

Dell SC Series Snapshots and SQL Server Backups Comparison

Dell Storage Engineering
January 2017

Revisions

Date	Description
August 2010	Initial release
January 2012	Revision – Replay 6.0.2
January 2017	Revised for updated tools, terminology, and practices

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2010-2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [1/4/2017] [Best practice guide] [CML1121]

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

- Revisions.....2
- Executive summary.....4
- 1 Native Microsoft SQL Server backup types5
 - 1.1 Full database backups.....5
 - 1.2 Differential database backups5
 - 1.3 Transaction log backups.....5
- 2 SC Series snapshots6
- 3 SC Series snapshot creation methods7
 - 3.1.1 Replay Manager for SQL Server7
 - 3.1.2 SC Series management tools7
- 4 Native SQL Server backups versus snapshots8
- 5 Conclusion10
- A Technical support and resources11
 - A.1 Related resources.....11

Executive summary

One of the most common and most important tasks of a database administrator or system administrator is data protection. Databases must be backed up and recovered in an effective and efficient manner. This document provides an overview of Dell™ SC Series snapshots and how snapshots can be used in conjunction with Native Microsoft® SQL Server® Backups to provide an extremely fast and space efficient point-in-time recovery for very large SQL Server databases. It also provides a comparison of native SQL Server backups and SC Series snapshots.

1 Native Microsoft SQL Server backup types

1.1 Full database backups

A full backup contains all the data in a specific database and its corresponding log files. For large databases, this is a resource intensive operation requiring a significant amount of time, disk space and I/O bandwidth, even when using compression. On their own, full database backups provide single point-in-time recovery. When combined with transaction log backups, a database can be recovered to any point in time.

1.2 Differential database backups

A differential backup is based on the latest full backup of the data. This is known as the base of the differential, or the differential base. A differential base is a full backup of read/write data. Differential database backup will create a copy all data pages that have changed since the last full backup. Together, a full backup and the differential backup can be used to recover a database to the point when the differential was taken or as a base for applying transaction log backups for point in time recovery. Depending on the change rate of the database, the differential backup can be more efficient to create than a full database backup but add an additional step to the recovery process.

1.3 Transaction log backups

A transaction log backup will take every page written to the transaction log since the last transaction log backup and write it to a backup file. A transaction log backup cannot be taken if the database is in simple recovery mode, or if a full database backup has not been completed. Transaction log backups, when used with a full and optionally a differential backup, allow a database to be recovered to any point in time. This provides the capability to restore a database to a point in time prior to a failure with minimal data loss. If the database is damaged, but the transaction log is intact, you can still run a final backup of the transaction log. This will allow a full recovery with no data loss. Transaction log backups are also a key component of log shipping.

There are additional backup types (partial database, differential partial database and file) designed to work with very large databases (VLDB).

A typical SQL Server backup strategy would consist of a full database backup in combination with differential database backups and/or transaction log backups. Strategies using these backup types are complex and are not detailed in this document. Additional information about these backup types can be found in [SQL Server Books Online](#).

All native SQL Server backup types use the same general approach. They read pages out of the database or transaction log, and write them to an external file on the file system. This generates additional I/O and requires additional disk space. For large databases, this additional I/O load can be significant.

2 SC Series snapshots

Taking a snapshot generates very little overhead. This makes database backups very fast, even for large databases. Snapshots are also space efficient, since each snapshot contains only the blocks that have changed since the last snapshot.

When using snapshots to protect a database, a snapshot is taken of each volume that contains database files while I/O within SQL Server is frozen. This preserves the state of each database file at the point in time the snapshot is created. It is strongly recommended to use Dell SC Series Replay Manager when taking snapshots of SQL Server database volumes. Through Microsoft Volume Shadow Copy Services (VSS) integration, this method guarantees application consistent snapshots of the database and transaction logs while the database is still online.

Replay Manager makes it easy to setup and schedule snapshots of all database volumes. Replay Manager provides the ability to schedule snapshots as often as every 15 minutes. In order to understand more on how SC Series snapshots store data, please refer to the *SC Series Administrators Guide*.

Note: The term *snapshot* refers to SC Series snapshots and should not be confused with SQL Server database snapshots or any other snapshot technology.

3 SC Series snapshot creation methods

3.1.1 Replay Manager for SQL Server

Dell EMC strongly recommends using Replay Manager to take snapshots of database volumes while SQL Server is online. Using Replay Manager is the only way to ensure transactional consistency, since writes are frozen inside of SQL Server while the snapshots are taken. Snapshots of active database volumes taken without Replay Manager may not be usable for database recovery. Replay Manager also allows snapshots to be used as the starting point of a restore chain that includes transaction log backups, as databases can be restored without recovery. In addition, restore points created by Replay Manager are fully logged in SQL Server providing complete backup integration.

Even though Replay Manager allows specific databases to be selected, snapshots are taken on the entire volume. If all databases on a given set of volumes are not backed up together, disk space could be wasted on the array. Consider placing the databases for each backup set on a separate set of volumes.

Placing a large database on its own set of volumes allows additional recovery flexibility when using Replay Manager. If a database is on its own volumes and in its own backup set, the database can be recovered very quickly using the Resync method in the Replay Manager Explorer or the Resync-RMRestorePoint PowerShell cmdlet without affecting other databases. This can be a big time saver when recovering a large database using a Replay Manager restore point and transaction log backups.

Take a snapshot of each user database volume at least once per day, with a retention period of 25 hours or longer. Data Progression moves frozen pages in the snapshot according to settings in the storage profile. If the default storage profile Recommended (All Tiers) is used, frozen pages in the snapshot are moved to RAID 5, while writes still occur at RAID 10. This allows Tier 1 space to be used more efficiently.

Tasks ideally suited for using snapshots with SQL Server include:

- Restoring a copy of a database to quickly recover data or objects after errant updates or deletes
- Restoring a copy of a database on the original server or another server for development, testing or reporting
- Replicating databases to a DR location; during a DR test, databases can be recovered without interrupting replication

For more information about using snapshots with SQL Server, see the *Replay Manager Administrator's Guide*.

3.1.2 SC Series management tools

Snapshots can be taken on the SC Series array using any of the SC Series management tools. However, it is strongly recommended to use Replay Manager to protect SQL Server databases. If snapshots are taken on SQL Server volumes outside of Replay Manager, group all the SQL volumes (Data and Logs) into a consistent snapshot profile. In order to learn how to use SC Series arrays to take snapshots and create consistent storage profiles, see the "Managing Snapshot Profiles" section in the *Dell Storage Manager Administrator's Guide*. Remember that snapshots taken outside of Replay Manager produce a crash-consistent snapshot which may or may not be recoverable by SQL Server in every case.

4 Native SQL Server backups versus snapshots

Table 1 Comparing snapshots and native backups

Factor	Replay Manager snapshots	Native backups
Backup Speed	Backups using SC Series Replay Manager are done within seconds (depending on the number of databases in the backup set). Snapshots are independent of the database/volume size.	The time native backups take to back up a database depends on database size. The larger the database, the more time to back it up. In some cases, a backup can take hours depending on the server speed and I/O capabilities of the system.
Space Utilization	Snapshots contain only the blocks of data that have changed since the last snapshot. When a snapshot takes place, the blocks associated have their metadata updated. Since snapshots do not actually create a copy of the data, they are extremely space efficient and smart.	Native backups create a copy of actual database. Hence they consume a lot more space as compared to snapshots, especially for very large databases (VLDB). With recent versions, the backup size can be reduced by compressing the database backups. However, the compressed backups are still large compared to snapshots.
I/O Operations	Initiating snapshots is not an I/O intensive operation as it does not involve creating a huge backup file or writing a lot of data. Since it is a metadata only operation the I/O operations are minimal.	Native backup is an extremely I/O intensive operation as it involves creating a copy of all the data and writing it to disk. Backup compression can reduce the I/O cost but requires additional CPU cycles during the backup and recovery process.
Recovery Speed	Recovering from a snapshot is a very easy and an extremely fast operation. An extremely large database (> 1 TB) can easily be recovered within minutes. This is a great feature especially for highly available databases that have a short RTO (Recovery Time Objective) and cannot tolerate high down-time.	Recovering from a native backup can be an extremely long operation since it depends on a lot of factors such as database size, backup size, and I/O capabilities of the system. A VLDB database (>1 TB) can sometimes take hours to restore and recover.
Recoverability in multiple environments	A single snapshot can be used to recover a database in multiple environments. A copy of the VLDB can easily be created in non-production environments within minutes using Replay Manager.	A single native backup can also be used to create copies in multiple environments. However as mentioned above, it can take a large amount of time for recovery.

Factor	Replay Manager snapshots	Native backups
Data Progression	Snapshots help Data Progression in SC Series arrays. Relative access of blocks within the volume help to determine what data moves to a lower tier (lower performing/higher density) ensuring that regularly accessed data is always available on the fastest disk.	Native backups do not have any impact on the Data Progression in SC Series arrays.
Backup Automation	Taking a snapshot and/or recovering can easily be automated using a PowerShell Script. SC Series Replay Manager provides an extremely useful PowerShell command set that provides powerful and flexible protection and recovery options.	Native backups can easily be automated by using T-SQL or PowerShell scripts.

5 Conclusion

Snapshots provide a method for backup and recovery that cannot be matched by native SQL Server backups for speed and space efficiency. Since snapshots also help Data Progression work better, all environments can greatly benefit from using snapshots. If the recovery point objective (RPO) of the application demands an exact point-in-time recovery, it is recommended to use a combination of SC Series Replay Manager snapshots in conjunction with Native SQL Server Transaction Log backups to comply with recovery requirements.

A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.

A.1 Related resources

Provide a list of documents and other assets that are referenced in the paper; include other resources that may be helpful.

[Dell Storage SC Arrays and SQL Server](#)

[Using Replay Manager with SQL Server](#)