



Backup strategy for SQL Server resources

SnapCenter Software 4.5

NetApp
August 07, 2024

Table of Contents

- Backup strategy for SQL Server resources 1
 - Define a backup strategy for SQL Server resources 1
 - Type of backups supported 1
 - Backup schedules for Plug-in for SQL server 2
 - Number of backup jobs needed for databases 3
 - Backup naming conventions for Plug-in for SQL server 3
 - Backup retention options for Plug-in for SQL Server 4
 - How long to retain transaction log backups on the source storage system 4
 - Multiple databases on the same volume 4
 - Backup copy verification using the primary or secondary storage volume for Plug-in for SQL Server 5
 - When to schedule verification jobs 5

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.

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.