



Backup configuration information

SnapCenter Software 6.0

NetApp
July 25, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapcenter/protect-sco/supported_backup_configs.html on July 25, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Backup configuration information 1
 - Supported Oracle database configurations for backups 1
 - Types of backup supported for Oracle databases 1
 - How SnapCenter discovers Oracle databases 2
 - Preferred nodes in RAC setup 3
 - How to catalog backups with Oracle Recovery Manager. 4
 - Predefined environment variables for backup specific prescript and postscript 6
 - Backup retention options 11
 - Backup schedules 11
 - Backup naming conventions 12

Backup configuration information

Supported Oracle database configurations for backups

SnapCenter supports backup of different Oracle database configurations.

- Oracle Standalone
- Oracle Real Application Clusters (RAC)
- Oracle Standalone Legacy
- Oracle Standalone Container Database (CDB)
- Oracle Data Guard standby

You can only create offline-mount backups of Data Guard standby databases. Offline-shutdown backup, archive log only backup, and full backup are not supported.

- Oracle Active Data Guard standby

You can only create online backups of Active Data Guard standby databases. Archive log only backup and full backup are not supported.

Before creating a backup of Data Guard standby or Active Data Guard standby database, the managed recovery process (MRP) is stopped and once the backup is created, MRP is started.

- Automatic Storage Management (ASM)
 - ASM standalone and ASM RAC on Virtual Machine Disk (VMDK)

Among all the restore methods supported for Oracle databases, you can perform only connect-and-copy restore of ASM RAC databases on VMDK.

- ASM standalone and ASM RAC on Raw device mapping (RDM)
You can perform backup, restore, and clone operations on Oracle databases on ASM, with or without ASMLib.
 - Oracle ASM Filter Driver (ASMFD)

PDB migration and PDB cloning operations are not supported.

- Oracle Flex ASM

For the latest information about supported Oracle versions, see the [NetApp Interoperability Matrix Tool](#).

Types of backup supported for Oracle databases

Backup type specifies the type of backup that you want to create. SnapCenter supports online and offline backup types for Oracle databases.

Online backup

A backup that is created when the database is in the online state is called an online backup. Also called a hot backup, an online backup enables you to create a backup of the database without shutting it down.

As part of online backup, you can create a backup of the following files:

- Data files and control files only
- Archive log files only (the database is not brought to backup mode in this scenario)
- Full database that includes data files, control files, and archive log files

Offline backup

A backup created when the database is either in a mounted or shutdown state is called an offline backup. An offline backup is also called a cold backup. You can include only data files and control files in offline backups. You can create either an offline mount or offline shutdown backup.

- When creating an offline mount backup, you must ensure that the database is in a mounted state.

If the database is in any other state, the backup operation fails.

- When creating an offline shutdown backup, the database can be in any state.

The database state is changed to the required state to create a backup. After creating the backup, the database state is reverted to the original state.

How SnapCenter discovers Oracle databases

Resources are Oracle databases on the host that are maintained by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available.

The following sections describe the process that SnapCenter uses to discover different types and versions of Oracle databases.

For Oracle versions 11g to 12cR1

RAC database

The RAC databases are discovered only on the basis of `/etc/oratab` entries. You should have the database entries in the `/etc/oratab` file.

Standalone

The standalone databases are discovered only on the basis of `/etc/oratab` entries.

ASM

The ASM instance entry should be available in the `/etc/oratab` file.

RAC One Node

The RAC One Node databases are discovered only on the basis of `/etc/oratab` entries. The databases should be either in `nomount`, `mount`, or `open` state. You should have the database entries in the `/etc/oratab` file.

The RAC One Node database status will be marked as `renamed` or `deleted` if the database is already discovered and backups are associated with the database.

You should perform the following steps if the database is relocated:

1. Manually add the relocated database entry in the `/etc/oratab` file on the failed-over RAC node.
2. Manually refresh the resources.
3. Select the RAC One Node database from the resource page, and then click Database Settings.
4. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database.
5. Perform the SnapCenter operations.
6. If you have relocated a database from one node to another node and if the `oratab` entry in the earlier node is not deleted, manually delete the `oratab` entry to avoid the same database being displayed twice.

For Oracle versions 12cR2 to 18c

RAC database

The RAC databases are discovered using the `srvctl config` command. You should have the database entries in the `/etc/oratab` file.

Standalone

The standalone databases are discovered based on the entries in the `/etc/oratab` file and the output of the `srvctl config` command.

ASM

The ASM instance entry need not be in the `/etc/oratab` file.

RAC One Node

The RAC One Node databases are discovered using the `srvctl config` command only. The databases should be either in `nomount`, `mount`, or `open` state. The RAC One Node database status will be marked as `renamed` or `deleted` if the database is already discovered and backups are associated with the database.

You should perform the following steps if the database is relocated: . Manually refresh the resources. . Select the RAC One Node database from the resource page, and then click Database Settings. . Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database. . Perform the SnapCenter operations.



If there are any Oracle 12cR2 and 18c database entries in the `/etc/oratab` file and the same database is registered with the `srvctl config` command, SnapCenter will eliminate the duplicate database entries. If there are stale database entries, the database will be discovered but the database will be unreachable and the status will be offline.

Preferred nodes in RAC setup

In Oracle Real Application Clusters (RAC) setup, you can specify the preferred nodes that SnapCenter uses to perform the backup operation. If you do not specify the preferred node, SnapCenter automatically assigns a node as the preferred node and backup is created on that node.

The preferred nodes might be one or all of the cluster nodes where the RAC database instances are present. The backup operation is triggered only on these preferred nodes in the order of the preference.

Example

The RAC database `cdbrac` has three instances: `cdbrac1` on `node1`, `cdbrac2` on `node2`, and `cdbrac3` on `node3`.

The node1 and node2 instances are configured to be the preferred nodes, with node2 as the first preference and node1 as the second preference. When you perform a backup operation, the operation is first attempted on node2 because it is the first preferred node.

If node2 is not in the state to back up, which could be due to multiple reasons such as the plug-in agent is not running on the host, the database instance on the host is not in the required state for the specified backup type, or the database instance on node2 in a FlexASM configuration is not being served by the local ASM instance; then the operation will be attempted on node1.

The node3 will not be used for backup because it is not on the list of preferred nodes.

Flex ASM setup

In a Flex ASM setup, Leaf nodes will not be listed as preferred nodes if the cardinality is less than the number nodes in the RAC cluster. If there is any change in the Flex ASM cluster node roles, you should manually discover so that the preferred nodes are refreshed.

Required database state

The RAC database instances on the preferred nodes must be in the required state for the backup to finish successfully:

- One of the RAC database instances in the configured preferred nodes must be in the open state to create an online backup.
- One of the RAC database instances in the configured preferred nodes must be in the mount state, and all other instances, including other preferred nodes, must be in the mount state or lower to create an offline mount backup.
- RAC database instances can be in any state, but you must specify the preferred nodes to create an offline shutdown backup.

How to catalog backups with Oracle Recovery Manager

You can catalog the backups of Oracle databases using Oracle Recovery Manager (RMAN) to store the backup information in the Oracle RMAN repository.

The cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. When you do not need these cataloged backups, you can remove the catalog information.

The database must be in mounted or higher state for cataloging. You can perform cataloging on data backups, archive log backups, and full backups. If cataloging is enabled for a backup of a resource group that has multiple databases, cataloging is performed for each database. For Oracle RAC databases, cataloging will be performed on the preferred node where the database is at least in mounted state.

If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

External catalog database

By default, the target database control file is used for cataloging. If you want to add external catalog database, you can configure it by specifying the credential and Transparent Network Substrate (TNS) name of the external catalog using the Database Settings wizard from the SnapCenter graphical user interface (GUI). You can also configure the external catalog database from the CLI by running the `Configure-SmOracleDatabase`

command with the `-OracleRmanCatalogCredentialName` and `-OracleRmanCatalogTnsName` options.

RMAN command

If you enabled the cataloging option while creating an Oracle backup policy from the SnapCenter GUI, the backups are cataloged using Oracle RMAN as a part of the backup operation. You can also perform deferred cataloging of backups by running the `Catalog-SmBackupWithOracleRMAN` command.

After cataloging the backups, you can run the `Get-SmBackupDetails` command to obtain the cataloged backup information such as the tag for cataloged datafiles, the control file catalog path, and the cataloged archive log locations.

Naming format

If the ASM disk group name is greater than or equal to 16 characters, from SnapCenter 3.0, the naming format used for the backup is `SC_HASHCODEofDISKGROUP_DBSID_BACKUPID`. However, If the disk group name is less than 16 characters, the naming format used for the backup is `DISKGROUPNAME_DBSID_BACKUPID`, which is the same format used in SnapCenter 2.0.

The `HASHCODEofDISKGROUP` is an automatically generated number (2 to 10 digit) unique for each ASM disk group.

Crosscheck operations

You can perform crosschecks to update outdated RMAN repository information about backups whose repository records do not match their physical status. For example, if a user removes archived logs from disk with an operating system command, the control file still indicates that the logs are on disk, when in fact they are not.

The crosscheck operation enables you to update the control file with the information. You can enable crosscheck by running the `Set-SmConfigSettings` command and assigning the value `TRUE` to the `ENABLE_CROSSCHECK` parameter. The default value is set to `FALSE`.

```
sccli Set-SmConfigSettings-ConfigSettingsTypePlugin-PluginCodeSCO-ConfigSettings  
"KEY=ENABLE_CROSSCHECK, VALUE=TRUE"
```

Remove catalog information

You can remove the catalog information by running the `Uncatalog-SmBackupWithOracleRMAN` command. You cannot remove the catalog information using the SnapCenter GUI. However, information of a cataloged backup is removed while deleting the backup or while deleting the retention and resource group associated with that cataloged backup.



When you force a deletion of the SnapCenter host, the information of the cataloged backups associated with that host are not removed. You must remove information of all the cataloged backups for that host before forcing the deletion of the host.

If the cataloging and uncataloging fails because the operation time exceeded the time out value specified for the `ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT` parameter, you should modify the value of the parameter by running the following command:

```
/opt/Netapp/snapcenter/spl/bin/sccli Set-SmConfigSettings-ConfigSettingsType  
Plugin -PluginCode SCO-ConfigSettings
```

```
"KEY=ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT,VALUE=user_defined_value"
```

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help command_name. Alternatively, you can refer to the [SnapCenter Software Command Reference Guide](#).

Predefined environment variables for backup specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while creating backup policies. This functionality is supported for all Oracle configurations except for VMDK.

SnapCenter predefines the values of the parameters that will be directly accessible in the environment where the shell scripts are executed. You do not have to manually specify the values of these parameters when executing the scripts.

Supported predefined environment variables for creating backup policy

- **SC_JOB_ID** specifies the job ID of the operation.

Example: 256

- **SC_ORACLE_SID** specifies the system identifier of the database.

If the operation involves multiple databases, the parameter will contain database names separated by pipe.

This parameter will be populated for application volumes.

Example: NFSB32|NFSB31

- **SC_HOST** specifies the host name of the database.

For RAC, host name will be the name of the host on which backup is performed.

This parameter will be populated for application volumes.

Example: scsmohost2.gdl.englobe.netapp.com

- **SC_OS_USER** specifies the operating system owner of the database.

The data will be formatted as <db1>@<osuser1>|<db2>@<osuser2>.

Example: NFSB31@oracle|NFSB32@oracle

- **SC_OS_GROUP** specifies the operating system group of the database.

The data will be formatted as <db1>@<osgroup1>|<db2>@<osgroup2>.

Example: NFSB31@install|NFSB32@oinstall

- **SC_BACKUP_TYPE** specifies the backup type (online full, online data, online log, offline shutdown, offline mount)

Examples:

- For full backup: ONLINEFULL
- data only backup: ONLINEDATA
- For log only backup: ONLINELOG

- **SC_BACKUP_NAME** specifies the name of the backup.

This parameter will be populated for application volumes.

Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1|AV@RG2_scspr2417819002_07-20-2021_12.16.48.9267

- **SC_BACKUP_ID** specifies the backup ID.

This parameter will be populated for application volumes.

Example: DATA@203|LOG@205|AV@207

- **SC_ORACLE_HOME** specifies the path of the Oracle home directory.

Example:

NFSB32@/ora01/app/oracle/product/18.1.0/db_1|NFSB31@/ora01/app/oracle/product/18.1.0/db_1

- **SC_BACKUP_RETENTION** specifies the retention period defined in the policy.

Examples:

- For full backup: Hourly|DATA@DAYS:3|LOG@COUNT:4
- For on-demand data only backup: Ondemand|DATA@COUNT:2
- For on-demand log only backup: Ondemand|LOG@COUNT:2

- **SC_RESOURCE_GROUP_NAME** specifies the name of the resource group.

Example: RG1

- **SC_BACKUP_POLICY_NAME** specifies the name of the backup policy.

Example: backup_policy

- **SC_AV_NAME** specifies the names of the application volumes.

Example: AV1|AV2

- **SC_PRIMARY_DATA_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for data files directory. It will be the name of the parent volume for luns and qtrees.

The data will be formatted as <db1>@<SVM1:volume1>|<db2>@<SVM2:volume2>.

Examples:

- For 2 databases in the same resource group:
NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA
- For single database with data files spread across multiple volumes:
buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA,herculus:/vol/scspr2417819002_NFS
- **SC_PRIMARY_ARCHIVELOGS_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for logs file directory. It will be the name of the parent volume for luns and qtrees.

Examples:

- For single database instance: buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO
- For multiple database instances:
NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO|NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO
- **SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG** specifies the list of Snapshots containing storage system name and volume name.

Examples:

- For single database instance:
buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For multiple database instances:
NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- **SC_PRIMARY_SNAPSHOT_NAMES** specifies the names of the primary Snapshots created during the backup.

Examples:

- For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For consistency group Snapshots that involves 2 volumes: cg3_R80404CBEF5V1_04-05-2021_03.08.03.4945_0_bfc279cc-28ad-465c-9d60-5487ac17b25d_2021_4_5_3_8_58_350
- **SC_PRIMARY_MOUNT_POINTS** specifies the mount point details which are part of the backup.

The details include the directory on which volumes are mounted and not the immediate parent of the file under backup. For an ASM configuration, it is the name of the disk group.

The data will be formatted as <db1>@<mountpoint1,mountpoint2>|<db2>@<mountpoint1,mountpoint2>.

Examples:

- For single database instance: /mnt/nfsdb3_data,/mnt/nfsdb3_log,/mnt/nfsdb3_data1
- For multiple database instances:
NFSB31@/mnt/nfsdb31_data,/mnt/nfsdb31_log,/mnt/nfsdb31_data1|NFSB32@/mnt/nfsdb32_data,/mnt/nfsdb32_log,/mnt/nfsdb32_data1
- For ASM: +DATA2DG,+LOG2DG

- **SC_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS** specifies the names of the snapshots created during the backup of each of the mount points.

Examples:

- For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data, RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log
- For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data, RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb31_data, RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb32_log

- **SC_ARCHIVELOGS_LOCATIONS** specifies the location of the archive logs directory.

The directory names will be the immediate parent of the archive log files. If the archive logs are placed in more than one location then all the locations will be captured. This also includes the FRA scenarios. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_log
- For multiple databases on NFS and for the NFSB31 database archive logs that are placed in two different locations: NFSB31@/mnt/nfsdb31_log1,/mnt/nfsdb31_log2|NFSB32@/mnt/nfsdb32_log
- For ASM: +LOG2DG/ASMDB2/ARCHIVELOG/2021_07_15

- **SC_REDO_LOGS_LOCATIONS** specifies the location of the redo logs directory.

The directory names will be the immediate parent of the redo log files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_data/newdb1
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/newdb32
- For ASM: +LOG2DG/ASMDB2/ONLINELOG

- **SC_CONTROL_FILES_LOCATIONS** specifies the location of the control files directory.

The directory names will be the immediate parent of the control files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_data/fra/newdb1,/mnt/nfsdb2_data/newdb1
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data/fra/newdb31,/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/fra/newdb32,/mnt/nfsdb32_data/newdb32
- For ASM: +LOG2DG/ASMDB2/CONTROLFILE
- **SC_DATA_FILES_LOCATIONS** specifies the location of the data files directory.

The directory names will be the immediate parent of the data files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb3_data1,/mnt/nfsdb3_data/NEWDB3/datafile
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data1,/mnt/nfsdb31_data/NEWDB31/datafile|NFSB32@/mnt/nfsdb32_data1,/mnt/nfsdb32_data/NEWDB32/datafile
- For ASM: +DATA2DG/ASMDB2/DATAFILE,+DATA2DG/ASMDB2/TEMPFILE
- **SC_SNAPSHOT_LABEL** specifies the name of the secondary labels.

Examples: Hourly, Daily, Weekly, Monthly, or custom label.

Supported delimiters

- **:** is used to separate SVM name and volume name

Example: buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

- **@** is used to separate data from its database name and to separate the value from its key.

Examples:

- NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- NFSB31@oracle|NFSB32@oracle
- **|** is used to separate the data between two different databases and to separate the data between two different entities for SC_BACKUP_ID, SC_BACKUP_RETENTION, and SC_BACKUP_NAME parameters.

Examples:

- DATA@203|LOG@205
- Hourly|DATA@DAYS:3|LOG@COUNT:4

- DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1

- / is used to separate the volume name from its Snapshot for SC_PRIMARY_SNAPSHOT_NAMES and SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG parameters.

Example: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

- , is used to separate set of variables for the same DB.

Example: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

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

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.



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

Backup schedules

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the

availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings > Policies**.

- Backup schedules

Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources > Resource Groups**.

Backup naming conventions

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

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.