



Technical Report

Best Practices Guide for Microsoft SQL Server using NetApp SnapCenter

Manohar Kulkarni, NetApp
February 2021 | TR-4714

Abstract

This best practice guide is intended to assist storage and application administrators successfully deploy Microsoft SQL Server 2008, R2, 2012, 2014, 2016, 2017, and 2019 on NetApp® storage using NetApp SnapCenter® technology for data protection.

TABLE OF CONTENTS

Executive summary	4
Purpose and scope	4
Audience	4
SnapCenter	4
Overview	4
SnapCenter architecture	4
SnapCenter features	6
SnapCenter Server requirements	6
Database storage layout	7
Considerations for setting Microsoft SQL Server database storage layout while backing up with SnapCenter	7
Considerations for database layout in a virtual environment	8
Upgrading SnapManager to SnapCenter Plug-in for Microsoft SQL Server	10
Upgrade methodology	10
General SnapCenter Server recommendations	11
Backup best practices	11
Using a resource group to back up multiple databases	12
Setting up a log backup	13
Enabling verification post database backup	13
Sizing considerations for SnapVault and SnapMirror	14
Considerations for an instance with a large number of small-to-large size databases	15
Restore best practices	16
ONTAP restore mechanism	16
Restore scenarios based on the recommended storage layout for SnapCenter	17
Restoring a database by restoring to an alternate host option	20
Clone best practices	21
Using a clone to create a database copy	21
Managing the clone copy	22
Managing clone copies in hybrid scenarios	22
Additional cloning guidelines	22
Microsoft SQL Server deployment for advanced setup	23
Always On Availability Group	23
Manage Microsoft SQL Server and SnapCenter backups for Asymmetric LUN Mapping in Windows clusters	24

Performance benchmarking	25
Lab setup 1: Using NetApp storage FAS 8020	25
Conclusion	26
Where to find additional information	26

LIST OF TABLES

Table 1) SnapCenter Server requirements.	6
Table 2) SnapCenter restore mechanism based on data file layout.	20
Table 3) SnapCenter Server configuration.	25
Table 4) Plug-in configuration.....	25
Table 5) Storage layout and number of SnapCenter backup groups.	25
Table 6) Workflow timings.	25

LIST OF FIGURES

Figure 1) SnapCenter Server architecture.....	5
Figure 2) Database layout for VMware using VMDK with NFS datastores.	10
Figure 3) Instance hosting multiple databases on the same volume.	13
Figure 4) Data and log files residing on separate LUNs per volume.	18
Figure 5) Multiple data files residing on same the LUN.	18
Figure 6) Data files residing on separate virtual disks (VMDK over VMFS).	19
Figure 7) Multiple data files residing on the same virtual disk (VMDK over VMFS).	19
Figure 8) SnapCenter managing data protection of resources across hybrid cloud setup.	22
Figure 9) SnapCenter backup policy settings.	23
Figure 10) SnapCenter managing SQL Server Availability Group between two failover cluster instance.	24

Executive summary

As data size and the number of databases is increasing, maintaining recovery time objective [RTO] and recovery point objective [RPO] and managing the backup operation has become very critical to enterprise applications. The requirements of a data protection product include providing efficient manageability for standalone and SQL Always On instances and on-demand quick availability of database copy while ensuring the storage space is not occupied with redundant copies.

By leveraging NetApp ONTAP® technologies to achieve RTO and RPO and lesser time to back up databases, SnapCenter allows you to modernize data protection.

Purpose and scope

Designing an optimal data protection and restore process is a vital strategy. An improper design can lead to longer backup and restore times, sprawling backup jobs, complicated and time-consuming restore operations, unnecessary consumption of redundant spaces, and many more scenarios.

Tactically deploying database files, offloading some of the background process to alternate instances, and grouping backup resources are some of the topics discussed in this document.

The best practices and recommendations described in this guide enable database architects and storage administrators to plan a highly available and easy-to-manage Microsoft SQL Server environment and meet stringent SLAs.

Audience

This guide assumes that the reader understands Microsoft SQL Server storage architecture and administration and data protection concepts of backup and restore. This guide also assumes the reader has a working knowledge of the following topics:

- NetApp ONTAP data management software
- NetApp SnapCenter software
- Microsoft SQL Server

To determine the configuration compatibility across the NetApp stack, see the [NetApp Interoperability Matrix Tool \(IMT\)](#).

SnapCenter

Overview

SnapCenter is the NetApp next-generation data protection software for tier 1 enterprise applications. SnapCenter, with its single pane of glass management interface, automates and simplifies the manual, complex, and time-consuming processes associated with the backup, recovery, and cloning of multiple databases and other application workloads.

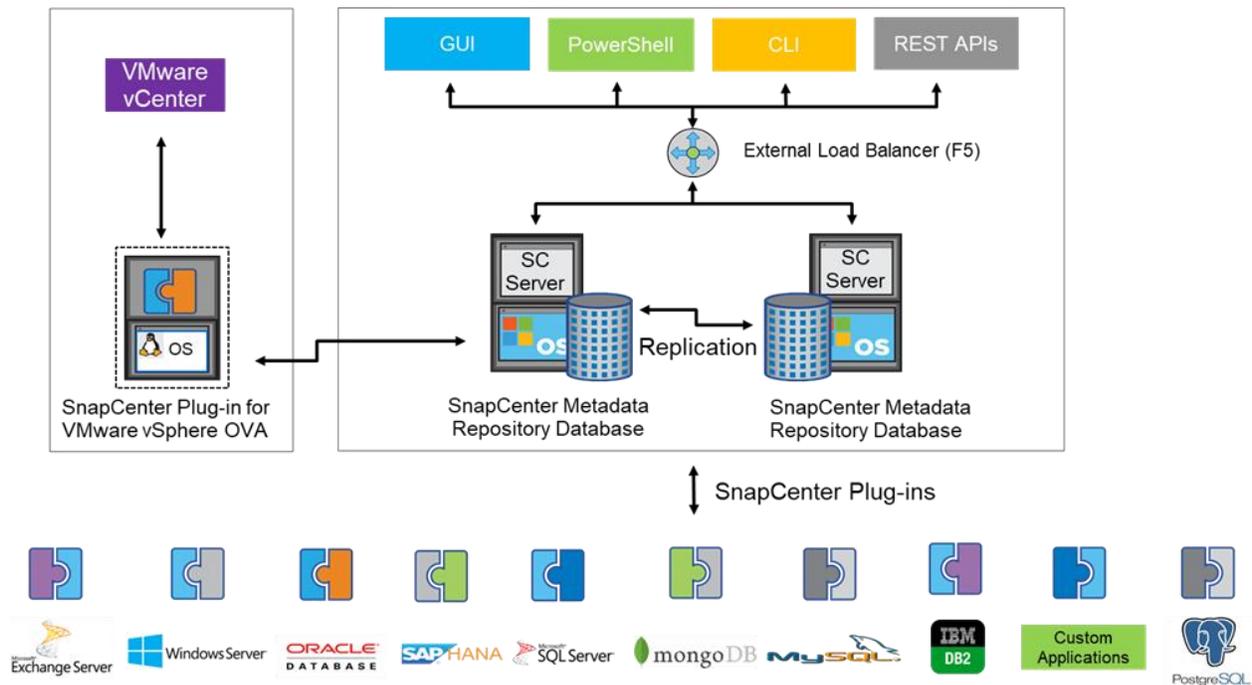
SnapCenter leverages NetApp technologies, including NetApp Snapshot™, NetApp SnapMirror®, SnapRestore®, and NetApp FlexClone®, which allows it to integrate seamlessly with technologies offered by Oracle, Microsoft, SAP, and VMware, and across FC, iSCSI, and NAS protocols. This integration allows IT organizations to scale their storage infrastructure, meet increasingly stringent SLA commitments, and improve the productivity of administrators across the enterprise.

SnapCenter architecture

SnapCenter is a centrally managed web-based application that runs on a Windows platform and remotely manages multiple servers that need to be protected.

Figure 1 is a high-level architecture of SnapCenter Server.

Figure 1) SnapCenter Server architecture.



SnapCenter Server has an HTML5-based GUI as well as Microsoft Windows PowerShell cmdlets and APIs. High availability can be set up leveraging external load balancer such as the F5 load balancer. High availability should be set up within the same data center. If one SnapCenter host is ever unavailable for any reason, then the second SnapCenter Server can seamlessly take over with minimal impact on operation.

SnapCenter Server can push out plug-ins to remote hosts. These plug-ins are used to interact with an application, database, or file system. For Windows, SQL Server, and Exchange, plug-ins are present on the remote hosts so that application- or database-level commands can be issued from the same host where the application or database is running.

To manage plug-ins and interaction between SnapCenter Server and the plug-in host, SnapCenter uses SM, service which is a SnapManager web service running on top of Windows Server Internet Information Services (IIS) on SnapCenter Server that takes all client requests such as backup, restore, clone and so on.

SnapCenter Server communicates those requests to SMCore, which is a service that runs colocated within SnapCenter Server and remote servers and plays a significant role in coordinating with the SnapCenter Plug-ins Package for Windows (which includes the SnapCenter Plug-in for Microsoft Windows Server and SnapCenter Plug-in for Microsoft SQL Server, to discover the host file system, gather database metadata, quiesce/thaw and lastly in managing SQL Server database during backup, restore, clone and verification).

SnapCenter Plug-in for VMware vSphere (SCV) is another SnapCenter virtualization plug-in that manages virtual servers running on VMWare and helps in discovering host file system and databases on VMDKs and RDMs. SCV is a separate installation with an Open Virtual Appliance (OVA) based setup on the Linux-based Debian OS. The SCV details must be registered in SnapCenter Server in order to discover VMWare virtual resources.

SnapCenter features

SnapCenter enables you to create application-consistent Snapshot copies and to complete data protection operations, including Snapshot copy-based backup, clone, restore, and backup verification operations. SnapCenter provides a centralized management environment, while using role-based access control (RBAC) to delegate data protection and management capabilities to individual application users across your SnapCenter Server and Windows hosts. SnapCenter includes the following key features:

- A unified and scalable platform across applications and database environments and virtual and nonvirtual storage, powered by SnapCenter Server
- Consistency of features and procedures across plug-ins and environments, supported by the SnapCenter user interface
- RBAC for security and centralized role delegation
- Application-consistent Snapshot copy management, restore, clone, and backup verification support from both primary and secondary destinations (NetApp SnapMirror and SnapVault®)
- Remote package installation from the SnapCenter GUI
- Nondisruptive, remote upgrades
- A dedicated SnapCenter repository that provides faster data retrieval
- High availability of SnapCenter server using external load balancer such as F5 load balancer
- Centralized scheduling and policy management to support backup and clone operations
- Centralized reporting, monitoring, and dashboard views

SnapCenter Plug-in for SQL Server features include:

- Automated discovery
- Wizard-based GUI for SQL Server data migration from third-party storage to NetApp SAN
- Support for full and log backups with granular restore options
- Offload verification to a remote SQL Server instance, even on SnapMirror destination
- Ability to restore to an alternate location
- Faster reseed option for SQL Server Always On
- Ability to clone Lifecycle Management
- Ability to clone a backup to any host
- Ability to refresh clones with production data on a schedule
- Ability to clone a clone
- Migration of long-term SnapVault copies from NetApp SnapManager for Microsoft SQL Server (SMSQL)

SnapCenter Server requirements

Table 1 lists the minimum requirements for installing SnapCenter Server and plug-in on Microsoft Windows Server.

Table 1) SnapCenter Server requirements.

Components	Requirement
Minimum CPU count	Four cores/vCPUs
Memory	Minimum: <ul style="list-style-type: none">• 8GB Recommended: <ul style="list-style-type: none">• 32GB

Components	Requirement
Storage space	Minimum space for installation: <ul style="list-style-type: none"> • 10GB Minimum space for repository: <ul style="list-style-type: none"> • 10GB
Supported operating system	<ul style="list-style-type: none"> • Windows Server 2012 • Windows Server 2012 R2 • Windows Server 2016 • Windows Server 2019
Software packages	<ul style="list-style-type: none"> • .NET 4.5.2 or later • Windows Management Framework (WMF) 4.0 or later • PowerShell 4.0 or later

For version compatibility for latest and other plug-ins please, see the [NetApp Interoperability Matrix Tool](#).

Database storage layout

Considerations for setting Microsoft SQL Server database storage layout while backing up with SnapCenter

For prescriptive best practices for the design consideration and for information about deploying Microsoft SQL Server on NetApp storage system, see [TR-4590: Best Practice Guide for Microsoft SQL Server with ONTAP](#).

SnapCenter best practice considerations for Microsoft SQL Server database layout are aligned with the recommended Microsoft SQL Server deployment best practices described in TR-4590. SnapCenter supports the backup of user and system databases residing on a NetApp storage system.

In addition to the performance benefit of segregating the user database layout into different volumes, the database also has high impact with the time required to back up and restore. Having separate volumes for data and log files significantly improves the restore time as compared to a volume hosting multiple user data files. Similarly, user databases with a high I/O intensive application are prone to an increase in the backup time. A more detailed explanation about backup and restore practices is provided later in this document.

Consider the following storage layout data points for backing up databases with SnapCenter:

- Isolate databases with I/O-intensive queries throughout the day in a different volume and should eventually have separate jobs to back it up.
- Place databases with a large DB size or databases having minimal RTO in separate volumes for faster recovery.
- Consolidate small-to-medium size databases that are less critical or having fewer I/O requirements to a single volume. Backing up such a large number of databases residing in the same volume leads to fewer Snapshot copies that need to be maintained. NetApp also suggests consolidating Microsoft SQL Server instances to use the same volumes to control the number of backup Snapshot copies taken.
- Create separate LUNs to store full text-related files and file streaming related files.
- Assign separate LUNS per host to store Microsoft SQL Server log backups.
- System databases store database server metadata; configuration and job details are not updated frequently. Place system databases/tempdb in separate drives/LUNs. Do not place system databases in the same volume as the user databases. User databases have a different backup policy and the frequency of the user database backup is not same for system databases.

- For Microsoft SQL Server Availability Group setup, place the data/log files for replicas in an identical folder structure on all nodes.

The following table lists additional NetApp recommendations on volume design.

Best Practices

- NetApp recommends having at least 10% free space available in an aggregate for optimal storage performance.
- Use flexible volumes to store Microsoft SQL Server database files and avoid sharing volumes between hosts.
- Use NTFS mount points instead of drive letters to surpass the 26-drive letter limitation in Microsoft Windows Server. When using volume mount points, a general recommendation is to give the volume label the same name as the mount point.
- Configure a volume auto size policy, when appropriate, to help prevent out-of-space conditions.
- On FAS systems, enable read reallocation on the volume when the Microsoft SQL Server database I/O profile consists of mostly large sequential reads, such as with decision support system workloads. Read reallocation optimizes the blocks to provide better performance.
- Set the Snapshot copy reserve value in the volume to zero for ease of monitoring from an operational perspective.
- Disable storage Snapshot copy schedules and retention policies. Instead, use SnapCenter Plug-in for Microsoft SQL Server (SCSQL) to manage Snapshot copies of the Microsoft SQL Server data volumes.
- Tempdb is a system database used by Microsoft SQL Server as a temporary workspace, especially for I/O intensive DBCC CHECKDB operations. Therefore, place this database on a dedicated volume with a separate set of spindles. In large environments in which volume count is a challenge, you can consolidate tempdb into fewer volumes and store it in the same volume as other system databases after careful planning. Data protection for tempdb is not a high priority because this database is recreated every time Microsoft SQL Server is restarted.
- Place user data files (.mdf) on separate volumes because they are random read/write workloads. It is common to create transaction log backups more frequently than database backups. For this reason, place transaction log files (.ldf) on a separate volume or VMDK from the data files so that independent backup schedules can be created for each. This separation also isolates the sequential write I/O of the log files from the random read/write I/O of data files and significantly improves Microsoft SQL Server performance.

Considerations for database layout in a virtual environment

NetApp recommends a similar approach to the recommended notes in the previous section on storage layout. There are additional points to consider when data files reside on VMDK or RDM; SnapCenter only supports database files on a VMware virtual disk.

Note: For other hypervisor environments such as Hyper-V, SnapCenter is supported for databases files residing on NetApp storage connected with iSCSI- or FC-based protocols.

Virtual disk on VMFS

VMware VMs typically include a set of files in two given formats: virtual machine file system (VMFS) or raw device mapping (RDM). Both formats enable you to access the virtual machine's disk (VMDK), but they differ in approach to storage, and VMware recommends VMFS for the majority of VMs. With VMFS, the VMDK files also hold the data, while with RDM, the data is stored on an external disk system. VMFS holds disk data from multiple VMs; RDM does not. VMFS was designed specifically to support virtualization. Although RDM is sometimes recommended for I/O-intensive operations, with VMFS, a storage volume can support one or many VMs. This volume can change without affecting network operations. Because they share storage volumes, VMs are easier to manage, and resource utilization

remains high. Various ESXi servers can read and write to the file system at once, because it stores information at the block level.

The following table lists the NetApp recommendations for VMDK:

Best Practices

- Use separate VMDKs for primary (.mdf) and log (.ldf) files for user databases. Make sure that these VMDKs reside in a datastore placed on a separate volume from the volume containing system databases and the operating system VMDKs.
- Use separate VMDKs for system databases (master, model, and msdb). Make sure that these VMDKs reside in a datastore placed on a separate volume from the volume containing user databases and the operating system VMDKs.
- Use separate VMDKs for the tempdb database.
- Data files (tables and indexes) are the primary files that are used by the Microsoft SQL Server storage engine. Each database might have multiple files and should be spread across multiple VMDKs.
- Avoid sharing volumes and datastores between different Microsoft Windows Server machines.

Note: Whether it's a single database file or multiple database files on VMDK (over VMFS), the database restore mechanism is a mount and copy restore.

Virtual disk on NFS

ESXi can access a designated NFS volume located on a NAS server, mount the volume, and use it for its storage needs. You can use NFS volumes to store and boot VMs in the same way that you use VMFS datastores.

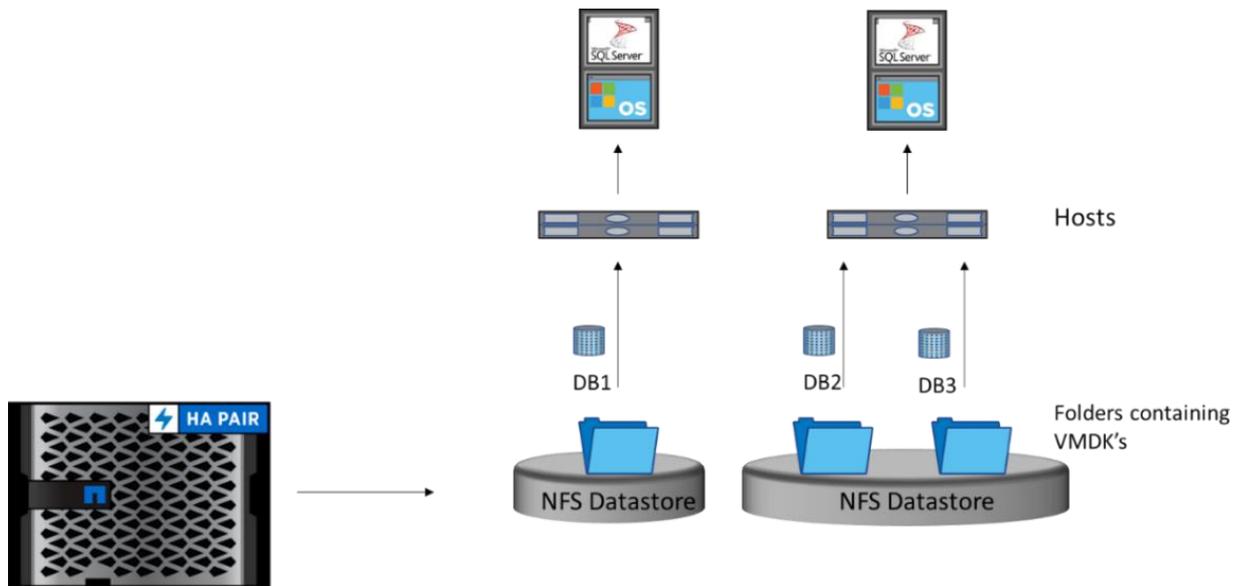
ESX 5.0 and later support up to 256 NFS datastores. The default value is eight, but this can be increased to the maximum number that is specific to the version of ESX or ESXi being used. The following table lists the NetApp recommendations for NFS datastores:

Best Practices

- Use one NFS datastore for multiple system databases from multiple instances.
- Use one NFS datastore per user database and user log; alternatively, separate the user database and the user log on different NFS datastores.
- Do not define a default gateway for the NFS storage network.
- Make sure that each NFS datastore is connected only once from each ESX or ESXi server using the same NetApp target IP address on each ESX or ESXi server.
- From VMware vCenter 6.7, SnapCenter supports 256 paravirtual SCSI disks only for RDM.

Figure 2 illustrates the database layout for VMware using VMDK with NFS datastores.

Figure 2) Database layout for VMware using VMDK with NFS datastores.



The database storage design has the following characteristics:

- If the user wants to restore a single database in the VMDK that is also holding single database, SnapCenter performs Single File SnapRestore to recover the data file and log. For example, recovering DB1 from the layout in Figure 2 will be a faster process. The Single File SnapRestore description is explained in the [ONTAP Restore Mechanism](#) section of this document.
- When placing multiple database files and logs in a single or separate VMDKs, restoring a single database includes a mount and copy restore operation, as in the case of DB2 and DB3 (shown in Figure 2).

Note: If there are no RDMs or VMDKs (VMFS) in the environment, NetApp recommends using hypervisor settings with a NFS or iSCSI protocol.

Upgrading SnapManager to SnapCenter Plug-in for Microsoft SQL Server

Upgrade methodology

The SnapManager product is at the end of lifecycle support; no new features or patches will be released.

NetApp recommends upgrading SnapManager to SnapCenter 4.4, if not already upgraded.

Use the following recommended process to migrate SnapManager to SnapCenter:

- Install SCSQL Plug-in with NetApp SMSQL.
- Get the current backup schedule information of SMSQL from PowerShell cmdlet Get-SmSchedule
- Create a new policy-based scheduled backup in SCSQL.
- If there are long-term Snapshot copies stored in the SnapVault destination, migrate those Snapshot copy details to SnapCenter by using the import command utility.

For more information about procedures and limitations, see [“Importing archived backups from SnapManager for Microsoft SQL Server to SnapCenter.”](#)

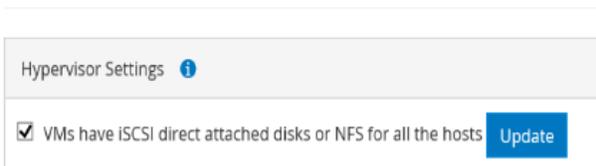
- Disable SMSQL schedules using cmdlet `Get-SmSchedule -Hostname <hostname> -PluginCode SMSQL -DisableCurrentGenSchedule`.
- Continue to run backup jobs from SnapCenter
- SMSQL and SCSQL can coexist but run the backup jobs only from SCSQL.
- After the backups taken through SMSQL have aged out, delete the Snapshot copies and uninstall SMSQL.

General SnapCenter Server recommendations

Consider the following recommendations for SnapCenter Server:

- Make sure the following global setting is selected:

Global Settings



- Enable hypervisor setting > Global Settings for iSCSI disk connected to VMs.
- When adding the Windows Failover Cluster host in SnapCenter, make the sure Add All Hosts in the Cluster or DAG option is selected.

Add all hosts in the cluster or DAG i

- Make sure that the SVM short name and ONTAP cluster name are resolvable from all the plug-in hosts and SnapCenter Server. Resolve the SVM short name by adding a DNS entry or etc\hosts entry.
- NetApp recommends uninstalling SnapCenter plug-ins from the SnapCenter GUI only.
- Databases deleted from SQL Server management studio but backups taken via SnapCenter are not cleaned up and remain on ONTAP. Snapshot copies can be manually cleaned by deleting them from SnapCenter or you can run following command to configure retention period of such deleted database. The setting is done at a global level and applied across all SQL Server hosts.

```
Set-SmConfigSettings -Server -configSettings @{"DeletedDatabaseRetentionDays"="2"}
```

By default, the value is set to 0, which means the Snapshot copies are not cleared.

- By default, the retention of log Snapshot copies on volume is set to seven days. The up-to-the-minute restore (UTM) retention set from the SnapCenter GUI is applicable only on the file system. To modify this settings on volume, use PowerShell cmdlet `Set-SmPolicy`.

For example:

```
Set-SmPolicy -PolicyName 'DBFullLog' -PolicyType 'Backup' -Description 'Full and log backup Policy' -scheduleType Hourly -retentionsettings @{"BackupType"="DATA";"RetentionCount"="3"}, @{"BackupType"="DATA";"ScheduleType"="HOURLY";"RetentionCount"="3"},@{"BackupType"="LOG";"RetentionCount"="3"},@{"BackupType"="LOG_SNAPSHOT";"RetentionCount"="3"} -pluginpolicytype 'SCSQL' -sqlbackuptype 'Fullbackupandlogbackup' -CreateLogFolderSnapshot
```

Backup best practices

This section provides various scenarios to configure backup policies to ensure that you recover your database with ease, with minimal downtime, and you do not encounter last-minute surprises.

The following guidelines should be considered before configuring a backup policy:

- Do you need a full backup of the database and a log backup of the database?
 - List the RPOs for your production and nonproduction systems:
 - In SnapCenter terms, RPO can be identified as the backup frequency; for example, how frequently you want to schedule the backup so that you can reduce the loss of data to up to few minutes. SnapCenter doesn't stop you from scheduling backup as minimum as five minutes; however, there might be few instances where a backup might not complete within five minutes during peak transactions or when the rate of change of data is more in the given time. A NetApp best practice is to schedule frequent transaction log backups instead of the full backups.
 - NetApp recommends that you schedule Snapshot copies only from one source, such as from SnapCenter or directly from ONTAP storage, to avoid reaching the Snapshot copy limit of 255 per volume or unwanted situations of breaching Snapshot copy retention on ONTAP 9.3 or earlier. NetApp recommends that you back up from SnapCenter because it provides an application-consistent backup.
- Note:** For ONTAP 9.4 and later, the Snapshot copy limit is extended to 1,024.
- There are numerous approaches to handle the RPOs and RTOs. One alternative to this backup approach is to have separate backup policies of data and log with different intervals. For example, from SnapCenter, schedule log backups in 15-minute intervals and data backups in intervals of 24-hour intervals.
- List the retention requirements.
 - How long do you want to retain these backups based on the category of backups (hourly, daily, weekly, or monthly)? If you want to protect the Snapshot copies for long term, then you can create a SnapVault destination. You can tag these backups in various categories such as hourly, daily, weekly, and monthly and set a retention for the same. SnapVault might not be an exact replacement for tape or cloud, but it can address the backup needs for long term.
 - For SnapCenter 4.4 and later, a backup triggered on demand is retained according to the scheduled backup policy. If an on-demand backup must be maintained with a different retention period, then create a separate on-demand policy with a required retention period and trigger backup.
 - Do you want to verify and validate each backup for compliance or an auditing need? The verification can be enabled in the policy and activated during the implicit resource group creation. You can verify the backups from both the primary or secondary disaster recovery or vault storage. You can also perform deferred verification by scheduling the backup either from the GUI or CLI. The reason for this verification process is that we clone the database volumes, mount on the host, and run the Microsoft SQL Server DBCC CHECKDB process across the mounted FlexClone files.

Using a resource group to back up multiple databases

Resource groups were introduced with SnapCenter. The advantage of having resource groups is that you can group databases of an instance in one job and reduce the number of Snapshot copies created per volume, if multiple databases reside on same volume. For example, if a Microsoft SQL Server instance is hosting multiple databases, such as 100 databases or more, on the same large volumes, SnapCenter won't create separate Snapshot copies for each database. Instead, it optimizes the Snapshot copies with a single application-consistent Snapshot copy for that volume. It therefore reduces the number of Snapshot copies on ONTAP and a consistent point-in-time (PiT) backup is available.

Note: Snapshot copies are created separately per host per volume.

In summary, resource groups help to reduce the storage overhead and control the number of Snapshot copies reaching limits.

Figure 3 is the example of an instance hosting multiple databases on the same volume (SC01040006500_5_MDML_Data_Log_Vol). This screenshot depicts a single Snapshot copy being created on ONTAP each time the Snapshot copy is taken through a resource group.

Figure 3) Instance hosting multiple databases on the same volume.

The screenshot shows the SnapCenter interface. At the top, there is a search bar and a resource group name: "RG_SC01040006500_5_INST_LVLDetails" with a red arrow pointing to "RG". Below this is a table with columns "Resource Name", "Type", and "Host". The table lists nine resources, all of type "SC01040006500_5_MDML_X (SC01040006500-5\NAMEDINST1)". A blue arrow points from this table to a detailed view of a volume: "Volume: SC01040006500_5_MDML_Data_Log_Vol". This view shows a table with columns "Snapshot Name" and "RG_SC01040006500_5_INST_LVL_SC01040006500-5_04-04-2018_13.35.17.5994_1". Two snapshots are listed with their respective names and timestamps.

Best Practice

Use a resource group for a backup configuration for Snapshot optimization and the number of jobs to be managed.

Setting up a log backup

Transaction log backups provide PiT recovery, which helps to attain the required RPO.

SnapCenter provides the ability to perform and manage log backup files. This capability helps in to automatically recover the complete database from SnapCenter without any manual intervention during the process.

Create a host log directory (for SnapCenter) on the dedicated FlexVol volume in which SnapCenter copies transaction logs. Set the log directory path based on the following scenario:

- If Microsoft SQL Server instance is running as a standalone instance, configure the host log directory per host.
- For an availability group, make sure to configure the host log directory for each node as part of the cluster group.
- In the event of failover cluster instance (FCI), configure the host log directory per instance and make sure the drive selected for the log directory is part of cluster disk group under Microsoft SQL Server role.

Best Practice

NetApp recommends backing up the transaction log backup from SnapCenter so that during the restoration process, SnapCenter can read all the backup files and restore in sequence automatically.

Enabling verification post database backup

Verification of a database is configured while configuring a backup job. In a consolidated environment where there are multiple databases running on the same instance, the verification process can be offloaded to alternate servers.

Verification is enabled in the resource group workflow after scheduling the backup. Select the SQL Server instances to run the verification process. If there are multiple databases in an instance, databases are distributed across different SQL Server hosts to run the verification process in parallel.

In the following example, Microsoft SQL Server Instance has 20 databases with a small-size database.

Here are the use cases:

1. Considering a single instance for verification.

(Job 86461) Verification of Resource Group 'DemoBackup' with verification policy 'DemoBackup_Verification'	22:54:53	0.00:02:30	
--	----------	------------	--

The total time taken was 00:2:30. Because one job was created to sequentially execute the verification job, the process took more time.

2. Considering three instances including the primary instance for verification. Jobs were distributed across various servers, thus reducing the verification time.

SC01040006500-5.cit1.local	22:04:55	0.00:00:58	
SC01040006500-4.cit1.local	22:04:55	0.00:00:46	
SC0104000650010.cit1.local	22:04:56	0.00:00:56	
SC01040006500-4.cit1.local	22:04:56	0.00:00:45	
SC01040006500-4.cit1.local	22:04:56	0.00:00:45	
SC01040006500-5.cit1.local	22:04:56	0.00:00:31	
SC01040006500-5.cit1.local	22:04:56	0.00:00:43	
SC0104000650010.cit1.local	22:04:56	0.00:00:42	
SC0104000650010.cit1.local	22:04:55	0.00:00:47	

The total time taken was 00:1:04.

Best Practice

To reduce the overall verification time, add SQL Server instances to the verification process if there are multiple databases. Consider those instances that are less utilized, are low-performance impact, and support an equal or later version or edition compared to the primary instance.

Sizing considerations for SnapVault and SnapMirror

Correctly sizing a SnapVault or SnapMirror solution is an important step in the design phase to make sure that backups complete within a planned backup window, RPOs can be met, and user I/O performance is not adversely affected. There are many variables that must be considered when sizing a SnapVault or SnapMirror solution.

- **Disk space.** The SnapVault or SnapMirror secondary, or target, must be sized so that adequate disk space is available to retain all of the planned backups. This space can be estimated fairly accurately by using known information about the primary data and the required RPOs. This calculation is independent of the systems used. The data you need to calculate the required disk space on the

SnapVault or SnapMirror target is the size of the primary data and log backup data, the daily/weekly/monthly data change rates, the number of daily/weekly/monthly backup copies to be kept, and the space savings that can be anticipated by using NetApp deduplication and compression

- **Data throughput.** It is also important to size a SnapVault or SnapMirror solution so that data can be transferred fast enough to complete backups in the amount of time available. To do this, first determine how much data, on average, is transferred in a SnapVault or SnapMirror incremental update and how much time must be allotted to complete the backup, then use these numbers to determine the required data throughput in units such as megabytes per second. If you have multiple SnapVault or SnapMirror relationships, remember to consider the size of all the backups that must be completed during a given backup window. After the required data throughput speed is determined, an appropriate NetApp AFF or FAS system can be selected that can handle the data speed required.
- **Client I/O impact.** Another factor to consider when sizing is the impact of the SnapVault or SnapMirror processes on the other workloads running on a system. Clients generally experience more latency with a greater number of concurrent SnapVault or SnapMirror streams. If a large number of transfers must take place during a given backup window, it might be best to stagger the updates so they don't all run at the same time and cause unacceptable latency to clients. For example, if 80 relationships must update within a 10-hour window, and it is determined that, on average, eight concurrent updates can finish in one hour, then a transfer schedule can be created that starts eight new updates every hour during the given 10-hour window.

Considerations for an instance with a large number of small-to-large size databases

SnapCenter is capable of backing up a large number of sizeable databases in an instance or group of instances within a resource group. The size of a database is not the major factor in backup time. The duration of a backup can vary depending on number of LUNS per volume, the load on Microsoft SQL Server, the total number of databases per instance, and specially, the I/O bandwidth and usage.

While configuring the policy to backup databases from an instance or resource group, NetApp recommends that you restrict the maximum database backed up per Snapshot copy to 100 per host. Make sure the total number of Snapshot copies does not exceed the 250-copy limit for ONTAP 9.3 or the 1,024-copy limit for ONTAP 9.4 and later.

NetApp strongly recommends that you schedule a backup during an off-peak time so the backup process does not affect the applications with a minor wait time.

NetApp also recommends that you limit the backup jobs running in parallel by grouping the number of databases instead of creating multiple jobs for each database or instance. For optimal performance of the backup duration, reduce the number of backup jobs to a number that can back up around 1,000 or less databases at a time.

As previously mentioned, I/O usage is an important factor in the backup process. The backup process must wait to quiesce until all the I/O operations on a database are complete. Databases with highly intensive I/O operations should be deferred to other backup time or should be isolated from other backup jobs to avoid affecting other resources within same resource group that are to be backed up.

For an environment that has six Microsoft SQL Server hosts, hosting 200 databases per instance, assuming four LUNs per host, and one LUN per volume are created, set the full backup policy with the maximum databases backed up per Snapshot copy to 100. Two hundred databases on each instance are laid out as 200 data files distributed equally on two LUNs, and 200 log files are distributed equally on two LUNs, which is 100 files per LUN and per volume.

Schedule three backup jobs by creating three resource groups, each grouping two instances that include a total of 400 databases.

Running all three backup jobs in parallel will back up 1,200 databases simultaneously. Depending on the load on the server and I/O usage, the start and end time on each instance can vary. In this instance, a total of 24 Snapshot copies are created.

In addition to the full backup, NetApp recommends that you configure a transaction log backup for critical databases. Make sure that the database property is set to full recovery model.

Best Practices

- Do not include the tempdb database in a backup because the data it contains is temporary. Place tempdb on a LUN or an SMB share that is in a storage system volume in which Snapshot copies will not be created.
- Always run SCSQL backups during off-peak hours to avoid performance impact on regular user activity.
- A Microsoft SQL Server instance with a high I/O intensive application should be isolated in a different backup job to reduce the overall backup time for other resources.
- Limit the set of databases to be simultaneously backed up to approximately 1,000 and stagger the remaining set of database backups to avoid a simultaneous process.
- Use the Microsoft SQL Server instance name in the resource group instead of multiple databases because whenever new databases are created in Microsoft SQL Server instance, SnapCenter automatically considers a new database for backup.
- If you change the database configuration, such as changing the database recovery model to full recovery model, perform a backup immediately to allow up-to-the-minute restore operations.
- SCSQL cannot restore transaction log backups created outside of SnapCenter.
- When cloning FlexVol volumes, make sure that you have sufficient space for the clone metadata.
- Create separate policy to manage and backup system databases at least once a week.
- The verification server must be connected to the NetApp storage system where the backup Snapshot copies are located. The connectivity can be FC, FCoE, or iSCSI. The connectivity of the verification server to the NetApp storage does not have to match the connectivity of the production Microsoft SQL Server host.
- Make sure that SCSQL backups on VMDK use a verification server that is running on a VM only.

Note: SnapCenter does not allow a VM and Microsoft SQL Server application to be backed up together. Both are backed up separately.

Restore best practices

Before detailing the SnapCenter restore best practices, it's important to understand the restore mechanism from ONTAP.

ONTAP restore mechanism

ONTAP provides various restore mechanisms that are leveraged by the SnapCenter.

Single File SnapRestore

Single File SnapRestore is the faster mode of restore where the process restores a complete LUN in a volume from a Snapshot copy to the read-write mode. A LUN is treated as a single file in a volume. During a single file restore process, SnapCenter replaces the current working copy of the LUN with a Snapshot copy.

In summary, Single File SnapRestore is a realignment of pointers to the blocks to the time when the Snapshot copies were taken.

Note: The Single File SnapRestore process works if the data files are residing on a disk connected with an iSCSI or FC protocol, physical RDM, or VMDKs on NFS only.

SubLun

The SubLun process restores a range of bytes in the underlying LUNs associated with the files being restored. For example, a Microsoft SQL Server instance might be hosting multiple user databases over a same LUN. Restoring a user database file from the LUN without touching other files is a partial restoration which in turn must restore specific range of bytes from the LUN.

Comparatively speaking, a SubLun restore is more time consuming than the Single File SnapRestore process but still faster than regular copy-restore process.

Note: The SubLun process works if the data files are residing on a disk connected with an iSCSI or FC protocol, physical RDM, or if compression is not enabled on the volume.

Mount and copy restore

Alternate to the restore option available in ONTAP; the mount and copy restore mechanism is the traditional method for restoration. If any of the above process do not viable then SnapCenter selects the traditional approach to mount the volume and copy the files to target host and attach the databases. This is obviously time-consuming process. In case of log file restoration, SnapCenter may choose to access the shared path if provided by end user to restore the logs thus eliminating the process to copy the log files.

During the mount and copy restore process, temporary storage will be visible for brief time until the required files are copied and then will be removed.

Depending upon the restoration scenario and the architecture deployed, SnapCenter picks up the restoration mechanism.

The following section describes how the restore mechanism is considered based on the storage layout.

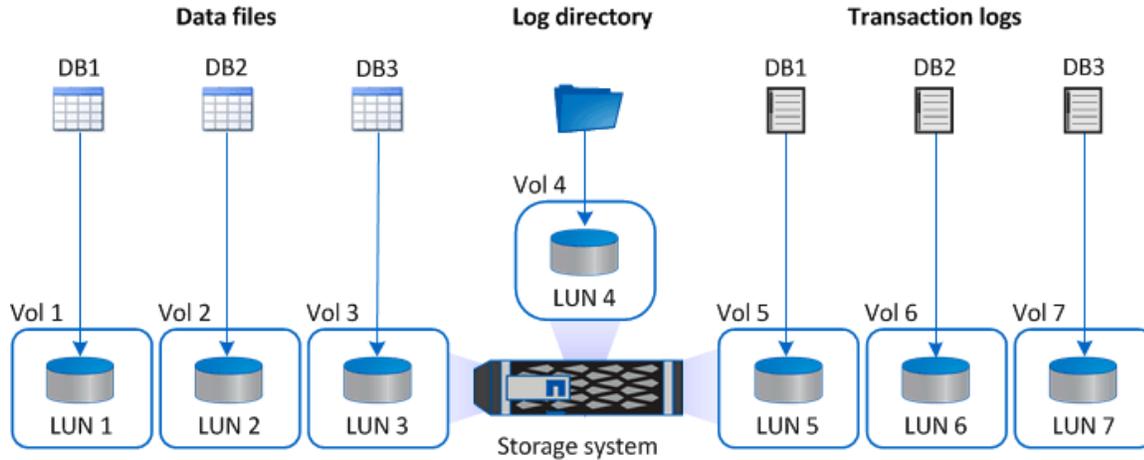
Restore scenarios based on the recommended storage layout for SnapCenter

With reference to the suggested storage layout in [SnapCenter 4.4 – Data Protection Guide for Microsoft SQL Server](#), this section describes the restoration path for each scenario.

Case 1: Storage layout for large databases on LUNs

In this case, each data file of a database is stored in a separate iSCSI/FC LUN. Therefore, when a database is restored, SnapCenter identifies whether all the data files residing on the LUN are associated with a single database to be restored and then uses the Single File SnapRestore mechanism to restore respective LUNs of the database from the Snapshot copy. Figure 4 illustrates the storage layout for large databases on LUNs.

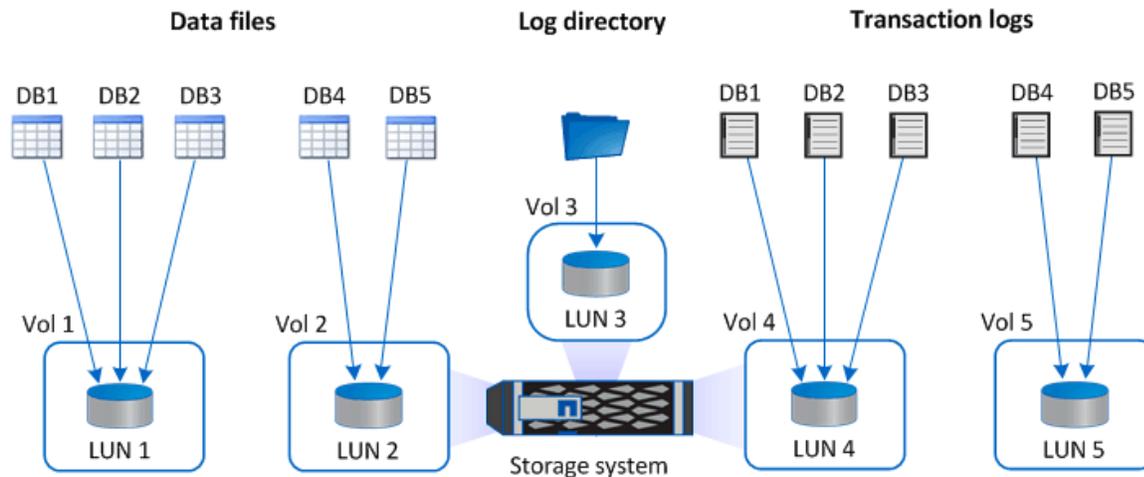
Figure 4) Data and log files residing on separate LUNs per volume.



Case 2: Storage layout for small to medium size databases on LUNs

Multiple databases are stored in an iSCSI/FC LUN, as shown in Figure 5. In this case, when a single database is to be restored, SnapCenter identifies the range of bytes used within LUN for that database. SnapCenter uses the SubLun mechanism to restore range of bytes to perform restore of a database.

Figure 5) Multiple data files residing on same the LUN.

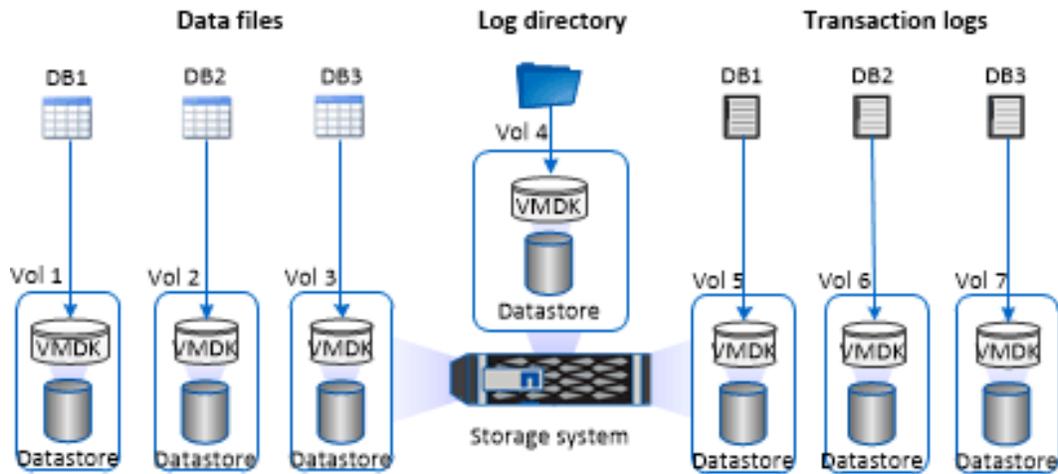


Case 3: Storage layout for large databases on VMDKs

In this case, database files reside on VMware VMDK files of VMFS. SnapCenter cannot recover VMFS as it does for LUN as a file, as shown in Figure 6.

SnapCenter uses mount and copy restore method of mounting the copy of volume as VMDK and then copy the files to desired location to perform database restoration by attaching database.

Figure 6) Data files residing on separate virtual disks (VMDK over VMFS).



Case 4: Storage layout for small-to-medium databases on VMDKs

SnapCenter cannot extract a range of bytes used for Microsoft SQL Server data files residing on VMDK of VMFS.

Similar to case 3, SnapCenter uses the mount and copy restore method to mount the copy of volume as VMDK and then copy the files to desired location to perform database restoration by attaching database, as shown in Figure 7.

Figure 7) Multiple data files residing on the same virtual disk (VMDK over VMFS).

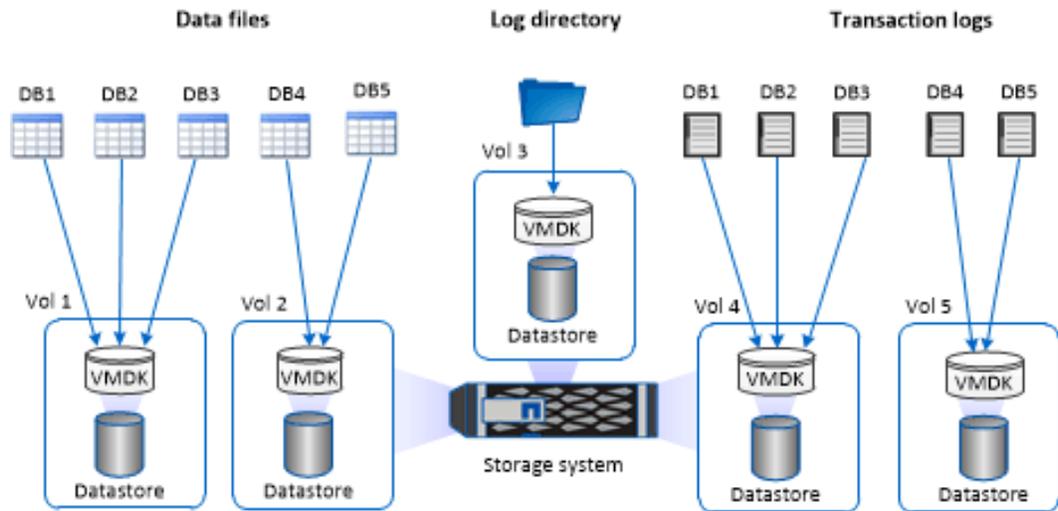


Table 2 shows the restore mechanism used by SnapCenter for various use cases.

Table 2) SnapCenter restore mechanism based on data file layout.

Serial Number	Scenario	In-Place Restore Mechanism	Restore Alternate Location
1	Single database spread across single or multiple LUNs/volume (no other data files on volume)	Single File SnapRestore	Mount and copy restore
2	Restoring single database from single LUN hosting multiple data files	SubLun (single database restore)	Mount and copy restore
3	Single database restore in VMDK on VMFS; single files or multiple files present on disk	Mount and copy restore	Mount and copy restore
4	Single database restore having files in multiple VMDKs (VMFS)	Mount and copy restore	Mount and copy restore
5	Single database restore from VMDK on NFS; no other files present	Single File SnapRestore	Mount and copy restore
6	Single database restore having files in multiple VMDKs (NFS)	Mount and copy restore	Mount and copy restore
7	VM hosted on Hyper-V (only iSCSI supported)	Based on #1 or 2	Based on #1 or 2
8	Different reseed scenarios such as reseeding from a local copy in the same location	Based on #1, 2, 3, 4, 5, and 6	Mount and copy restore
9	Restoring database to public cloud (Azure, AWS); iSCSI-connected storage only	–	Mount and copy restore
10	In the event of a server/instance rebuild (different host in case original server is affected by virus)	–	Mount and copy restore
11	In the event of a server/instance rebuild (with same configuration/setup including name)	Single File SnapRestore	–
12	Restoration from VMware vCenter to another vCenter group (only for VMDK)	Not supported	Not supported
13	Restoration from VMware vCenter to another vCenter group (iSCSI)	--	Mount and copy restore
14	Database restoration from VM to the physical instance (only iSCSI supported)	--	Mount and copy restore

Restoring a database by restoring to an alternate host option

Restoring to an alternate host is a copy-based restore mechanism. Large databases might take more time to perform a restore operation because it's a streaming process.

A NetApp best practice is to set the timeout value to avoid backup failure at the end. The default setting is three hours (10,800,000ms).

A 1.5TB database restore operation generally takes about nine hours. NetApp recommends setting the following timeout value in the `SMCoreServiceHost.exe.config` file, located in plug-in server in `C:\Program Files\NetApp\SnapCenter\SMCore\`.

```
Search for keyword RESTTimeout and alter the
<add key="RESTTimeout" value="10800000" />
To
<add key="RESTTimeout" value="32400000" /> -- to set 9 hours restoration timeout value.
```

If the database size is too large and you expect more time for the restore operation, then set the restore time accordingly (in ms).

Best Practices

- NetApp recommends altering the `RESTTimeout` value setting before you restore a large database.
- NetApp recommends using the custom log directory option when restoring a database with restore to alternate host if the volume is not accessible in alternate host. This help to read log backup files directly from given location rather than mounting the directory

Clone best practices

Using a clone to create a database copy

To restore a database onto another location on a dev or test environment or to create a copy for business analysis purposes, the NetApp best practice is to leverage the cloning methodology to create a copy of the database on the same instance or alternate instance.

The cloning of databases that are 500GB on an iSCSI disk hosted on a FAS environment typically takes about five minutes. After which user can perform all the require read/write operation on a database.

The cloning of a database can be achieved by dual-method: one is to create a clone from the latest backup and the second is to use a clone life cycle management through which the latest copy can be made available on the secondary instance.

The clone lifecycle management, together with a resource group, can perform cloning of a single database or multiple databases in the same instance. Starting SnapCenter 4.1, clone lifecycle management is supported for standalone and FCI instance.

At a later time, a user can perform a clone split during a peak period or maintenance window to isolate the clone copy from the Snapshot copy and avoid any dependency on either of them in future. To set the restoration timeout value so the large databases do clone split completely successfully, see the section titled, "Restoring a database by restoring to an alternate host option."

Note: A split operation consumes similar time as a restore operation.

SnapCenter allows you to mount the clone copy on the required disk to maintain the format of folder structure on the secondary instance and continue to schedule backup jobs.

Best Practice

For a recurring database restoration job, schedule a clone lifecycle management job from SnapCenter. This option reduces the overhead of managing a job from Microsoft SQL Server to perform a stream backup and restore job.

Managing the clone copy

As previously described, cloning is a simplified and quicker process to restore databases. Managing a clone is a notable task to ensure there are no orphaned storage or ghost entries in the SnapCenter inventory.

SnapCenter maintains a list of clone copies because the copies are created by SnapCenter. Deleting the database from Microsoft SQL Server can make the drives on the host redundant and unaccounted for. Thus leaving the logical drive empty and volume unused.

Best Practice

Always delete the clone copy of a database from SnapCenter

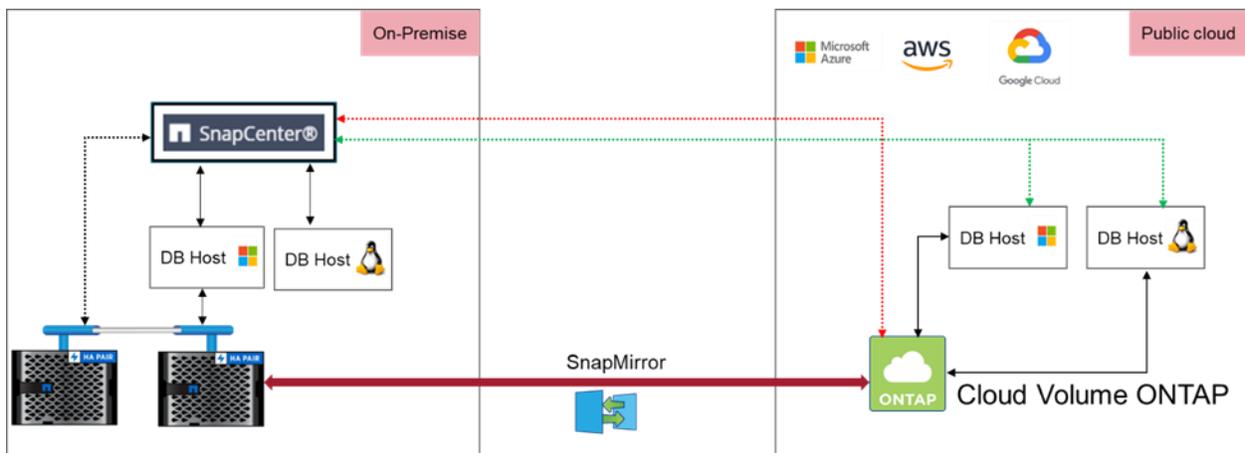
Managing clone copies in hybrid scenarios

SnapCenter supports clone and restore operations of a Microsoft SQL Server database onto another hypervisor built-in on-premises with Hyper-V or on a public cloud such as Azure, AWS, and GCP. The cloning is only possible if the source and target disks are mapped with the iSCSI protocol. Microsoft SQL Server host running on Azure, AWS, and GCP infrastructure can be added as a host to SnapCenter and perform the task.

Figure 8 illustrates a high-level, small-scale hybrid architecture: SnapCenter managing servers on-premises and on the cloud. Cloud Volume ONTAP is used as secondary storage for replication.

NetApp Private Storage (NPS) can also be leveraged for storage through cloud-connected colocation facilities such as Equinix and use hyperscalers such as AWS or Azure for compute operations. This way data can remain private outside of the cloud.

Figure 8) SnapCenter managing data protection of resources across hybrid cloud setup.



Additional cloning guidelines

Follow these additional cloning guidelines:

- Create a backup of the database by using SnapCenter.
- Make sure that the aggregate used by the Microsoft SQL Server database is included in the SVM's list of assigned aggregates for the clone operation to be successful.
- Create a clone on the same host as that of the source database. While creating the clone on an alternate host, ensure that the alternate host meets the following requirements:

- SnapCenter Plug-in for Microsoft Server should be installed on the alternate host.
- The clone host should be able to discover LUNs from primary or secondary storage.
- If you are cloning from primary storage or secondary (Vault or Mirror) storage to an alternate host, then make sure that an iSCSI session is established between the secondary storage and the alternate host or zoned properly for FC.
- If you are cloning from Vault or Mirror storage to the same host, then make sure that an iSCSI session or FC zoning is established between the Vault or Mirror storage and the Microsoft SQL Server host.
- If you are cloning in a virtualized environment, ensure that an iSCSI session is established between the primary or secondary storage and the ESXi server hosting the alternate host or zoned properly for FC:
 - Use the same Microsoft SQL Server version/edition or later as that of the source database host.
 - Use the clone lifecycle management to schedule recurring cloning process.

Microsoft SQL Server deployment for advanced setup

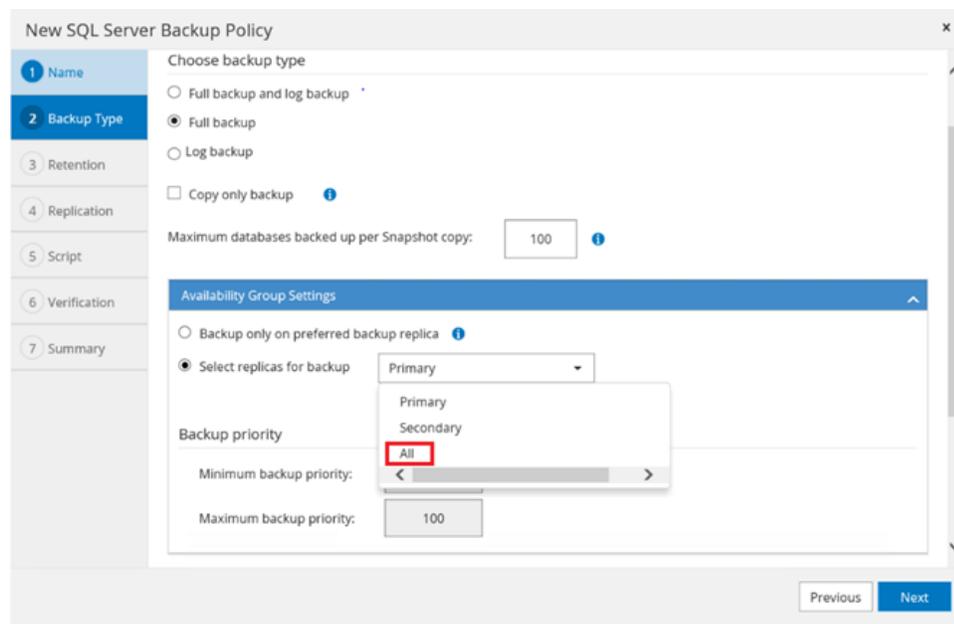
Always On Availability Group

SnapCenter provides various settings for protecting databases in the Availability Group. The data protection policy set through SnapCenter overrides the settings on Microsoft SQL Server Availability Group property. For example, the Reseeding Availability Group database. In the event of a database corruption on either of the nodes, databases are recovered faster by restoring databases from locally available backup.

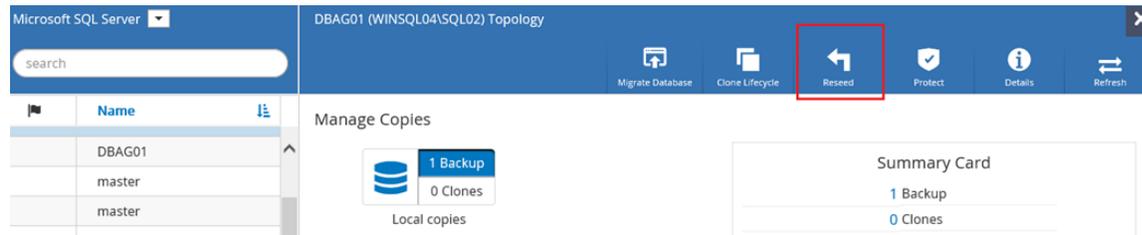
NetApp recommends that you back up databases from all available nodes to ensure faster recovery on each node.

Figure 9 shows the setting to back up databases in the Availability Group on all available nodes. The settings are visible when configuring the policy.

Figure 9) SnapCenter backup policy settings.



Reseeding a database does not require you to rebuild the Availability Group configurations. To restore the database and add it back to the Availability Group, click Reseed under SnapCenter > Resources > Databases and then select the database.



Note: The Reseed option is available on Microsoft SQL Server Availability Group running on Windows Failover cluster instance only.

With SnapCenter 4.4, the SQL Server Availability Group database reseed and restore operations access other replicas log backups as network share instead of mounting the storage.

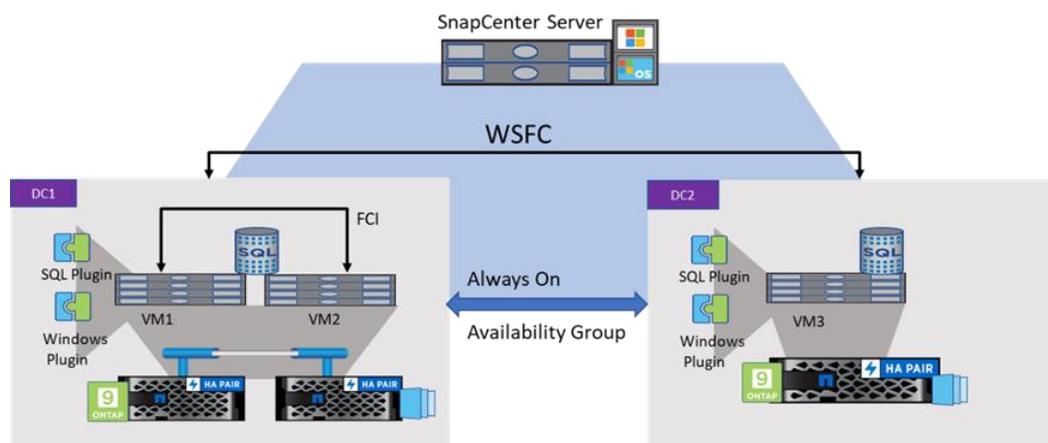
Best Practice

NetApp recommends that you back up databases from all available nodes so that the reseed operation can be performed across all nodes.

Manage Microsoft SQL Server and SnapCenter backups for Asymmetric LUN Mapping in Windows clusters

SCSQL supports the discovery of Microsoft SQL Server and Asymmetric LUN Mapping (ALM) configuration for Microsoft SQL Server high availability and disaster recovery. Such architectures are usually designed to achieve business continuity in the event of a disaster where nodes of Windows cluster can fail over to another data center using Availability Group.

Figure 10) SnapCenter managing SQL Server Availability Group between two failover cluster instance.



In Figure 10, the architecture is supported with Availability Group between two subsets of FCI or between a subset of FCI and a single node. This setup is useful when one to many FCI subset on production site is map to one node in the secondary site with databases residing on respective drives. This setup helps in many to one relationship where minimal servers are present at the disaster recovery location.

A NetApp best practice is to have the same mount point or drive letter across nodes in different data centers within the same Windows Server Failover Cluster (WSFC). Having the same drive letter ensures the reseeding works well without having to alter the restore location.

Note: Provisioning is not supported in an ALM type of architecture.

Performance benchmarking

The data in this section shows the scale of plug-ins that SnapCenter Server can handle. The test performed in the NetApp lab accounted for a customer-centric environment.

Lab setup 1: Using NetApp storage FAS 8020

SnapCenter Server configuration

Table 3) SnapCenter Server configuration.

OS Server	Memory	Processors
Windows 2012 R2	16GB	Four cores

Plug-in configuration

Table 4) Plug-in configuration.

SCSQL Plug-ins	SCO Plug-ins	Databases	DB size	Protection	SVMs	Geo Specific	Host in Primary DC	Host in Secondary DC
60	20	1,200	Up to 2TB	SnapVault/SnapMirror	6	Multisite data center	60	20

Table 5) Storage layout and number of SnapCenter backup groups.

Resource groups	40
LUNS per host	6
LUN/volume mapping	1:1

Schedules and workflows

- Daily backup: Full backup with SnapVault/SnapMirror update with inline verification (one time a day):
 - Ten concurrent backups are scheduled
- Log backup every six hours
- On-demand clone (clone from backup)

Workflow timings

Table 6) Workflow timings.

Serial Number	Workflow	Scale	Timing
1	Full + log backup with inline verification (SM and SV update)	200 DBs/10 hosts	11 minutes

Serial Number	Workflow	Scale	Timing
2	Log backup with SM and SV update	200 DBs/10 hosts	3 minutes
3	Restore a Microsoft SQL Server Database (SnapRestore/ Single File SnapRestore)	1TB	<1 minute
4	Restore a Microsoft SQL Server Database (SubLunFileClone)	1TB	6 minutes
5	Back up with verification	1TB	4 hours 40 minutes
6	Clone from backup	1TB	<2 mins

Conclusion

Microsoft SQL Server is an enterprise-class product. Multiple configuration options are available to suit most of a customer's needs. NetApp storage and data management software is built in a similar fashion, providing users with the flexibility to manage Microsoft SQL Server in a manner that most closely meets their business requirements. With high-performance, easy-to-manage storage systems and robust software offerings, NetApp offers the flexible storage and data management solutions to support Microsoft SQL Server.

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- SnapCenter 4.4 Concepts Guide
<https://docs.netapp.com/ocsc-44/topic/com.netapp.doc.ocsc-con/SnapCenter%20concepts.pdf>
- SnapCenter 4.4 Data Protection Guide for Microsoft SQL Server
<https://docs.netapp.com/ocsc-44/topic/com.netapp.doc.ocsc-dpg-mssql/Protecting%20Microsoft%20SQL%20Server%20databases.pdf>
- SnapCenter 4.4 Administration Guide
<https://docs.netapp.com/ocsc-44/topic/com.netapp.doc.ocsc-ag/Performing%20administrative%20tasks.pdf>

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2021 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.

TR-4714-0221