

SQL Server High Availability and Disaster Recovery with Dell EMC SC Series

Abstract

This paper discusses options available with Dell EMC™ SC Series storage and Microsoft® SQL Server® and how they can be used to architect the right solution for even the most challenging environments.

November 2018

Revisions

Date	Description
11/14/2011	Initial release
07/26/2012	Technical Review
10/5/2012	Updated for Enterprise Manager 6.2
December 2016	Complete document update
November 2018	Format updates

Acknowledgements

Updated by: Doug Bernhardt

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.

© 2011–2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. Published in the USA. [11/21/2018] [Best Practices] [CML1119]

Dell 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
Acknowledgements.....	2
Executive summary.....	4
1 Overview of Dell EMC SC Series tools for HA/DR.....	5
1.1 Snapshots.....	5
1.2 Replay Manager	5
1.3 Live Migration	6
1.4 Online Thin Import.....	6
1.5 Replication.....	6
1.5.1 Asynchronous replication	7
1.5.2 Semi-synchronous replication	7
1.5.3 Synchronous replication	7
1.5.4 Understanding replication requirements.....	7
1.6 Synchronous Live Volume w/Auto Failover.....	8
2 Comparing Microsoft SQL Server and SC Series availability features	9
2.1 High availability features.....	9
2.1.1 Asynchronous replication or log shipping.....	9
2.1.2 Asynchronous replication or Always On Availability Groups.....	9
2.1.3 Synchronous replication or Always On Availability Groups.....	9
2.2 Other considerations	10
2.2.1 Expertise.....	10
2.2.2 Licensing.....	10
2.2.3 Compatibility	10
2.2.4 Network.....	11
3 Preparing SQL Server for SC Series Replication.....	12
3.1.1 Volume and file selection.....	12
3.1.2 Tempdb considerations	12
3.1.3 Replay Manager	12
3.1.4 Database level or instance level protection.....	12
3.1.5 SQL Server tuning	12
3.1.6 Testing.....	13
3.1.7 Scripting.....	13
A Additional resources.....	14
A.1 Technical support and resources	14
A.2 Related resources	14

Executive summary

Microsoft® SQL Server® is used to host databases that can have a wide variety of high availability (HA) and disaster recovery (DR) requirements. Both Microsoft SQL Server and Dell EMC™ SC Series storage arrays have a number of capabilities that can be used to design solutions to fit these requirements. Since designing for stringent Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) is complex and takes many business and technology requirements into account, there is no single solution that works best in all cases. The purpose of this paper is to discuss options available with both Dell EMC SC Series and Microsoft SQL Server and how they can be used to architect the right solution for even the most challenging environments. While hypervisors also provide HA and DR features, those features are not discussed in this document.

1 Overview of Dell EMC SC Series tools for HA/DR

Several types of events can have an effect on the availability of SQL Server databases and the applications that use them. These events range from unplanned events, such as natural disasters that may impact an entire data center, to planned events such as hardware or software upgrades. Dell EMC SC Series arrays have several features that enable customers to deploy architectures that are resilient to these types of events and increase the overall availability of SQL Server databases. This section gives a brief overview of each feature, as well as best practices and use cases when implementing these features in SQL Server environments.

1.1 Snapshots

One of the fundamental features that SC Series arrays provide is the ability to perform storage-based snapshots at the volume level. Since snapshots are a metadata operation, they allow a recovery point to be created in seconds, no matter how large the volume, while requiring minimal space. The speed and space efficiency of array-based snapshots simply cannot be matched by traditional protection methods. This unique capability is the building block for other SC Series availability and protection features.

When the boot volume is also hosted on the array, the entire server can be protected with snapshots. This is extremely useful when patching systems and/or performing software upgrades. The system can be quickly reverted to its prior state if something goes wrong. Typically, full system backups would take many minutes or even hours to create and a similar duration to restore. Since reverting the system from snapshots simply involves metadata operations and volume remapping, the entire recovery process can be completed in a fraction of the time, minimizing downtime for critical applications.

1.2 Replay Manager

Replay Manager is the recommended way to take snapshots of SQL Server database volumes. Replay Manager is an application that creates and manages application consistent snapshots on an SC Series array. Application consistency is provided using the Microsoft Volume Shadow Copy Service (VSS). The Replay Manager VSS provider, along with VSS communicate between SQL Server and the SC Series array to ensure application consistency. This provides reliable SQL Server database recovery from snapshots that can be used for same server recovery, recovery on a different server, or recovery from a different array.

While snapshots can be taken on the array without Replay Manager, SQL Server is not aware of snapshots taken without Replay Manager and in-flight data that has not been written to storage will be excluded from the snapshot. Using Replay Manager ensures that data is flushed to disk and writes are held before taking a snapshot.

Replay Manager provides several recovery options that simplify the process of snapshot-based database recovery. When recovering databases on the original server, Replay Manager restore points can be used with SQL Server log backups to perform point-in-time recovery, using either the Restore or Resync method. The Resync method provides an extremely fast way to recover a large SQL Server database without taking the time required by a full database restore from traditional backups. The Expose feature in Replay Manager can also be used to perform a side by side recovery of a database for scenarios where object-level recovery or further data analysis is needed. The [Using Replay Manager with Microsoft SQL Server](#) video presents an excellent overview of Replay Manager functionality. For a complete description of Replay Manager features, consult the Replay Manager Administrators Guide. There are also several videos and white papers on [SC Series technical documents](#) that cover Replay Manager features and recovery scenarios in detail.

1.3 Live Migration

Occasionally, there is a need to migrate storage between SC Series arrays for hardware refreshes, change of location, or other reasons. Typically, this would involve replicating the database to new volumes, using either SQL Server log shipping or database mirroring, and then taking a brief outage as the database on the source volumes is swapped out for the database copy on the new volumes.

Live Migration is a feature that allows the migration of volumes between SC Series arrays with zero downtime. This allows planned migration, maintenance, and upgrades to occur without interrupting access to SQL Server databases. In addition to avoiding the outage that would be necessary with SQL Server tools, features such as Quality of Service (QoS) and deduplication can be used for greater control and efficiency. QoS allows bandwidth limits and time windows to be set for data transfer. Replication deduplication can make the data transfer more efficient by reducing the amount of redundant data that needs to be transferred, resulting in a smaller payload.

To prepare for a Live Migration, the SQL Server host is configured to have storage mapped to both arrays involved in the migration. Once the Live Migration is started, the data is synchronized in the background and using Live Volume technology, I/O requests are sent to the proper array. During Live Migration, I/O is uninterrupted and disk signatures are maintained so that SQL Server applications are unaware that any storage changes are taking place. Performance degradation to some extent is to be expected and will depend on connectivity and data transfer volume. Once the Live Migration is complete, volume ownership and I/O requests are transferred to the new array. Live Migration is a great tool to use when performing planned storage maintenance or upgrades at the same site or different sites when storage connectivity needs to be maintained to avoid an outage. For more information consult the [Dell EMC SC Series Live Migrate and Volume Advisor best practices guide](#).

1.4 Online Thin Import

Thin Import is a feature that works similar to Live Migration and for the same use cases, but instead of migrating volumes between SC Series arrays, Thin Import allows the migration of volumes from Dell EMC PS to SC Series with minimal downtime. During the online Thin Import process, a short outage is taken to configure Thin Import on the PS Series volume and re-present it to the host as an SC Series volume. Once this configuration is complete, the volume can be brought back online. The data import and final cutover will complete in the background while the applications are online. Similar to alternative SQL Server methods such as log shipping and database mirroring a brief outage is required. Thin Import also supports QoS, allowing bandwidth and schedule settings for data transfer. For more information on Online Thin Import consult the [Data Migration from Dell PS Series or PowerVault MD3 to Dell EMC SC Series Storage using Thin Import deployment and configuration guide](#).

1.5 Replication

Creating a high availability solution for SQL Server often involves creating a copy of the data on another storage device and synchronizing that data in some manner. Dell EMC SC Series replication provides data synchronization between multiple arrays. Data is replicated at the volume level allowing a choice of replication settings on a per volume basis. Using SC Series replication can be an effective way to protect SQL Server databases due to the flexibility and configuration options that it provides. The variety of options provide a robust way to develop a replication scheme that provides the proper mix of performance and bandwidth efficiency while still meeting RTO and RPO requirements.

Note: The term replication in this document refers to SC Series replication. While Microsoft SQL Server does offer various types of data replication, they are not typically used to satisfy high availability requirements and therefore not discussed in this document.

1.5.1 Asynchronous replication

Dell EMC SC Series asynchronous replication offers a great amount of flexibility. When the Replicate Active Snapshot option is not used, replication takes place whenever a snapshot is taken. In this mode, changes are sent in a batch on a schedule that is defined by the snapshot frequency. When the snapshot is taken (preferably with Replay Manager) those changes are forwarded to the volumes at the destination.

1.5.2 Semi-synchronous replication

Semi-synchronous replication is a term used to describe asynchronous replication with the Replication Active Snapshot option. Changes are queued and replicated as soon as they are written on the source. Assuming the storage network and destination array can keep up, this provides a near zero data loss solution. Semi-synchronous replication is a very practical solution since it can provide near-real-time synchronization without the additional latency cost of a synchronous replication solution. Since write order is not guaranteed between replicated volumes, the recovery point for databases hosted on multiple replicated volumes will likely be the latest replicated snapshot and recovery based on the active snapshot would be best-effort depending on the crash recovery mechanism of SQL Server to bring the database to a consistent state.

1.5.3 Synchronous replication

For replication solutions where zero data loss is needed, SC Series synchronous replication may be used. One of the key differences between synchronous replication and asynchronous replication is that with synchronous replication, writes to the source storage center must also be replicated to the destination storage center and acknowledged before the write is acknowledged to the host. This guarantees that the data is protected. This behavior is provided in the High Consistency mode of synchronous replication. The guaranteed consistency is provided at the cost of higher overall write latency since the total write latency must also include the write to and acknowledgment from the secondary array.

High Availability mode allows the source volume to accept writes even if the destination volume is unavailable. In High Availability mode, if the destination volume is unavailable, replication will essentially switch to asynchronous replication. When the destination volume is available and resynchronized, it switches back to synchronous replication. The tradeoff for increased replication resiliency is the potential for data loss if a catastrophic failure on the primary volume occurs while the secondary volume is not up to date. Since this potential exists, High Availability mode does not fulfill a zero-data-loss requirement. The paper [Dell EMC SC Series Storage: Synchronous Replication and Live Volume](#) is a great resource to learn more about SC Series replication.

1.5.4 Understanding replication requirements

When planning a replication architecture, it is important to understand what the costs are to determine the replication type and the infrastructure required. In some cases, the RPO requirements are set and the infrastructure is designed to support them. In other cases, the infrastructure may be set and the replication features that provide the best RPO are chosen. Either way, understanding the data change rate and the amount of data to replicate is important in determining the design.

SC Series arrays have a replication feature called Simulate Replication that help determine what the actual cost of replication will be. This feature allows setup and configuration of replication and then tracks and

reports on the amount of data that would be replicated given the selected options. This information can be used to make better replication architecture decisions. If you are satisfied with the simulated replication, it can be converted into an actual replication.

1.6 Synchronous Live Volume w/Auto Failover

[Always On Failover Cluster Instances](#) (FCIs) are one of the key technologies recommended by Microsoft for providing High Availability for SQL Server. The recommendation is to use a SAN for the shared cluster storage. Storage Center has long supported FCIs for local clusters utilizing the same array. However, there are times when using separate arrays is required. This may be for increased redundancy within the same data center or stretch cluster configurations within a campus or across separate locations. This is sometimes referred to as a metro-cluster or geo-cluster. For these scenarios, SC Series Arrays include a feature called Live Volume that allows the same volume to be presented from two different arrays to the cluster nodes, fulfilling the shared storage requirement for Always On FCIs.

When using the Live Volume feature with SQL Server, a non-uniform configuration is usually desired, especially when multiple sites are involved to help ensure that both the primary Live Volume and the SQL Server instance are running at the same site and front end I/O traffic is located within the site. To support this configuration, there are some best practices that should be followed. First, Live Volumes should be synchronous in High Availability mode. High Availability mode allows Auto Failover to be used for the Live Volume. If the primary volume fails without Auto Failover enabled, Live Volumes will be unavailable until the secondary is manually recovered. Second, the Swap Roles Automatically option should also be enabled. This allows the primary and secondary to swap roles so the clustered disk resources are writeable.

It's important to understand that the SQL Server Failover Cluster Instance is reacting to the loss of storage visibility to trigger the failover and it is not coordinated in any way by the array or Dell EMC Storage Manager. Also, there are several configuration options for customizing Live Volume to meet various manual and automatic failover requirements. The [Live Volume with Auto Failover Support for Microsoft](#) video is a great resource for in-depth understanding of Synchronous Live Volume with Auto Failover.

Note: Replay Manager is not compatible with Live Volume.

2 Comparing Microsoft SQL Server and SC Series availability features

2.1 High availability features

When selecting an architecture for high availability, it may be beneficial to compare SQL Server features to SC Series array features. [High Availability Solutions \(SQL Server\)](#) recommended by Microsoft include Always On Failover Cluster Instances (discussed above), Always On Availability Groups and log shipping. While Microsoft SQL Server does offer various types of data replication they are not listed in this section, as they are not typically used for to satisfy high availability requirements.

[Log shipping](#) automates the process of sending transaction log backups and restoring them on the destination server. The RPO for log shipping is typically measured in minutes based on the backup frequency and backup file availability.

The [Always On Availability Groups](#) feature provides database protection by replicating all database changes for a set of primary databases to one or more sets of secondary databases. Availability groups can be configured to replicate asynchronously or synchronously.

2.1.1 Asynchronous replication or log shipping

Asynchronous replication offers a great amount of flexibility due to the various options that it supports. When the Replicate Active Snapshot option is not used, replication takes place whenever a snapshot is taken. In this mode, it operates somewhat like SQL Server log shipping in that changes are sent in batch on a schedule that is defined by the backup/snapshot frequency. When a snapshot is taken (preferably with Replay Manager), those changes are forwarded to the volumes at the destination. For example, configuring SQL Server log shipping with an hourly backup schedule would provide a similar RPO as asynchronous replication with a Replay Manager Job on an hourly schedule.

2.1.2 Asynchronous replication or Always On Availability Groups

When using asynchronous replication with the Replicate Active Snapshot option, writes are queued and replicated as soon as they are written on the source. Assuming the storage network and destination array can keep up, this provides a near zero data loss solution. Although there are design factors to take into account when comparing the two, asynchronous replication could provide a similar RPO to Always On Availability Groups in asynchronous mode.

2.1.3 Synchronous replication or Always On Availability Groups

For situations where zero data loss is required, synchronous replication is a great solution. The key difference between synchronous and asynchronous replication, is that with synchronous replication, any writes to the source volume must also be written to the destination volume before the write is acknowledged to the host. In High Consistency mode of synchronous replication if the destination volume is unreachable, then the source volume will not accept writes until the issue is resolved. This prevents data loss in the event that the destination volume is unreachable for any reason. High Availability mode of synchronous replication allows the primary volume to remain online in the event the secondary volume is unavailable. The tradeoff is that higher availability is offered at the risk of data loss, if the primary volume were to fail before the issue with the secondary volume could be resolved.

Always On Availability Groups in synchronous mode operates similar to SC Series Synchronous Replication in High Availability mode. Data is replicated synchronously. In the event the connection to the secondary exceeds the timeout threshold, the primary continues to accept writes, without replicating to the secondary, until the issue is resolved.

When using either synchronous replication mechanism there is an additional cost that is introduced due to the nature of the feature. Since both require acknowledgment from the secondary location prior to being committed at the primary, the round-trip latency from the primary to the secondary is added to the transaction. This is probably the largest deterrent from using a synchronous data replication approach.

There may not be a 1:1 relationship between a business transaction and a database transaction. It completely depends on application design. For example, examine a simple business transaction such as saving new customer data. To implement this as a database transaction, it could be a single database insert with all the information, or it could be a separate database insert for customer name, home address, street address, home phone, work phone, cell phone, etc. Even in this simple example, the additional latency impact to the application varies greatly depending on the implementation. A synchronous replication implementation, regardless of the method chosen, should be tested to assess the actual cost.

2.2 Other considerations

Both Microsoft SQL Server and SC Series arrays offer several features that can be used to satisfy high availability and disaster recovery requirements for SQL Server databases. Snapshots, mirroring, and replication features for database protection are available in both products but with very different implementations largely due to their layer in the architecture stack. This sometimes makes it difficult to do a direct technical comparison of these features. In addition to comparing them on a feature level, there are other factors that are worth considering.

2.2.1 Expertise

SQL Server HA and DR features can be complex and require expert level DBA skills to architect and maintain. Not only is a dedicated DBA resource likely required, but the design and day to day troubleshooting of HA and DR features are typically duties of a senior SQL Server DBA. This may increase the number of staff and/or skillsets that are required to setup, maintain, and recover these environments. With SC Series, configuring HA and DR features is a very simple process using the same tools and concepts that storage administrators are already familiar with. While DBAs need to be aware of the process, this presents the opportunity to offload this workload and allow them to focus on other tasks. If a question arises, CoPilot support is available 24x7 to assist with questions or issues that may arise.

2.2.2 Licensing

Software licensing costs can be an important part of the architecture decision process and should be considered. In large architectures with hundreds of database instances, licensing costs could be substantial. If availability and/or recovery features are a determining factor in choosing a SQL Server version, it may be more cost-effective to use array-based features instead. Examine the license costs for different architectures to determine if there is a significant cost difference that may play a role in the architecture decision.

2.2.3 Compatibility

While Microsoft SQL Server contains a robust set of features for database high availability (HA) and disaster recovery (DR), those features only apply to SQL Server databases. In a database application environment, there are typically other components external to SQL Server such as file storage, configuration files, other

database types, other applications, etc. that require a protection solution as well. Data protection at the storage layer provides broad capabilities and recovery features of the SC Series arrays can be used for a wide range of products and data protection needs.

2.2.4 Network

Another difference to keep in mind when considering SQL Server synchronization features versus Dell EMC Storage synchronization features is that the SQL Server traffic will take place on the regular network while the Dell EMC Storage features will use the storage network. Depending on your network architecture these may be separate physical networks or separate VLANs on the same physical network. Consider the capabilities and capacities of each to determine if this may be a factor.

3 Preparing SQL Server for SC Series Replication

When using SC Series replication for data protection, in addition to choosing the proper type of replication there are other considerations that should be taken into account. Paying careful attention to the following items will assist in performance, efficiency, and recoverability of the replication solution.

3.1.1 Volume and file selection

Since replication occurs at the volume level, all volumes that contain SQL Server data will need to be replicated. This includes all volumes that contain data files, log files, and any other files that comprise the database set. Make sure only data that needs to be replicated exists on replicated volumes. If databases exist on those volumes that do not require replication, they can create unnecessary replication traffic.

3.1.2 Tempdb considerations

Tempdb is used as a temporary work area for SQL Server and is rebuilt, rather than recovered when SQL Server restarts. This means that while the availability of tempdb is important, data recovery is not. Therefore, there is no need to replicate tempdb data or log files.

3.1.3 Replay Manager

Replay Manager should be used to take snapshots at regular intervals on replicated volumes. This will enable reliable recovery from snapshots in addition to the replicated copy of the volume in the event that replication is out of sync or a prior recovery point is needed. Snapshots taken with Replay Manager are preferred to regular array-based snapshots since VSS flushes dirty pages from memory and then holds writes while taking the snapshots, ensuring reliable recovery from snapshots. With array-based snapshots, there is a risk that the database may not be in a recoverable state.

3.1.4 Database level or instance level protection

It is important to understand that by protecting the database volumes, only selected databases are being protected, not the entire SQL Server Instance. It is no different than using the SQL Server HA/DR features, such as Database Mirroring or Always On Availability Groups, in that regard. The message here is that there are other critical components in the SQL Server Instance that may need to be protected, such as logins or SQL Agent jobs. A process will be required to accomplish that. While these components are stored in the system databases that can be backed up, the recovery steps can be somewhat complicated unless the target SQL Server instance will have identical hardware configuration, instance name, server configuration and version. Typically, these things may be different at the target site so the most common approach is to protect these items with deployment or automated scripts.

Another approach to providing instance protection is to replicate all volumes and use the boot from SAN feature. In this manner the entire server is protected and the replicated volumes could be recovered on an identical server at the target site.

3.1.5 SQL Server tuning

There are a couple of items that can be considered from a SQL Server perspective to optimize the Dell EMC SC Series data replication process. The first one is to enable data compression inside the database. Compressing the SQL Server data pages will reduce overall I/O and therefore reduce the amount of data that needs to be replicated. Compression benefit and cost will vary depending on the data so while it's easy to implement, the performance impact will need to be assessed. Second, evaluate the database applications if possible to see if there are any efficiencies to be gained. While it can be a difficult task, application and

database tuning offer the largest gains in terms of performance and scalability not only for replication but for the entire system and the importance can't be emphasized enough.

3.1.6 Testing

Testing an HA/DR architecture is critical to ensure that if an issue arises the system performs as expected. Often this is difficult to do and requires careful scheduling and coordination to perform an actual failover to perform verification. Dell Storage Manager includes a feature called Test Activate Disaster Recovery that allows the test of a disaster recovery solution to be performed online without disrupting replication or requiring a live failover to the secondary site. Since the replication features are built on snapshot technology, the Test Activate Disaster Recovery feature automates the process of creating and mounting view volumes from snapshots on the secondary site. Since these view volumes provide full read/write access, the databases can be brought online to exercise the failover procedure. Test Activate Disaster Recovery is a great feature that can be used to verify that the DR plan works as expected, providing confidence in the HA/DR architecture.

3.1.7 Scripting

PowerShell is a great tool for scripting tasks in a Windows environment. This can reduce the time it takes to perform complex tasks while also minimizing mistakes. SC Series arrays have deep PowerShell integration. All of the SC Series features discussed in this document can be managed through PowerShell. Using PowerShell, management of SQL Server, Windows, and storage resources can all be performed in a single script. For more information and samples on getting started with PowerShell and SC Series Arrays see the [Dell Storage PowerShell SDK Cookbook](#).

A Additional resources

A.1 Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

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

A.2 Related resources

[Microsoft High Availability Solutions \(SQL Server\)](#)

[Microsoft SQL Server Replication Features and tasks](#)

[Using Replay Manager with Microsoft SQL Server](#)

[Dell EMC SC Series Live Migrate and Volume Advisor best practices guide](#)

[Data Migration from Dell PS Series or PowerVault MD3 to Dell EMC SC Series Storage using Thin Import](#)

[Dell EMC SC Series Storage: Synchronous Replication and Live Volume](#)

[Microsoft SQL Server Always On Failover Cluster Instances](#)

[Live Volume with Auto Failover Support for Microsoft](#)

[Microsoft SQL Server Log Shipping](#)

[Microsoft SQL Server Always On Availability Groups](#)

[Dell Storage PowerShell SDK Cookbook](#)