



SnapCenter Plug-in for Microsoft SQL Server

SnapCenter Software 4.8

NetApp
February 12, 2024

Table of Contents

- SnapCenter Plug-in for Microsoft SQL Server 1
 - SnapCenter Plug-in for Microsoft SQL Server overview 1
 - What you can do with the SnapCenter Plug-in for Microsoft SQL Server 1
 - SnapCenter Plug-in for Microsoft SQL Server features 2
 - Support for Asymmetric LUN Mapping in Windows clusters 3
 - Storage types supported by SnapCenter Plug-ins for Microsoft Windows and for Microsoft SQL Server. . . . 4
 - Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server 6
 - Minimum ONTAP privileges required for SQL plug-in 8
 - Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for SQL server 12
 - Backup strategy for SQL Server resources 13
 - Restoration strategy for SQL Server 17
 - Define a cloning strategy for SQL Server. 20

SnapCenter Plug-in for Microsoft SQL Server

SnapCenter Plug-in for Microsoft SQL Server overview

The SnapCenter Plug-in for Microsoft SQL Server is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Microsoft SQL Server databases. The Plug-in for SQL Server automates SQL Server database backup, verification, restore, and clone operations in your SnapCenter environment.

When the Plug-in for SQL Server is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance or archival purposes.

What you can do with the SnapCenter Plug-in for Microsoft SQL Server

When the SnapCenter Plug-in for Microsoft SQL Server is installed in your environment, you can use SnapCenter to back up, restore, and clone SQL Server databases.

You can perform the following tasks that support backup operations, restore operations, and clone operations of SQL Server databases and database resources:

- Back up SQL Server databases and associated transaction logs

You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.

- Restore database resources
 - You can restore master system databases, msdb system databases, and model system databases.
 - You cannot restore multiple databases, instances, and availability groups.
 - You cannot restore the system database to an alternate path.
- Create point-in-time clones of production databases

You cannot perform backup, restore, clone, and clone lifecycle operations on tempdb system databases.

- Verify backup operations immediately or defer verification until later

Verification of SQL Server system database is not supported. SnapCenter clones the databases to perform verification operation. SnapCenter cannot clone SQL Server system databases, and therefore verification of these databases is not supported.

- Schedule backup operations and clone operations
- Monitor backup operations, restore operations, and clone operations



The Plug-in for SQL Server does not support backup and recovery of SQL Server databases on SMB shares.

SnapCenter Plug-in for Microsoft SQL Server features

The Plug-in for SQL Server integrates with Microsoft SQL Server on the Windows host and with NetApp Snapshot copy technology on the storage system. To work with the Plug-in for SQL Server, you use the SnapCenter interface.

The Plug-in for SQL Server includes these major features:

- **Unified graphical user interface powered by SnapCenter**

The SnapCenter interface provides you with standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup and restore processes across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins. SnapCenter also offers centralized scheduling and policy management to support backup and clone operations.

- **Automated central administration**

You can schedule routine SQL Server backups, configure policy-based backup retention, and set up point-in-time and up-to-the-minute restore operations. You can also proactively monitor your SQL Server environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot copy technology**

The Plug-in for SQL Server uses NetApp Snapshot copy technology with the NetApp SnapCenter Plug-in for Microsoft Windows. This enables you to back up databases in seconds and restore them quickly without taking SQL Server offline. Snapshot copies consume minimal storage space.

In addition to these major features, the Plug-in for SQL Server offers the following benefits:

- Backup, restore, clone, and verification workflow support
- RBAC-supported security and centralized role delegation
- Creation of space-efficient, point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology

A FlexClone license is required on the storage system holding the clone.

- Nondisruptive and automated backup verification
- Ability to run multiple backups at the same time across multiple servers
- PowerShell cmdlets for scripting of backup, verify, restore, and clone operations
- Support for AlwaysOn Availability Groups (AGs) in SQL Server to accelerate AG setup, backup, and restore operations
- In-memory database and Buffer Pool Extension (BPE) as part of SQL Server 2014
- Support for backup of LUNs and virtual machine disks (VMDKs)
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, raw device mapping (RDM), and VMDK over NFS and VMFS



NAS volumes should have a default export policy in storage virtual machine (SVM).

- Support for FileStream and file group in SQL Server standalone databases.


Support for Asymmetric LUN Mapping in Windows clusters

SnapCenter Plug-in for Microsoft SQL Server supports discovery in SQL Server 2012 and later, Asymmetric LUN Mapping (ALM) configurations for high availability, and availability groups for disaster recovery. When discovering resources, SnapCenter discovers databases on local hosts and on remote hosts in ALM configurations.

An ALM configuration is a single Windows server failover cluster that contains one or more nodes in a primary data center and one or more nodes in a disaster recovery center.

Following is an example of an ALM configuration:

- Two failover cluster instances (FCI) in a multi-site datacenter
- FCI for local high availability (HA) and Availability Group (AG) for disaster recovery with a stand-alone instance at the disaster recovery site



WSFC—Windows Server Failover Cluster

The storage in the primary datacenter is shared between the FCI nodes present in the primary datacenter. The storage in the disaster recovery datacenter is shared between the FCI nodes present in the disaster recovery datacenter.

The storage on the primary datacenter is not visible to the nodes on the disaster recovery datacenter, and vice versa.



ALM architecture combines two shared storage solution used by FCI, with non-shared or dedicated storage solution used by SQL AG. The AG solution uses identical drive letters for shared disk resources across data centers. This arrangement of storage, where a cluster disk is shared between a subset of nodes within a WSFC, is referred to as ALM.


Storage types supported by SnapCenter Plug-ins for Microsoft Windows and for Microsoft SQL Server

SnapCenter supports a wide range of storage types on both physical machines and virtual machines. You must verify whether support is available for your storage type before installing the package for your host.

SnapCenter provisioning and data protection support is available on Windows Server. For the latest information about supported versions, see the [NetApp Interoperability Matrix Tool](#).

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
Physical server	SMB3 (CIFS) shares residing on a storage virtual machine (SVM)	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. You cannot use SnapCenter to back up any data or shares using the SMB protocol.
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
VMware VM	Virtual Machine File Systems (VMFS) or NFS datastores	VMware vSphere	
VMware VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. You cannot use SnapCenter to back up any data or shares using the SMB protocol.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	<p>You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	<div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>

Machine	Storage type	Provision using	Support notes
Hyper-V VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	<p>Support for provisioning only.</p> <p>You cannot use SnapCenter to back up any data or shares using the SMB protocol.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>

Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server

A well-designed storage layout allows SnapCenter Server to back up your databases to meet your recovery objectives. You should consider several factors while defining your storage layout, including the size of the database, the rate of change of the database, and the frequency with which you perform backups.

The following sections define the storage layout recommendations and restrictions for LUNs and virtual machine disks (VMDKs) with SnapCenter Plug-in for Microsoft SQL Server installed in your environment.

In this case, LUNs can include VMware RDM disks and iSCSI direct-attached LUNs that are mapped to the guest.

LUN and VMDK requirements

You can optionally use dedicated LUNs or VMDKs for optimum performance and management for the following databases:


- Master and model system databases
- Tempdb
- User database files (.mdf and .ndf)
- User database transaction log files (.ldf)
- Log directory

To restore large databases, the best practice is to use dedicated LUNs or VMDKs. The time taken to restore a complete LUN or VMDK is less than the time taken to restore the individual files that are stored in the LUN or VMDK.


For the log directory, you should create a separate LUN or VMDK so that there is sufficient free space in the data or log file disks.

LUN and VMDK sample layouts


The following graphic shows how you can configure the storage layout for large databases on LUNs:




The following graphic shows how you can configure the storage layout for medium or small databases on LUNs:



The following graphic shows how you can configure the storage layout for large databases on VMDKs:



The following graphic shows how you can configure the storage layout for medium or small databases on VMDKs:



Minimum ONTAP privileges required for SQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later
event generate-autosupport-log
job history show
job stop

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

lun

lun create

lun delete

lun igroup add

lun igroup create

lun igroup delete

lun igroup rename

lun igroup show

lun mapping add-reporting-nodes

lun mapping create

lun mapping delete

lun mapping remove-reporting-nodes

lun mapping show

lun modify

lun move-in-volume

lun offline

lun online

lun resize

lun serial

lun show

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

snapmirror policy add-rule

snapmirror policy modify-rule

snapmirror policy remove-rule

snapmirror policy show

snapmirror restore

snapmirror show

snapmirror show-history

snapmirror update

snapmirror update-ls-set

snapmirror list-destinations

version

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

volume clone create

volume clone show

volume clone split start

volume clone split stop

volume create

volume destroy

volume file clone create

volume file show-disk-usage

volume offline

volume online

volume modify

volume qtree create

volume qtree delete

volume qtree modify

volume qtree show

volume restrict

volume show

volume snapshot create

volume snapshot delete

volume snapshot modify

volume snapshot rename

volume snapshot restore

volume snapshot restore-file

volume snapshot show

volume unmount

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

vserver cifs

vserver cifs share create

vserver cifs share delete

vserver cifs shadowcopy show

vserver cifs share show

vserver cifs show

vserver export-policy

vserver export-policy create

vserver export-policy delete

vserver export-policy rule create

vserver export-policy rule show

vserver export-policy show

vserver iscsi

vserver iscsi connection show

vserver show

network interface

network interface show

vserver

metrocluster show

Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for SQL server

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot copy operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.



If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.



SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary > Mirror > Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the [ONTAP documentation](#).



SnapCenter does not support **sync_mirror** replication.

Backup strategy for SQL Server resources

Define a backup strategy for SQL Server resources

Defining a backup strategy before you create your backup jobs helps ensure that you have the backups that you require to successfully restore or clone your databases. Your Service Level Agreement (SLA), Recovery Time Objective (RTO), and Recovery Point Objective (RPO) largely determine your backup strategy.

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. The RTO is the time by when a business process must be restored after a disruption in 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, RTO, and RPO contribute to the backup strategy.

Type of backups supported

Backing up SQL Server system and user databases using SnapCenter requires that you choose the resource type, such as databases, SQL server instances, and Availability Groups (AG). Snapshot copy technology is leveraged to create online, read-only copies of the volumes on which the resources reside.

You can select the copy-only option to specify that the SQL Server does not truncate transaction logs. You should use this option when you are also managing the SQL Server with other backup applications. Keeping the transaction logs intact enables any backup application to restore the system databases. Copy-only backups are independent of the sequence of scheduled backups, and they do not affect the backup and restore procedures of the database.

Backup type	Description	Copy-only option with backup type
Full backup and log backup	<p>Backs up the system database and truncates the transaction logs.</p> <p>The SQL Server truncates the transaction logs by removing the entries that are already committed to the database.</p> <p>After the full backup is complete, this option creates a transaction log that captures transaction information. Typically, you should choose this option. However, if your backup time is short, you can choose not to run a transaction log backup with full backup.</p> <p>You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.</p>	<p>Backs up the system database files and the transaction logs without truncating the logs.</p> <p>A copy-only backup cannot serve as a differential base or differential backup, and does not affect the differential base. Restoring a copy-only full backup is the same as restoring any other full backup.</p>
Full database backup	<p>Backs up the system database files.</p> <p>You can create full database backup for master, model, and msdb system databases.</p>	<p>Backs up the system database files.</p>
Transaction log backup	<p>Backs up the truncated transaction logs, copying only the transactions that were committed since the most recent transaction log was backed up.</p> <p>If you schedule frequent transaction log backups alongside full database backups, you can choose granular recovery points.</p>	<p>Backs up the transaction logs without truncating them.</p> <p>This backup type does not affect the sequencing of regular log backups. Copy-only log backups are useful for performing online restore operations.</p>

Backup schedules for Plug-in for SQL server

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**.

Number of backup jobs needed for databases

Factors that determine the number of backup jobs that you need include the size of the database, the number of volumes used, the rate of change of the database, and your Service Level Agreement (SLA).

For database backups, the number of backup jobs that you choose typically depends on the number of volumes on which you placed your databases. For example, if you placed a group of small databases on one volume and a large database on another volume, you might create one backup job for the small databases and one backup job for the large database.

Backup naming conventions for Plug-in for SQL server

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

The Snapshot copy 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 copy 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 copy name.

Backup retention options for Plug-in for SQL Server

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.

How long to retain transaction log backups on the source storage system

SnapCenter Plug-in for Microsoft SQL Server needs transaction log backups to perform up-to-the-minute restore operations, which restore your database to a time between two full backups.

For example, if Plug-in for SQL Server took a full backup at 8:00 a.m. and another full backup at 5:00 p.m., it could use the latest transaction log backup to restore the database to any time between 8:00 a.m. and 5:00 p.m. If transaction logs are not available, Plug-in for SQL Server can perform point-in-time restore operations only, which restore a database to the time that Plug-in for SQL Server completed a full backup.

Typically, you require up-to-the-minute restore operations for only a day or two. By default, SnapCenter retains a minimum of two days.

Multiple databases on the same volume

You can put all databases on the same volume, because the backup policy has an option to set the maximum databases per backup (default value is 100).

For example, if you have 200 databases in the same volume, two Snapshot copies are created with 100 databases in each of the two Snapshot copies.

Backup copy verification using the primary or secondary storage volume for Plug-in for SQL Server

You can verify backup copies on the primary storage volume or on either the SnapMirror or SnapVault secondary storage volume. Verification using a secondary storage volume reduces load on the primary storage volume.

When you verify a backup that is either on the primary or secondary storage volume, all the primary and the secondary Snapshot copies are marked as verified.

SnapRestore license is required to verify backup copies on SnapMirror and SnapVault secondary storage volume.

When to schedule verification jobs

Although SnapCenter can verify backups immediately after it creates them, doing so can significantly increase the time required to complete the backup job and is resource intensive. Hence, it is almost always best to schedule verification in a separate job for a later time. For example, if you back up a database at 5:00 p.m. every day, you might schedule verification to occur an hour later at 6:00 p.m.

For the same reason, it is usually not necessary to run backup verification every time you perform a backup. Performing verification at regular but less frequent intervals is usually sufficient to ensure the integrity of the backup. A single verification job can verify multiple backups at the same time.

Restoration strategy for SQL Server

Define a restoration strategy for SQL Server

Defining a restoration strategy for SQL Server enables you to restore your database successfully.

Sources and destinations for a restore operation

You can restore a SQL Server database from a backup copy on either primary or secondary storage. You also can restore the database to different destinations in addition to its original location, enabling you to choose the destination that supports your requirements.

Sources for a restore operation

You can restore databases from primary or secondary storage.

Destinations for a restore operation

You can restore databases to various destinations:

Destination	Description
The original location	By default, SnapCenter restores the database to the same location on the same SQL Server instance.
A different location	You can restore the database to a different location on any SQL Server instance within the same host.
Original or different location using different database names	You can restore the database with a different name to any SQL Server instance on the same host where the backup was created.



Restore to alternate host across ESX servers for SQL databases on VMDKs (NFS and VMFS datastores) is not supported.

SQL Server recovery models supported by SnapCenter

Specific recovery models are assigned to each database type by default. The SQL Server database administrator can reassign each database to a different recovery model.

SnapCenter supports three types of SQL Server recovery models:

- Simple recovery model

When you use the simple recovery model, you cannot back up the transaction logs.

- Full recovery model

When you use the full recovery model, you can restore a database to its previous state from the point of failure.

- Bulk logged recovery model

When you use the bulk logged recovery model, you must manually re-execute the bulk logged operation. You must perform the bulk logged operation if the transaction log that contains the operation's commit record has not been backed up before restore. If the bulk logged operation inserts 10 million rows in a database, and the database fails before the transaction log is backed up, then the restored database will not contain the rows that were inserted by the bulk logged operation.

Types of restore operations

You can use SnapCenter to perform different types of restore operations on SQL Server resources.

- Restore up-to-the-minute
- Restore to a previous point in time

You can restore up to the minute or restore to a previous point in time in the following situations:

- Restore from SnapMirror or SnapVault secondary storage

- Restore to alternate path (location)



SnapCenter does not support volume-based SnapRestore.

Restore up to the minute

In an up-to-the-minute restore operation (selected by default), databases are recovered up to the point of failure. SnapCenter accomplishes this by performing the following sequence:

1. Backs up the last active transaction log before restoring the database.
2. Restores the databases from the full database backup that you select.
3. Applies all the transaction logs that were not committed to the databases (including transaction logs from the backups from the time the backup was created up to the most current time).

Transaction logs are moved ahead and applied to any selected databases.

An up-to-the-minute restore operation requires a contiguous set of transaction logs.

Because the SnapCenter cannot restore SQL Server database transaction logs from log-shipping backup files (log-shipping enables you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances), you are not able to perform an up-to-the-minute restore operation from the transaction log backups. For this reason, you should use the SnapCenter to back up your SQL Server database transaction log files.

If you do not need to retain up-to-the-minute restore capability for all backups, you can configure your system's transaction log backup retention through the backup policies.

Example of an up-to-the-minute restore operation

Assume that you run the SQL Server backup every day at noon, and on Wednesday at 4:00 p.m. you need to restore from a backup. For some reason, the backup from Wednesday noon failed verification, so you decide to restore from the Tuesday noon backup. After that, if the backup is restored, all the transaction logs are moved forward and applied to the restored databases, starting with those that were not committed when you created Tuesday's backup and continuing through the latest transaction log written on Wednesday at 4:00 p.m. (if the transaction logs were backed up).

Restore to a previous point in time


In a point-in-time restore operation, databases are restored only to a specific time from the past. A point-in-time restore operation occurs in the following restore situations:

- The database is restored to a given time in a backed-up transaction log.
- The database is restored, and only a subset of backed-up transaction logs are applied to it.



Restoring a database to a point in time results in a new recovery path.

The following image illustrates the issues when a point-in-time restore operation is performed:



In the image, recovery path 1 consists of a full backup followed by several transaction log backups. You restore the database to a point in time. New transaction log backups are created after the point-in-time restore operation, which results in recovery path 2. The new transaction log backups are created without creating a new full backup. Due to data corruption or other problems, you cannot restore the current database until a new full backup is created. Also, it is not possible to apply the transaction logs created in recovery path 2 to the full backup belonging to recovery path 1.

If you apply transaction log backups, you can also specify a particular date and time at which you want to stop the application of backed up transactions. To do this, you specify a date and time within the available range and the SnapCenter removes any transactions that were not committed prior to that point in time. You can use this method to restore databases to a point in time before a corruption occurred, or to recover from an accidental database or table deletion.

Example of a point-in-time restore operation

Suppose you make full database backups once at midnight and a transaction log backup every hour. The database crashes at 9:45 a.m., but you still back up the transaction logs of the failed database. You can choose from among these point-in-time restore scenarios:

- Restore the full database backup made at midnight and accept the loss of the database changes made afterward. (Option: None)
- Restore the full database backup and apply all the transaction log backups until 9:45 a.m. (Option: Log until)
- Restore the full database backup and apply transaction log backups, specifying the time you want the transactions to restore from the last set of transaction log backups. (Option: By specific time)

In this case, you would calculate the date and time at which a certain error was reported. Any transactions that were not committed prior to the date and time specified are removed.

Define a cloning strategy for SQL Server

Defining a cloning strategy enables you to clone your database successfully.

1. Review the limitations related to clone operations.
2. Decide the type of clone you require.

Limitations of clone operations

You should be aware of the limitations of clone operations before you clone the databases.

- If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the clone operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.
- Cloning of databases from a LUN that is directly attached to a host (for instance, by using Microsoft iSCSI Initiator on a Windows host) to a VMDK or an RDM LUN on the same Windows host, or another Windows host, or vice versa, is not supported.
- The root directory of the volume mount point cannot be a shared directory.
- If you move a LUN that contains a clone to a new volume, the clone cannot be deleted.

Types of clone operations

You can use SnapCenter to clone either a SQL Server database backup or a production database.

- Clone from a database backup

The cloned database can serve as a baseline for developing new applications and help isolate application errors that occur in the production environment. The cloned database can also be used for recovery from soft database errors.

- Clone lifecycle

You can use SnapCenter to schedule recurring clone jobs that will occur when the production database is not busy.

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.