



Restoration strategy for SQL Server

SnapCenter Software

NetApp
October 26, 2021

This PDF was generated from https://docs.netapp.com/us-en/snapcenter/protect-scsql/task_define_a_restoration_strategy_for_sql_server.html on October 26, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Restoration strategy for SQL Server 1
 - Define a restoration strategy for SQL Server 1
 - Sources and destinations for a restore operation 1
 - SQL Server recovery models supported by SnapCenter 1
 - Types of restore operations 2

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.

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.