation directory.
- 2. Add the server.log.level parameter in the platform.override text file.
- 3. Assign a value (TRACE, DEBUG, INFO, WARN, ERROR, FATAL, or PROGRESS) to the server.log.level parameter.

For example, to change the log level to ERROR, set the value of server.log.level to ERROR.

server.log.level=ERROR

4. Restart the SnapManager server.



If the additional log information is not required, you can delete the server.log.level parameter from the platform.override text file.

SnapManager manages the volume of server log files based on the user-defined values of the following parameters in the smo.config file:

- log.max_log_files
- log.max_log_file_size
- log.max_rolling_operation_factory_logs

Troubleshooting clone issues

You can find information about that might occur during clone operations and how you can resolve them.

Symptom	Explanation	Workaround
---------	-------------	------------

The clone operation fails when the archive destination is set to USE_DB_RECOVERY_FILE_DES T.	When the archive destination is referring to USE_DB_RECOVERY_FILE_DES T, Flash recovery area (FRA) actively manages the archive log. SnapManager does not use the FRA location during clone or restore operations and thus the operations fail.	Change the archive destination to actual archive log location instead of the FRA location.

The clone operation fails with the following error message: Cannot perform operation: Clone Create.	This issue occurs if Oracle listener fails to connect to the database.	If you are using SnapManager GUI to clone a backup, perform the following actions:
executing SQL: [ALTER DATABASE OPEN RESETLOGS;]. The command returned: ORA-		 From the Repository tree, click Repository > Host > Profile to display the backups.
01195: online backup of file 1 needs more recovery to be consistent		2. Right-click the backup that you want to clone and select Clone .
		3. On the Clone Initialization page, enter the mandatory values and select the clone specification method.
		 On the Clone Specification page, select Parameters.
		5. Click the +Parameter tab.
		 In the Parameter Name field, enter the name as local_listener and click OK.
		7. Select the Override Default check box for the local_listener row.
		 Click any parameter, then double-click the local_listener parameter, and enter the following value:(ADDRESS=(PROTOCO L=TCP)(HOST=<your_host_na me>)(PORT=<port#>))</port#></your_host_na
		9. Click Save To File.
		10. Click Next and continue with the clone create wizard.
		If you are using CLI to clone a backup, you must include the following information in the <parameters></parameters> tag of the clone specification file:

The clone operation fails with an error message saying that the mountpoint you are using is already in use.	SnapManager does not let you mount a clone over an existing mountpoint. So an incomplete clone did not remove the mountpoint.	Specify a different mountpoint to be used by the clone, or unmount the problematic mountpoint.
The clone operation fails with an error message about data files not having a .dbf extension.	Some versions of the Oracle NID utility do not work with data files unless the files use a .dbf extension.	 Rename the data file to give it a .dbf extension. Repeat the backup operation. Clone the new backup.
The clone operation fails due to unmet requirements.	You are attempting to create a clone; however, some of the prerequisites have not been met.	Proceed as described in <i>Creating a clone</i> to meet the prerequisites.
SnapManager for Oracle fails to clone Oracle 10gR2 (10.2.0.5) physical Oracle Data Guard Standby databases.	SnapManager for Oracle does not disable the managed recovery mode while performing an offline backup of the Oracle 10gR2 (10.2.0.5) physical standby databases created using Oracle Data Guard services. Due to this issue, the offline backup taken is inconsistent. When SnapManager for Oracle tries to clone the offline backup, it does not even try to perform any recovery on the cloned database. Because the backup is inconsistent, the cloned database requires recovery, and thus Oracle fails to create the clone successfully.	Upgrade the Oracle database to the Oracle 11gR1 (11.1.0.7 patch).

Troubleshooting graphical user interface issues

You can find information about some common known graphical user interface (GUI) issues that might help you resolve them.

Explanation Workaround	Issue	Explanation	Workaround
------------------------	-------	-------------	------------

While accessing the SnapManager GUI to perform an operation, the following error message might be displayed: SMO-20111 : Authentication failed for user on host.	This issue occurs if the password of the user is changed in the host on which the SnapManager server is running.After the password is changed, the credential cache that is created for the user who launched the GUI becomes invalid. SnapManager GUI still uses the credentials in the cache to authenticate and thus the authentication fails.	 You must perform one of following tasks: Delete the credentials of the user whose password was changed and then add the new credentials in the cache by running the following commands: a. smo credential delete b. smo credential delete b. smo credential set Clear the entire cache by running the smo credential clear command. Re-open the GUI and set the credentials, if prompted.
Security warning is displayed while using Java Web Start to access SnapManager GUI.	While accessing SnapManager GUI using Java Web Start, a security warning is displayed.This issue occurs because JNLP jars are self-signed and Java version used by SnapManager does not allow self-signed jars at high security level.	Either change the security settings to medium in the java control panel or add the SnapManager GUI URL to the exception list.
The SnapManager web start GUI displays the incorrect version.	After downgrading SnapManager from a later version to an earlier version when you launch the web start GUI, the later version of the SnapManager web start GUI is launched.	 You must also clear the cache by performing the following steps: 1. Click Start and select Run. 2. Enter the following: javaws -viewer 3. On the Java cache viewer screen, right-click the SnapManager application and select Delete.

When you restart the GUI and try to check the backups for a certain profile, you see only the names of the profiles.	SnapManager does not display any information about a profile until you open it.	Perform t 1. Right selec Snap Profil box. 2. Enter passy Snap	the following steps: -click the profile and t Open from the menu. Manager displays the e Authentication dialog the host user name and word. Manager displays the
		backı i	You only need to authenticate the profile once as long as the credentials are valid and remain in the cache.
The GUI installation on Windows succeeds, but with errors.	The user account used to install the GUI does not have enough permissions to set up the icons and shortcuts for all the users. The user account does not have permissions to modify the C:\Documents and Settings\All Users directory.	 Reins settin In Ch clear short users Log in that is reinst 	stall the GUI with different gs. oose Icon Availability, the Make these tcuts available to all s of this PC check box. In using a user account s not restricted and tall the GUI.
When you open the first repository in the GUI, an error message similar to the following is displayed: The Profile name XXXX clashes with previously loaded repository.	Identically named profiles cannot exist in a repository. Also, you can open only one repository at a time.	Reference from two (OS) use issuing and repository SMO_33 'NEW_NA 'OLD_NA	e the conflicting profiles different operating system rs or rename the profile by n SQL statement for the y: UPDATE _PROFILE SET NAME = AME' WHERE NAME = ME'

An error message similar to the following is displayed: SMO-01092: Unable to initialize repository repo1@ does not exist:repo1SMO- 11006: Cannot resolve host does not exist	The repository is inaccessible, perhaps because it no longer exists. The GUI initializes the list of repositories from the credentials file.	The error message asks if you would like to remove this repository so that no attempt is made to load it in the future. If you do not need to access this repository, click Delete to remove it from the GUI view. This removes the reference to the repository in the credentials file and the GUI does not attempt to load the repository again.
SnapManager takes a longer time to load the database tree structure and results in a timeout error message being displayed on the SnapManager GUI.	When you try to perform a partial backup operation from the SnapManager GUI, SnapManager tries to load the credentials for all the profiles, and if there are any invalid entries, SnapManager tries to validate the entry and this results in a timeout error message being displayed.	Delete the credentials of the unused host, repository, and profile by using the credential delete command from the SnapManager command-line interface (CLI).
The custom scripts for the preprocessing or postprocessing activity to occur before or after the backup, restore, or clone operations, are not visible from the SnapManager GUI.	When you add custom scripts in the custom backup, restore, or clone script location after you start the respective wizard, the custom scripts are not displayed under the Available Scripts list.	Restart the SnapManager host server and then open the SnapManager GUI.
You cannot use the clone specification XML file created in SnapManager (3.1 or earlier) for the clone operation.	From SnapManager 3.2 for Oracle, the task specification section (task- specification) is provided as a separate task specification XML file.	If you are using SnapManager 3.2 for Oracle, you must remove the task specification section from the clone specification XML or create a new clone specification XML file.SnapManager 3.3 or later does not support the clone specification XML file created in SnapManager 3.2 or earlier releases.

SnapManager operation on the GUI does not proceed after you have cleared user credentials by using the smo credential clear command from the SnapManager CLI or by clicking Admin > Credentials > Clear > Cache from the SnapManager GUI.	The credentials set for the repositories, hosts, and profiles are cleared. SnapManager verifies user credentials before starting any operation.When user credentials are invalid, SnapManager fails to authenticate. When a host or a profile is deleted from the repository, the user credentials are still available in the cache. These unnecessary credential entries slow down the SnapManager operations from the GUI.	 Restart the SnapManager GUI depending on how the cache is cleared. Note: If you have cleared the credential cache from the SnapManager GUI, you do not need to exit the SnapManager GUI. If you have cleared the credential cache from the SnapManager CLI, you must restart the SnapManager GUI. If you have deleted the encrypted credential file manually, you must restart the SnapManager GUI. Set the credentials that you have given for the repository, profile host, and profile. From the SnapManager GUI, if there is no repository mapped under the Repositories tree, perform the following steps: Click Tasks > Add Exisiting repository Right-click the repository, click Open, and enter the user credentials in the Repository Credentials Authentication window. Right-click the host under the repository, click Open, and enter the user credentials in Host Credentials Authentication. Right-click the profile under the host, click Open, and enter the user credentials in Profile Credentials Authentication.
You cannot open the SnapManager GUI by using Java Web Start GUI due to weaker Secure Sockets Layer (SSL) cipher strength of the browser.	SnapManager does not support SSL ciphers weaker than 128 bits.	Upgrade the browser version and check the cipher strength.

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: C:\Program Files\NetApp\SnapManager for Oracle\properties\smo.config

You assign the following values:

- SMO Server.port=27214
- SMO Server.rmiRegistry.port=27215
- remote.registry.ocijdbc.port= 27215

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

To start the SnapManager server, perform the following steps:

- 1. Click Start > Control Panel > Administrative Tools > Services.
- 2. You can start the server in one of three ways:
 - In the left panel, click Start.
 - Right-click NetApp SnapManager 3.3 for Oracle and select Start from the drop-down menu.
 - Double-click NetApp SnapManager 3.3 for Oracle and in the Properties window that opens, click Start.

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: E:\\arch, G:\\arch, and H:\\arch. While creating the archive log file backup, if you exclude E:\\arch by using the command

```
smo backup create -profile almsamp1 -data -online -archivelogs -exclude
-dest E:\\arch
```

, SnapManager for Oracle excludes all the destinations starting with E:\\arch.

Workaround

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

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.

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.

Running multiple parallel operations fails in SnapManager

When you run multiple parallel operations on separate databases that reside on the same storage system, the igroup for LUNs associated with both the databases might get deleted because of one of the operations. Later, if the other operation attempts to use the deleted igroup, SnapManager displays an error message.

For example, if you are running the backup delete and backup create operations on different databases almost at the same time, the backup create operation fails. The following sequential steps show what happens when you run backup delete and backup create operations on different databases almost at the same time.

- 1. Run the backup delete command.
- 2. Run the backup create command.
- 3. The backup create command identifies the already existing igroup and uses the same igroup for mapping the LUN.
- 4. The backup delete command deletes the backup LUN, which was mapped to the same igroup.
- 5. The backup delete command then deletes the igroup because there are no LUNs associated with the igroup.
- 6. The backup create command creates the backup and tries to map to the igroup that does not exist, and therefore the operation fails.

What to do

You must create igroup for each storage system used by the database by using the following command: sdcli igroup create

Unable to restore RAC database from one of the RAC nodes where the profile was not created

In an Oracle RAC environment where both nodes belong to the same cluster, if you attempt a restore operation from a node which is different from the node where the backup was created, the restore operation fails.

For example, if you create a backup in Node A and try to restore from Node B, the restore operation fails.

What to do

Before performing restore operation from node B, perform the following in node B:

- 1. Add the repository.
- 2. Sync the profile by running the command smo profile sync.
- 3. Set the credential for the profile to be used for restore operation by running the command smo credential set.
- 4. Update the profile to add the new host name and the corresponding SID by running the command smo profile update.

Where to go for more information

You can find information about the basic tasks involved in installing and using SnapManager.

Document	Description
SnapManager description page	This page provides information about SnapManager, pointers to online documentation, and a link to the SnapManager download page, from which you can download the software.
Data ONTAP SAN Configuration Guide for 7-Mode	This document is available at mysupport.netapp.com. It is a dynamic, online document that contains the most up-to-date information about the requirements for setting up a system in a SAN environment. It provides the current details about storage systems and host platforms, cabling issues, switch issues, and configurations.
SnapManager and SnapDrive Compatibility Matrix	This document is available in the Interoperability section at mysupport.netapp.com/matrix. It is a dynamic, online document that contains the most up-to-date information specific to SnapManager and its platform requirements.
SnapManager Release Notes	This document comes with SnapManager. You can also download a copy from mysupport.netapp.com. It contains any last-minute information that you need to get the configuration up and running smoothly.
NetApp host attach and support kits documentation	mysupport.netapp.com.
System Configuration Guide	mysupport.netapp.com.
Data ONTAP Block Access Management Guide	mysupport.netapp.com
Host operating system and database information	These documents provide information about your host operating system and database software.

Copyright information

Copyright © 2024 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.