Technical Report

Microsoft SharePoint and SnapManager 8.0 for SharePoint with Clustered Data ONTAP: Best Practices Guide

Cheryl George, Rob Barker, NetApp

November 2013 | TR-4243

Executive Summary

This document discusses the planning considerations and best practices when deploying Microsoft® SharePoint® 2013 and Microsoft SharePoint 2010 on NetApp® storage systems running clustered Data ONTAP®. It also covers the best practices for the NetApp enterprise data management solution for SharePoint, which is called SnapManager® 8.0 for SharePoint.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

This document describes the best practices and design considerations when deploying SharePoint Foundation 2013 and SharePoint Server 2013, as well as SharePoint Foundation 2010 and SharePoint Server 2010 (this includes all currently released service packs) and SnapManager for SharePoint (SMSP 8.0) on NetApp storage systems running clustered Data ONTAP. This document provides guidance for achieving effective and efficient storage planning, end-to-end data protection, and SharePoint system performance. The scope of this guide is limited to technical design guidelines based on the design principles and preferred standards that NetApp recommends for storage infrastructure when deploying SharePoint. The end-to-end implementation is out of scope of this report.

The best practices and recommendations described in this guide enable Microsoft SharePoint Server architects and NetApp storage administrators to plan a highly performing, available, and easy-to-manage SharePoint environment and to meet stringent service-level agreements (SLAs). It is assumed that the reader has working knowledge of the following:

- NetApp clustered Data ONTAP operating system
- NetApp SnapDrive® for Windows® (SDW) data management software
- NetApp SnapManager for SQL Server® (SMSQL)
- NetApp SnapManager for SharePoint (SMSP)
- Microsoft SharePoint Server 2010, 2013 architecture and administration
- Microsoft SQL Server 2012, 2008 R2, or 2008

To know more about the configuration compatibility with NetApp and Microsoft software, refer to the NetApp Interoperability Matrix Tool.

1.1 Intended Audience

This document is for experienced SharePoint administrators, IT managers, and storage administrators who have reviewed the following NetApp product documentation:

- NetApp SnapDrive for Windows (SDW)
- SnapManager for Microsoft SQL Server (SMSQL)
- SnapManager for Microsoft SharePoint (SMSP)
- Clustered Data ONTAP

2 Servers in SharePoint 2013 Farm

The roles of a Microsoft SharePoint Server 2013 farm can be deployed as a standalone server or across many servers. The roles of the servers include:

- **Web server role**. This server responds to user requests for web pages. The web front end (WFE) can typically be load balanced by using either the Windows Server® network load-balancing feature or other third-party software or hardware.
- **Application server role**. Provides features/services for SharePoint, which can also be load balanced. A few examples of the features/services provided include central administration, access services, Excel® services, user profile service, and secure store service. For a full list of services that can be configured with SharePoint 2013, refer to Configure services and service applications in SharePoint 2013.
- **Database server role**. Stores content, configuration, administration, and service data. For details on high-availability and disaster recovery options for various SharePoint 2013 databases, refer to Microsoft TechNet article Supported high availability and disaster recovery options for SharePoint databases (SharePoint 2013).

For details on logical and physical architectures for SharePoint farms, refer to Architecture design for SharePoint 2013 IT pros.

3 Storage Efficiency and Manageability

Storage efficiency is the ability to store and manage Microsoft SharePoint data efficiently to consume the least amount of space with negligible or no impact on the overall server performance. Storage efficiency goes beyond just data deduplication; it is a combination of a RAID, thin or thick provisioning (overall layout and utilization), disk mirroring, and other data protection technologies.

3.1 NetApp Technology for Storage Efficiency

The NetApp technologies described in this section help to implement storage efficiency and to reap cost-savings benefits by optimizing existing storage in the infrastructure as well as deferring or avoiding future storage expenditures. The more these technologies are used in conjunction, the larger the savings.

Flash Cache

To help improve the storage efficiency and read I/O performance and latency of SATA-based deployments, Flash Cache™ should be used. Flash Cache is a controller-wide read cache that can be installed on certain models of NetApp storage controllers.

Flash Cache technology is suitable:

- When housing search index data
- For workloads that are 95% read intensive

When binary large object (BLOB) content is externalized to NetApp CIFS shares on SATA, NetApp recommends Flash Cache for random read workloads such as browser-based page visits in SharePoint.

For details on Flash Cache, refer to Flash Cache Best Practices Guide.

For more details on Flash Cache modes, refer to Flash Cache: Modes of Operation.

Flash Pool

NetApp Flash Pool™ is an aggregate-level read-write cache option with solid-state drives (SSDs) and hard disk drives (HDDs) in a single storage pool (aggregate), with the SSDs providing a fast response time cache for volumes that are provisioned on the Flash Pool aggregate. In addition to read caching, Flash Pool also provides write caching. This technology is well suited for the following databases:

- **tempdb (SQL Server system database.** Needs to be considered as a working area for SharePoint operations. In SharePoint, every action that you take is staged in the tempdb before it is committed.
- **Transaction log.** Critical component of a database because it records all transactions and database modifications made by each transaction.
- **Search service application database.** Important in SharePoint 2013 search in that it is a combination of SharePoint 2010 search and fast search for SharePoint 2010.
- **SharePoint content database.** Collaboration for document publishing workloads.

Snapshot

NetApp Snapshot™ technology provides low-cost, fast-backup, point-in-time copies of the volume for SMB shares or LUNs that host the SharePoint databases.

The cost for Snapshot copies is only at the rate of block-level changes, not 100% for each backup as with mirror copies. It also allows the database clones to be created near instantaneously, because no data is actually copied. It can result in savings in storage costs for backup and restore purposes and opens up a number of efficient data management possibilities. SMSP creates Snapshot copies of the production SharePoint databases stored in LUN/SMB shares, through integration with Microsoft Volume Shadow Copy Service (VSS) technology built into Windows Server.

Best Practices

- For better Snapshot copy management, do not create LUNs on the same storage system volume if those LUNs have to be connected to different hosts.
- Avoid scheduling Snapshot copies at the same time as SnapMirror® updates or SnapVault® activity, driven from SnapManager for SharePoint using SnapManager for SQL Server and SnapDrive for Windows (SDW). If these schedules conflict, Snapshot copies may not be created.

To understand further about the operation of NetApp Snapshot, refer to Operational How-To Guide NetApp Snapshot Management.

Storage Thin Provisioning

NetApp thin provisioning deployed with NetApp FlexVol® technology allows you to allocate storage on demand as data is written to disk, instead of preallocating all of the capacity ahead of time (sometimes referred to as thick provisioning), thereby optimizing utilization of available storage. Thin provisioning eliminates almost all whitespace, which helps avoid poor utilization rates. FlexVol volume (flexible volume) is the enabling technology behind NetApp thin provisioning, which can be considered as the virtualization layer of Data ONTAP.

When storage consumption is unpredictable or highly volatile, it is best to reduce the level of storage overcommitment so that storage is available for any growth spikes, which in the case of SharePoint happens more often than not. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much overcommitment is acceptable, if any. Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable.

If the time required to procure new storage is very long, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

The potential risk when configuring the SharePoint environment for thin provisioning is a LUN going offline when there is not enough space to write further data.

NetApp thin provisioning allows the storage to be purchased as the application grows, preventing the need to purchase multiple years of future storage up front. As SharePoint storage needs grow, physical disk space can be added without taking it offline or adversely affecting performance.

Best Practices

- When you enable NetApp thin-provisioned LUNs, NetApp recommends deploying these LUNs in FlexVol volumes that are also thin provisioned with a capacity that is 2x the size of the LUN. When a LUN is deployed in this manner, the FlexVol volume acts merely as a quota. The storage consumed by the LUN is reported in FlexVol and its containing aggregate.
- NetApp recommends using thin-provisioned LUNs for maximum storage efficiency. However, when enabling NetApp thin provisioning, administrators should also configure:
  - SMSP retention policies for removing Snapshot copies from these thin-provisioned LUNs
  - Policies on the volumes for automatic sizing of a volume and LUN fractional reserve
- Thin provisioning is recommended for both SharePoint content database and its transaction log volumes.
- Thin provisioning requires additional storage management to make sure there is no impact to overall storage, especially if multiple workloads (SharePoint, Exchange, and so on) are being

hosted on the same system.

- When additional space is required, you can add more disks to the aggregates and provision storage to the user. For more efficient use of disk space in a SnapMirror configuration, use thin provisioning to overcommit aggregates, because SnapMirror requires the destination volume to be the same size as or greater than the source volume.
- Monitor and manage the environment when using thin provisioning.

Space Guarantee

Space guarantee enables thin provisioning. The space guarantee option can be set at the aggregate, volume, or LUN level, depending on the requirements of a SharePoint environment. If the space guarantee at the volume level is set to “volume,” the amount of space required by the FlexVol volume is always available from its aggregate. This means the space is subtracted, or reserved, from the aggregate’s available space at volume creation time. This is the default setting for FlexVol volumes.

If the space guarantee for the volume is set to “none,” the volume reserves no space from the aggregate during volume creation. Space is first taken from the aggregate when data is actually written to the volume. Write operations to space-reserved LUNs in a volume with guarantee=none will fail if the containing aggregate does not have enough available space. LUN reservation makes sure that the LUN has space in the volume, but setting ‘guarantee=none’ does not make sure that the volume has space in the aggregate. When the space guarantee for the volume is set to “file,” the aggregate makes sure that space is always available for overwrites to space-reserved LUNs.

Space guarantee needs to be implemented correctly. The SMSP backups will fail when you run out of space, hence the need to have this monitored.

Space Reclamation

Space reclamation is a feature in SnapDrive for Windows that is used to communicate with the NetApp storage system in order to free up blocks in a LUN that are marked as “free” by the NTFS metadata. In addition to the ability to reclaim space using SnapDrive for Windows, it is also possible to use the Data ONTAP PowerShell Toolkit using `Invoke-NaHostVolumeSpaceReclaim or Invoke-NcHostVolumeSpaceReclaim`. Space reclamation must be initiated from time to time to recover the unused space in a LUN.

Note: Normal data traffic to the LUN can continue while the space reclamation process runs. However, certain operations cannot be performed during the space reclamation process:

- Creating or restoring a Snapshot copy stops space reclamation.
- The LUN may not be deleted, disconnected, or expanded.
- The mount point cannot be changed.
- Running Windows defragmentation is not recommended.

Note: Space reclamation is not supported on VMDKs hosted on VMFS datastores.

Autodelete

Autodelete typically allows Data ONTAP to automatically delete Snapshot copies when a threshold/trigger is met. Autodelete works at the volume level and not on individual LUNs. This means that LUNs will not automatically grow and must be handled separately using NetApp SnapDrive for Windows.

Best Practices

- NetApp recommends that the Snapshot backup retention functionality in SnapManager for SharePoint be used instead of the autodelete feature built into Data ONTAP. When properly configured, SnapManager will delete Snapshot copies as per the retention. If Data ONTAP autodelete is used, SnapManager will not be able to perform its retention duties, and backup sets could break. Snapshot copies should only be deleted through SnapManager, either with the retention or with the delete backup wizard. Snap autodelete needs to be off on all volumes that host SharePoint and SMSP content.
- NetApp recommends not using the autodelete option in NAS/SMB environments to avoid potential deletion of BLOB Snapshot copies.
- NetApp recommends setting Snap reserve to 0 for SAN environments because it simplifies space management, allowing maximum usable volume space by either the LUNs or the Snapshot copies within the volume. It is advised not to keep Snap reserve to the default value of 20% because user writes are already limited by the LUN size.
- When using SnapMirror or SnapVault technology to replicate a SharePoint database, NetApp recommends not using the “disrupt” option for commitment, because SnapMirror baseline Snapshot copies can be destroyed by autodelete. In many configurations, deleting the last SnapMirror Snapshot copy is not desirable because a new full baseline copy is required to resume mirroring operations. For example, if the source and destination are at different sites, recreating this baseline can be a time-consuming and an expensive process.

Autosize

This volume setting for FlexVol volumes defines whether a volume should automatically grow to avoid filling up to capacity. It is possible to define how quickly the volume should grow by using the “-i” option. The default growth increment is 5% of the volume size at creation. It is also possible to define how large the volume is allowed to grow by using the -m option. If volume autosize is enabled, the default maximum size to grow to is 120% of the original volume size. AutoSize works at the volume level and not on individual LUNs. This means that LUNs will not automatically grow and must be handled using NetApp SnapDrive for Windows (SDW).

Best Practices

- NetApp recommends planning for additional buffer space of 20% in the aggregate containing the volume to allow for the expansion of the SMB/CIFS volume used for storing SharePoint data when implementing volume autosize policies.
- NetApp recommends prioritizing autosize over autodelete because deletions occur at the Data ONTAP level, and it is possible to have a backup set of a transaction log and database where one of the Snapshot copies has been automatically deleted or orphaned.

Fractional_reserve

Fractional_reserve or fractional overwrite reserve or LUN overwrite reserve is a volume option that specifies how much space Data ONTAP reserves for Snapshot overwrite after all other space in the volume is used. The default value for `fractional_reserve` is 0.

NetApp FlexClone

A FlexClone® volume is a writable point-in-time Snapshot copy of a FlexVol volume or another FlexClone volume. FlexClone uses space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and the clone. FlexClone volumes are ideal for any situation that involves:

- Testing or development
- Progress being made by locking in incremental improvements
- Need to distribute data in changeable form without endangering the integrity of the original

A common scenario is to use FlexClone volumes in an environment before committing a Microsoft SQL Server rollup or hotfix into production.

FlexClone technology can be leveraged both at the primary storage system and at the SnapMirror destinations for effective utilization of resources. The cloned volume inherits compression and deduplication attributes of the parent; any new data written to the cloned volume can be compressed/deduplicated. FlexClone volumes can also be used for disaster recovery testing without affecting the operational continuity of the Microsoft SharePoint environment.

For detailed information on how FlexClone works and on command line references, refer to Clustered Data ONTAP 8.2 Logical Storage Management Guide.

Best Practice
Use SnapDrive to connect to the required Snapshot copy and to automatically execute the FlexClone operation.

NetApp Deduplication

The deduplication process only stores unique blocks of data in the volume and creates additional metadata in this process. Each 4k block in the storage system has a digital fingerprint, which is compared to other fingerprints in the volume. If two fingerprints are found to be the same, a byte-for-byte comparison is done of all bytes in the block. If they are an exact match, the duplicate block is discarded, and the space is reclaimed. The core enabling technology of deduplication is fingerprints.

Deduplication consumes system resources and can alter the data layout on disk. Due to the application I/O pattern and the effect of deduplication on the data layout, the read and write I/O performance can vary. Dedupe scheduling can have a huge impact on controller performance. Hence, make sure you do not have all dedupe jobs running on default schedule at midnight to avoid system halt during this process. Deduplication might also have a positive increase on performance because the system memory and Flash Cache modules are both aware of the deduplicated blocks. As a block is read, it is inserted into memory or cache. If this deduplicated block is accessed by another operation, it will be accessed from physical memory as opposed to the spinning disk, resulting in a much improved access time.

Note: Deduplication is transparent to SharePoint, which does not recognize the block changes, so the SQL Server database remains unchanged in size from the host, even though there are capacity savings at the volume level. Data compression and deduplication can provide significant space savings, but thorough testing should be done to determine the savings for your environment.

There is a limit on the maximum size of the volume for deduplication, because deduplication depends primarily on the amount of system memory, which varies based on the storage platform. Before opting for deduplication in SharePoint environments, be sure to consider the preceding factor when sizing the volume layout.

For additional information, refer to TR-3966: Compression and Deduplication for Clustered Data ONTAP.

3.2 Best Practice Configurations When Using Thin Provisioning for SharePoint Environments

There are many ways to configure the NetApp storage appliance for LUN thin provisioning; each has advantages and disadvantages. It should be noted that it is possible to have thinly provisioned volumes and non–thinly provisioned volumes on the same storage system or even the same aggregate. The following are considered to be best practice configurations when using thin provisioning for Microsoft SharePoint.

Option 1: Volume Guarantee Set to 'None'

Volume guarantee	= none
LUN reservation	= enabled
fractional_reserve	= 0%
snap_reserve	= 0%
autodelete	= volume/ oldest_first
autosize	= off
try_first	= snap_delete

This configuration has the advantage of free space in the aggregate being used as a shared pool of free space. The disadvantages of this configuration are the high level of dependency between volumes and that the level of thin provisioning cannot easily be tuned on an individual volume basis. When using this configuration, the total size of the volumes is greater than the actual storage available in the host aggregate. With this configuration, storage administrators can generally size the volume so that they only need to manage and monitor the used space in the aggregate. This option does not affect the space for hosting the live data, but rather allows the backup space to dynamically change.

Option 2: Using Autogrow/Autodelete

Volume guarantee	= volume
LUN reservation	= disabled
fractional_reserve	= 0%
snap_reserve	= 0%
autodelete	= volume/ oldest_first
autosize	= on
try_first	= autogrow

This configuration allows the administrator to fine-tune the level of thin provisioning for Microsoft SharePoint environments. With this configuration, the volume size defines and allocates an amount of space that is only available to LUNs within that volume. The aggregate provides a shared storage pool of available space for all the volumes contained within it. If the LUNs or Snapshot copies require more space than available in the volume, the volumes will automatically grow, taking more space from the containing aggregate. Additionally, the advantage of having the LUN space reservation disabled is that Snapshot copies can then use the space that is not needed by the LUNs. The LUNs themselves are also not at risk of running out of space because the autodelete feature will remove the Snapshot copies consuming space.

Note: Snapshot copies used for creating FlexClone volumes will not be deleted by the autodelete option.

Best Practice

NetApp recommends using autogrow for most common deployment configurations.

3.3 Monitoring

When using NetApp efficiency features, the volumes should be appropriately sized so that autosize and/or autodelete policies are not triggered unless there is an abnormal rate of change or a problem with Snapshot copy retention. NetApp OnCommand® Unified Manager Core Package management software that includes Operations Manager is the recommended tool to monitor SharePoint volumes for these events and to send notifications to the storage administration team to follow up further with the SQL Server administration team. SNMP can also be used to monitor these events.

After a notification for a volume autogrow or Snapshot autodelete event has been received by the storage administration team, the recommended action is for the storage administration team to examine the affected storage controllers and then follow up with the SharePoint administration team for further administrative actions.

A typical cause of volume autosize events is that the rate of change greatly surpassed the rate of change assumption used in sizing the volume. Typically, collaboration workloads experience increased data change rates. Another cause for volume autosize events is that older Snapshot copies created by SnapManager for SharePoint Server are not being deleted. As Snapshot copies age, they can grow in size and consume more capacity than originally allocated in the volume.

Reasons for SnapManager for SharePoint not deleting backups are:

- SMSP backups are failing.
- SMSP backup retention policies are not being enforced correctly because Snapshot copies were manually removed outside of SnapManager for SQL Server on the controller itself.

Monitor SnapManager for SharePoint event IDs to understand the health of SnapManager for SharePoint. To monitor the health of SnapManager retention, use Windows PowerShell® commands from the Data ONTAP PowerShell Toolkit and native Windows PowerShell commands.

4 Planning Storage Layout (Aggregates, Volumes, LUNs, and SMB Shares)

The combination of NetApp storage solutions and Microsoft SharePoint enables the creation of enterprise-level database storage designs that can meet most demanding application requirements. To optimize both technologies, the appropriate layout of SharePoint databases is necessary for performance, faster access, recoverability, and management of the SharePoint infrastructure.

Business Requirements for SharePoint Farms and Services

To define business requirements, determine the following for each farm and service in the environment:

- **Recovery point objective (RPO)** is the objective for the maximum time period between the last available backup and any potential failure point. It is determined by how much data that the business can afford to lose if a failure were to occur.
- **Recovery time objective (RTO)** is the objective for the maximum time that a data recovery process will take. It is determined by the time that the business can afford for the site or service to be unavailable.
- **Recovery level objective (RLO)** is the objective that defines the granularity with which you must be able to recover data: whether you must be able to recover the whole farm, web application, site collection, site, list or library, item, or item version level.

Shorter RPO and RTO and finer granularity of RLO all typically cost more. Planning for SharePoint in the context of NetApp is required for optimal performance and backup/restore, keeping in mind key factors such as RPO, RTO, and RLO. A good storage design allows for meeting both business recovery requirements and data growth.

4.1 Aggregates

Aggregates are the lowest level container in the storage stack, containing of physical disks from which volumes are carved out. Fewer large aggregates will maximize performance; however, they might not meet the data availability requirements set forth in the SLA. In SharePoint environments with multiple database copies, Microsoft no longer requires separating database and transaction log files to separate sets of disks. This means that SharePoint database and respective transaction log volumes can be placed in the same aggregate.

NetApp recommends using one large aggregate for all SharePoint databases. There are two reasons for this:

- One aggregate makes the I/O abilities of all spindles available to all files.
- One aggregate enables the most efficient use of disk space.

NetApp has performed various tests using both approaches, testing shared and dedicated aggregates with data and log separated as well as workload types separation (web content management, social, search, mobile access, and so on). The conclusion is that one large aggregate yields significant performance benefits and is easier for administrators to manage. Also, with an increase in size of physical disks, efficient space management using multiple aggregates becomes even more challenging.

When creating and sizing the aggregate for SharePoint environments, consider the following:

- The total size of all the SharePoint databases
- The I/O load generated by all users accessing all the databases that will share the same aggregate
- The projected storage space growth
- The projected user count growth
- Plans for adding new content databases to SharePoint web applications
- Any content other than SharePoint that might use the same aggregate; for example, other SQL Server databases, Exchange, or line-of-business databases

Best Practices

- NetApp recommends placing all the SharePoint content on the same aggregate.
- Make sure your aggregates containing SharePoint data are not larger than they need to be, which in turn affects the time, disk space, and bandwidth required for disk sanitization.
- NetApp recommends having at least 10% free space available in an aggregate hosting SharePoint data for optimal storage performance. For additional information, refer to TR-3929: Reallocate Best Practices Guide.
- If you are creating a FlexClone volume from a SnapMirror destination volume and expanding the aggregates containing the source and destination volumes, expand both source and destination, and use a base Snapshot copy that was created after the source volume was expanded.

4.2 Volume

Data ONTAP enables the creation of FlexVol volumes for managing data without the need to assign physical disks to the volumes. Many volumes can be created in a single aggregate, and each volume can be expanded or shrunk or moved between aggregates. Volume layout is critical in creating and sustaining a highly available SharePoint environment. Careful consideration of backup groups, disaster recovery scenarios, and archiving solutions helps to determine the placement of volumes onto aggregates and the corresponding LUNs onto those volumes.

Volume Design Considerations

Before a database volume design can be created, the backup and recovery requirements must be defined. They provide the “specs” needed for the volume design process.

Following are best practice considerations to apply during the volume design process:

- Place the SQL Server system databases (master, model, msdb) on a dedicated volume to provide separation from the SharePoint databases.
- Place the TempDB on a separate LUN. For optimal performance the TEMPDB data file and log file should be placed on separate LUNs versus sharing a LUN.
- Make sure that the database (mdf and ldf) residing on LUN within a volume is separate from that used by the SnapInfo LUN to avoid stream-based backup.

If each database has a unique backup requirement:

- Either create separate FlexVol volumes for each database and transaction log, or
- Place databases with similar backup and recovery requirements on the same FlexVol volume. This can be a good option in cases of many small databases in a SQL Server instance.

Figure 1) SMSP volume layout.

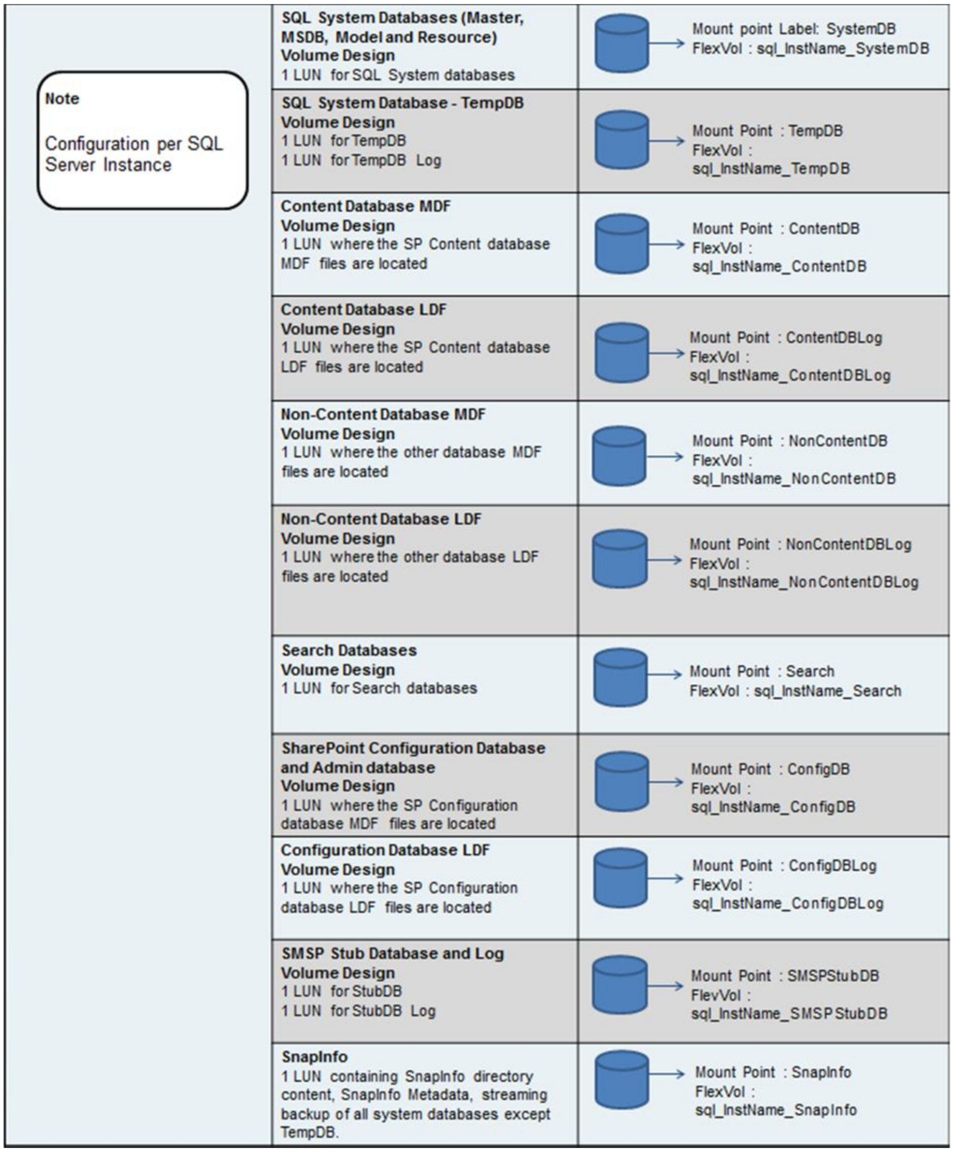

Note Configuration per SQL Server Instance	**SQL System Databases (Master, MSDB, Model and Resource)** **Volume Design** 1 LUN for SQL System databases	Mount point Label: SystemDB FlexVol : sql_InstName_SystemDB
	SQL System Database - TempDB **Volume Design** 1 LUN for TempDB 1 LUN for TempDB Log	Mount Point : TempDB FlexVol : sql_InstName_TempDB
	Content Database MDF **Volume Design** 1 LUN where the SP Content database MDF files are located	Mount Point : ContentDB FlexVol : sql_InstName_ContentDB
	Content Database LDF **Volume Design** 1 LUN where the SP Content database LDF files are located	Mount Point : ContentDBLog FlexVol : sql_InstName_ContentDBLog
	Non-Content Database MDF **Volume Design** 1 LUN where the other database MDF files are located	Mount Point : NonContentDB FlexVol : sql_InstName_NonContentDB
	Non-Content Database LDF **Volume Design** 1 LUN where the other database LDF files are located	Mount Point : NonContentDBLog FlexVol : sql_InstName_NonContentDBLog
	Search Databases **Volume Design** 1 LUN for Search databases	Mount Point : Search FlexVol : sql_InstName_Search
	SharePoint Configuration Database and Admin database **Volume Design** 1 LUN where the SP Configuration database MDF files are located	Mount Point : ConfigDB FlexVol : sql_InstName_ConfigDB
	Configuration Database LDF **Volume Design** 1 LUN where the SP Configuration database LDF files are located	Mount Point : ConfigDBLog FlexVol : sql_InstName_ConfigDBLog
	SMSP Stub Database and Log **Volume Design** 1 LUN for StubDB 1 LUN for StubDB Log	Mount Point : SMSPStubDB FlevVol : sql_InstName_SMSPStubDB
	SnapInfo 1 LUN containing SnapInfo directory content, SnapInfo Metadata, streaming backup of all system databases except TempDB.	Mount Point : SnapInfo FlexVol : sql_InstName_SnapInfo

If a database uses multiple volumes, then Snapshot copies are made sequentially. However, if it uses multiple LUNs on the same volume, then all Snapshot copies of these LUNs are made simultaneously, because Snapshot copies are volume-based.

In SharePoint 2013, databases created using the SharePoint central administration website use the model database, which has specific configuration settings, such as file location, growth settings, and more, as a template; it will apply these settings from the model database instead of getting created with default server-configured settings. By default, these databases will be placed on the same LUN as the SQL Server system databases. You need to manually migrate any newly created databases to their respective NetApp LUNs. For example, creating a new stub database for use with SMSP remote BLOB storage (RBS) provider will need to be migrated postcreation.

Best Practices

- Use Windows PowerShell and the SharePoint Windows PowerShell cmdlets to create databases (content and other) to avoid having a global unique identifier (GUID) in the name.
- Use FlexVol volumes to store SharePoint database files and avoid sharing volumes/datastores between different Windows host machines.
- Disable opportunistic locking (oplocks) on volumes hosting SMB shares where SharePoint data is stored to avoid corruption due to caching.
- For performance improvement for SQL Server used by the SharePoint farm:
  - Disable minimum read ahead (minra) on the volume where SharePoint data is stored.
  - Enable no updates of access times on inodes when a file is read (no_atime_update).
- Enable checking of NVRAM at controller boot (nvfail) to alert administrators to shut down databases if there is a problem with the NVRAM. This setting helps prevent SharePoint data corruption.
- Configure volume autosize policy, where appropriate, to help avoid out-of-space conditions.
- Make sure that Unicode is enabled on the SMB/CIFS volumes (create_ucode, convert_ucode).
- Make sure to use "NTFS" security style for the volumes used by LUNs and SMB shares.
- Make sure the SharePoint databases and the BLOB data reside on separate volumes.
- For SMSP Storage Manager, the I/O performance has no difference when using multiple farm/web app sharing one volume or each volume per farm/web app. This plays importance from backup/restore point of view; the BLOB storage volume should not be mixed between farms to make the backup data retention management easier. Also, if users plan to create multiple backup plans to group some web applications together, it is recommended to use one dedicated volume for the web applications in the backup plan, so BLOB backup/restore can be easier to manage with retention.

For complete details, refer to the SnapManager 8.0 for SharePoint Installation Guide.

4.3 LUNs

NetApp storage can be presented to Windows hosts as logical units called LUNs, which appear as local hard disks to the server. LUNs can be used in the same way as the local disks are used on the host. NetApp LUNs can be created using SnapDrive for Windows or Windows PowerShell, accessed by using the Fibre Channel (FC) or iSCSI protocols.

Table 1) Information on SMSP and SharePoint LUN layout.

Content	LUN	Description
SQL Server system databases	/vol/sql_Inst_Name_SystemDB/lunSQLSystemDB For master, model, and so on.	Place the SQL Server system databases on a dedicated volume, separate from the volume hosting the user databases.

Content	LUN	Description
	/vol/sql_Inst_Name_TempDB/lunTempDB /vol/sql_Inst_Name_TempDB/lunTempDBLog For optimal performance, separate the TempDB data and log files into separate LUNs within the TempDB volume.	Backed up using job scheduled through SMSQL directly and not SMSP. TempDB should not be included in a backup because the data it contains is temporary. Place tempdb on a LUN/SMB share that is in a storage system volume where Snapshot copies will not be created; otherwise, large amounts of valuable Snapshot space could be consumed.
SharePoint content databases	/vol/sql_Inst_Name_ContentDb/lunSPContentDB /vol/sql_Inst_Name_ContentDBLog/lunSPContentDBLog	These databases are backed up using SMSP. The layout of the content databases will be determined by the RTO of the databases. When you place multiple databases on the same LUN, the restore of individual databases will be done through the SnapDrive sub-LUN restore feature.
SharePoint configuration database	/vol/sql_Inst_Name_ConfigDB/lunSPCoreDBs /vol/sql_Inst_Name_ConfigDBLog/lunSPCoreDBLogs	These are not very read/write-intensive. Hence, you can also choose to: • Store the SharePoint central admin databases and service application databases. • Also, host the SMSP control and archive databases. These databases can be backed up using SMSP by adding the preceding databases as custom databases.
SMSP stub database	/vol/sql_Inst_Name_StubDb/lunSMSPStubDBs /vol/sql_Inst_Name_StubDb/lunSMSPStubDBLogs	SMSP stub database is highly read-write intensive in a collaboration environment; place the stub database and log on separate LUNs in its own volume. This way, you can host all the stub databases created per web application within the SharePoint farm.
SnapInfo	/vol/sql_Inst_Name_SnapInfo/lunSnapInfo	Used to store backup metadata for SMSQL. Make sure the databases residing on LUN/SMB shares within a volume are separate from that used by the SnapInfo volume to avoid stream-based backup and instead leverage the NetApp Snapshot technology.
Other databases	/vol/sql_Inst_Name_genDb/lunOtherDbs /vol/sql_Inst_Name_genLog/lunOtherDbLog	Databases for third-party-related apps or not related to SharePoint, but hosted on the SharePoint instance that can be backed up in SMSP using custom database option.
SharePoint search index	/vol/sql_Inst_Name_SPSearchIndex/lunSPSearchIndex	SMSP manages Snapshot copy of LUN using SDW.
SMSP 8.0 index (contains job metadata, SMSP backup index, WFE IIS metadata backup	/vol/sql_Inst_Name_SMSPIndex/lunSMSPIndex	SMSP manages Snapshot copy of LUN using SDW. The SMSP backup index data is not copied through Snapshot during a backup, hence requires an SDCLI script to take a Snapshot copy after the backup

Content	LUN	Description
data)		completes. If update SnapMirror option in the backup is selected, this volume does not have its mirror updated using SnapDrive. This can be placed on a NetApp LUN, VMDK, or CIFS volume. This location is configured using SMSP device manager.

Note: The preceding LUN names are provided as examples and can be replaced with business naming policies as necessary.

Note: SharePoint databases need to be migrated to NetApp storage system LUNs using the SMSP Migrate database tool in order to be backed up by SMSP. Refer to SnapManager 8.0 for Microsoft SharePoint Platform Backup and Restore User's Guide. SMSP checks the storage of BLOB data for connector and storage manager to be on NetApp storage (physical device on NetApp storage system CIFS share) to be able to be backed up using SMSP.

Note: Before LUNs are provisioned on a Storage Virtual Machine (SVM), you must add your storage system name and IP address to the Windows Server `etc\hosts` file to be able to resolve the virtual storage server management LIF IP address to the virtual storage server name.

Best Practices

- Use SnapDrive for Windows to create LUNs that are properly aligned on disk for the specific Windows file system that is used. Also, creating LUNs with SnapDrive for Windows makes sure that LUNs are of the correct type and with the correct disk offset, which is crucial for best disk performance.
- Do not create LUNs on the root aggregate of the node, root storage system volume /vol/vol0, or root volume of the SVM.
- Do not place user databases on the root LUN of the mount point.
- It is common to take transaction log backups more frequently than database backups, so place the transaction log on a separate LUN from the data files so independent backup schedules can be created for each. This also separates the random I/O of the data files from the sequential I/O to the log files and can improve performance of SQL Server used in the SharePoint farm.
- For clustered instances of SQL Server of the SharePoint farm:
  - The SnapInfo LUN must be a cluster disk resource in the same cluster group as the SQL Server instance being backed up through SMSQL.
  - Place SharePoint databases onto shared LUNs that are physical disk cluster resources assigned to the cluster group associated with the SQL Server instance.
- Make sure the database and SnapInfo LUNs/SMB shares are on separate volume to avoid retention policy from overwriting Snapshot copies when used with SnapVault.
- If the RPO of the content database demands a quick recovery, then that content database could have moved to a dedicated LUN.
- Make sure the DNS of the storage system is set up correctly for SnapDrive for Windows to be able to correctly resolve DNS name and allow creating and displaying LUNs/SMB shares and SMSQL to identify these LUN/SMB shares.
- Make sure to leave automatic Snapshot scheduling off as configured by SnapDrive for Windows.

4.4 SMB Shares

One of the major components added to clustered Data ONTAP 8.2 is support for the SMB 3.0 NAS protocol, a feature introduced with Windows Server 2012.

Best Practices

- Do not place the database files for the same SharePoint database on LUN and SMB shares.
- Configure SDW Transport Protocol Setting dialog with which SVM management LIF to connect (by providing SVM IP address, username, and password) to view all SMB shares on its CIFS server, which then becomes visible to SMSQL.
- The SMB share path used for database file paths has to be `\\<CIFS server name>\<share name>` for SMSQL to be able to recognize these database file paths as valid file paths hosted by NetApp storage.
- If you are running a failover cluster instance (FCI) system that only has SMB share and you want to use that FCI as the verification server, use any UNC path for default mount point directory to tell SMSQL that this is an FCI SMB-only configuration.
- Set the security style of the volume or qtree to NTFS to verify that permissions work correctly. Without this setting, Windows CIFS clients cannot access the shares correctly.
- Make sure no antivirus scanning is performed on the SMB/CIFS shares where SharePoint data is stored to avoid failed transactions due to latency.
- Make sure Windows host caching is disabled on the SMB/CIFS share where SharePoint data is stored to avoid corruption due to caching.
- Do not enable large maximum transmission unit (MTU) or “multicredit” support on SQL Server 2012. NetApp SMB/CIFS does not support SMB large MTU configurations.

5 Sizing for SnapManager for SharePoint

SharePoint Server 2013 Planning Considerations

Although SharePoint farms vary in complexity and size, a combination of careful planning and a phased deployment that includes ongoing testing and evaluation significantly reduces the risk of unexpected outcomes. Sizing is bound by capacity and performance, which will decide the number of disks and type of disks depending on required I/O. There are many factors that need to be considered when planning a SharePoint environment to size it correctly. Some of those factors include workload type, I/O operations per second (IOPS), requests per second (RPS), latency, read/write ratios, and working set size.

It is also important to have a well-thought-out information architecture (IA) and taxonomy, which will go a long way in helping SharePoint to be more discoverable, logical, and manageable. When you have good appreciation and understanding of capacity planning and management, you can apply your knowledge to system sizing. Sizing is the term used to describe the selection and configuration of appropriate data architecture, logical and physical topology, and hardware for a solution platform. There is a range of capacity management and usage considerations that affect how you should determine the most appropriate hardware and configuration options.

There are certain “by design” SharePoint limits that cannot be exceeded and some whose default values may be changed by the farm administrator. It is not a best practice to operate at or near an established limit because acceptable performance and reliability targets are best achieved when a farm’s design provides for a reasonable balance of limit values. For a comprehensive list, refer to Software boundaries and limits for SharePoint 2013.

In Microsoft SharePoint, the service architecture model provides a framework in which you deploy and manage services across a farm or across multiple farms. A service application represents a deployed instance of a service that you can configure and manage centrally and that many web applications can consume.

For additional information, refer to Plan service deployment in SharePoint 2013.

Table 2 details the databases that are created as part of the SharePoint deployment, based on the product version and edition. As part of sizing a SharePoint solution it is important to understand that each of the databases listed in Table 2 are not created equal. Each of the following databases has different requirements on:

- Location
- Growth factors
- Read/write characteristics
- Scaling strategy
- Recovery model

When planning a SharePoint solution, each of these factors plays into the decisions made on types of disks to use; for example, place TEMPDB on SSD or Flash Pool for performance versus BLOB content being placed on lower tier storage (SATA).

Storage and SQL Server capacity planning and configuration (SharePoint Server 2013) provides many of the details for calculating requirements. The most important details are that Microsoft recommends 2 IOPS per GB for optimal performance and at the low end 0.25 IOPS per GB. NetApp recommends always sizing based on 2 IOPS per GB so that there will always be IOPS available. Sizing to the minimal requirements in many cases will lead to poor performance and require a reactive versus proactive approach. The databases listed in Table 2 will not at all use 2 IOPS per GB, so for those individual databases IOPS can be gained for other critical databases.

Best Practices

- Do not enable auto-create statistics.
- Set max degree of parallelism (MAXDOP) to 1 for optimal query plans.
- Consider the following databases as Flash Pool candidates: TEMPDB, search, and usage (SharePoint 2010 only).
- Set new database files to a fixed number versus percentage.
- Always monitor database growth even with autogrowth set.
- Collaboration-oriented sites have a database priority as follows:
  - TEMPDB (mdf/ldf)
  - Content database (ldf)
  - Search databases
  - Content database (mdf)
- Read-oriented sites have a database priority as follows:
  - TEMPDB (mdf/ldf
  - Content database (mdf)
  - Search databases
  - Content database (ldf)

Table 2) SharePoint 2013 database details.

Product	Databases
SharePoint Server 2013	SharePoint system databases: • Configuration: SharePoint_Config • Central administration content: SharePoint_Admin_Content • Content (one or more): WSS_Content Service application databases:

Product	Databases
	• SharePoint search service application: – Search administration: Search_Service_Application_DB_<GUID> – Analytics reporting: Search_Service_Application_AnalyticsReportingStoreDB_<GUID> – Crawl: Search_Service_Application_CrawlStoreDB_<GUID> – Link: Search_Service_Application_LinkStoreDB_<GUID> • SharePoint user profile service databases: – Profile: User Profile Service Application_ProfileDB_<GUID> – Synchronization: User Profile Service Application_SyncDB_<GUID> – Social tagging: User Profile Service Application_SocialDB_<GUID> • App management: AppManagement • Secure store service: Secure_Store_Service_DB_<GUID> • Usage: SharePoint_Logging • Subscription settings service: SettingsServiceDB • Business data connectivity: Bdc_Service_DB_<GUID> Standard and enterprise editions: • Project Server 2013: ProjectWebApp • SQL Server PowerPivot service application: DefaultPowerPivotServiceApplicationDB_<GUID> • PerformancePoint services: PerformancePoint Service _<GUID> • State service: SessionStateService_<GUID> • Word automation services: WordAutomationServices_<GUID> • Managed metadata service: Managed Metadata Service Application_Metadata_<GUID> • Taxonomy: Managed Metadata Service_<GUID> • Machine translation services: SharePoint Translation Services_<GUID> • Apps for SharePoint: Apps_<GUID>
SharePoint Foundation 2013	SharePoint system databases: • Configuration: SharePoint_Config • Central administration content: SharePoint_Admin_Content • Content (one or more): WSS_Content Service application databases: • SharePoint search service application: – Search administration: Search_Service_Application_DB_<GUID> – Analytics reporting: Search_Service_Application_AnalyticsReportingStoreDB_<GUID> – Crawl: Search_Service_Application_CrawlStoreDB_<GUID> – Link: Search_Service_Application_LinkStoreDB_<GUID> • App management: AppManagement • Secure store service: Secure_Store_Service_DB_<GUID> • Usage: SharePoint_Logging • Subscription settings service: SettingsServiceDB

Product	Databases
	• Business data connectivity: Bdc_Service_DB_<GUID>The Business Data

To understand details about the databases for SharePoint 2010, refer to Database types and descriptions (SharePoint Foundation 2010).

For additional information, refer to:

- Database types and descriptions (SharePoint 2013)
- SharePoint Infrastructure planning and design process

The deployment of SharePoint Server components and objects on NetApp systems in general requires careful planning.

Sizing the Control Service Database

The control service has one control database, which contains the SMSP configuration data and backup plans, storage optimization (storage manager, archive manager, and connector) rules, and the job records. The data growth rate on control database is relatively small, and with retention on jobs, the job record can be automatically pruned from the control database.

Best Practices

- Always use retention policy with backup and archive plan so that the job record is pruned from control database.
  - Normally, the backup job data is automatically pruned by the data retention policy when you create a backup plan and also assign a retention policy.
  - If you have backup data for backup job not handled by retention policy, you need to manually delete it. (Navigate to job monitor to find the backup job and its ID > the logical device to find the directory for data of the backup job with the backup plan name and backup job name [ID], delete the folder > job monitor and delete the job monitor record.) Deleting the line in job monitor only deletes the record in SMSP database, so you need to manually delete the real data folder first.
- Change the recovery model for SMSP control database to full because it is changed frequently, causing the log size to grow.
- When using SQL 2012 availability groups, make sure all SharePoint and SMSP databases are set to full recovery model.
- It is highly recommended to configure a job-pruning policy if you are running backups frequently to make sure the control database is not overloaded with job data.

Sizing the Media Service Server

SnapManager for SharePoint provides a media service that manages the following:

- Backup job metadata
- Backup index
- IIS metadata backup
- SharePoint 15 or 14 hives files backup

The size of the backup job indexes created by the media service depends on the level of granularity chosen when creating the plan. In a normal SharePoint web application, as the level of granularity becomes finer, the number of objects that must be indexed increases, and therefore the size of the index increases. Normally, it is difficult to get a count of the number of objects at each level of granularity, which makes sizing the index very difficult. Installation of media service requires 1GB free space on the system drive that installs SMSP media service.

The media service cache is used for granular index generating buffer space. If the media service is a single node, set it as LUN (using the UNC path `\\host\driver$\path` will actually use the local disk access). In the case of using media service high availability, set a file share UNC path or shared disk path.

The data growth rate is small for SMSP platform because the backup data is written directly to physical device without using media service cache, and it is only for index creating temporary data usage.

Best Practices

- Reserve enough space for backup granular index data space if the backup plan has schedules. The media storage size needs to be enough to hold all data between the two retention cycles.
- NetApp recommends placing the media service on a dedicated physical host or virtual machine. This is necessary to cope with the additional processing power needed for managing the backup job data (metadata and index).
- For largely distributed deployments, it is recommended to have the media service deployed within close proximity of the web servers and physical storage, but not on the same hardware. Host the media service on hardware with high reliability in addition to high availability to prevent backups from being interrupted due to hardware failure.
- NetApp also recommends not installing the media service on WFE for security, monitoring, and scalability purposes.

Sizing for Media Service Server Datastore

Depending on the backup options in the backup plan, the following data and size can be estimated for the backup job (the SMSP backup will not save database content to media service):

1. The backup job will save a catalog file on media service, which is small and typically less than 10MB.
2. If the backup job is run with granular index enabled, the index data will be storage on media service. The size of index data is related to the number of items in the content database. Based on the test results:

   Est. index size (MB) = 1.5*(53*Nsc+ 27*Ns + 0.45*(D*1024*1024/S)/L + 0.35*(D*1024*1024/S))/1024

   Where:

   a. Nsc: Number of site collections in content database

   b. Ns: Number of subsites in content database

   c. L: Number of lists/folders in content database

   d. D: Total data size

   e. S: Average document size

	Nsc	Ns	L	D (GB)	S (KB)	Index DB Size (MB)
My sites	10	2,500	50	2,000	150	7,459.0
Intranet team sites	1	5,000	100	1,000	256	2,324.8
Department sites	5	500	250	1,000	1,024	547.9

Example calculation:

– Nsc = 20,000 site collections

– Ns = 100,000 sites

- L = 100 items per list
- D = 20000MB (20GB) of total data size
- S = 256KB for average document size

3. If the backup plan selected is to back up a WFE, the WFE data will be streamed to the media service LUN; each backup includes the SharePoint hive, global assembly cache, web parts, IIS metadata, and custom solutions. This size can vary depending how many custom SharePoint solutions are deployed from third-party ISV developers or internal development efforts. For example, a SharePoint 2013 farm with two WFE servers and one application server required 2.76GB on the media service LUN; in this example there were no customizations installed. The contents of these backups are stored as SQLite (sqlite.org).db files.

Estimation of Backup Data Size

1. The number of backup jobs data for one backup plan that needs to be placed on the media service LUN is determined by the backup retention policy defined in the storage policy. Because a NetApp volume has a maximum of 255 Snapshot copies per volume, it is important that a retention policy is set up to never reach this limit. The backup job saves a backup index on the media service LUN (for example: C:\SMSPCatalog\data_platform\Farm(SQL1#SHAREPOINT_CONFIG)\PLAN2013100115290038794 7\FB20131001152930531944), depending on the backup options in backup plan, and does not store the database content itself. If the backup job is run with granular index enabled, the index data will be stored on the media service LUN. The size of index data is related to the number of items in the SharePoint content database. For example, one document item takes about 1Kb for index, and if the documents have multiple versions, each version takes approximately 300 bytes.
2. If BLOB backup is selected in the backup plan, the index of BLOB is also saved to media service LUN. The WFEs are responsible for processing all RBS processes using the SMSP RBS provider installed on the different WFE servers. The SMSP RBS provider creates records in the stub database that correlates the content on SMB (CIFS) with the content contained within the SharePoint content database. The SharePoint content database only contains the RBS auxiliary table containing the BLOB ID; the stub DB contains the information of the RBS BLOB storage and how BLOB ID is mapped to the real BLOB storage location. The SMSP BLOB stub database keeps record of each BLOB, with each BLOB record using about 300 bytes.

Estimation of Archive Data Size

The media service is also used to save the archive data. Assuming the archive rule is created without compression enabled, the storage space used by archive data is basically same as the data size used in SharePoint. We can estimate the archive data storage size on media using the size of archive data in SharePoint plus 5% (for metadata and archive manager index usage).

Best Practices

- The BLOB stub database is set to simple recovery model to avoid big transaction log size. If user needs to change to full recovery model, make sure that the LUN has enough space for transaction log and shrink the log if it is necessary.
- The minimum unit for assigning stub database is content database. Keep one stub database per web application instead of one stub database for the entire SharePoint farm. This gives you the flexibility to have multiple backup plans for various web applications. Also, if BLOB data is part of backup plan, the restore will not overwrite the stub DB from a different backup plan.
- The SMSP agent uses stub DB for BLOB access, hence:
  - This stub DB should be close to the SharePoint WFE.
  - Have the stub DB in the same SQL Server instance as the SharePoint content database.
- BLOB stub database size might increase if there is large set of documents uploaded or a synchronization using Connector to connect NetApp SMB shares with large number of files.

6 Performance

Accurately sizing NetApp storage controllers for SharePoint workloads is essential for good performance. Consult a local NetApp SharePoint expert to provide accurate performance sizing along with capacity requirements in the previous section and layout for environments using SharePoint 2013/2010.

Best Practices

- Leverage the SMSP storage optimization modules to externalize BLOB data in order to help increase SQL Server performance by offloading write-intensive operations.
- An SMSP synchronization job is fairly resource-intensive; hence, running multiple synchronization jobs simultaneously might affect the performance of the server on which the control service is installed. To avoid this condition, configure SMSP processing pool where synchronization jobs that are added into the processing pool become threads. The number of jobs you allow in the processing pool is the maximum number of synchronization jobs that can be run simultaneously; the remaining jobs are placed in a queue.
- If the media service server is going to be virtualized, make sure to have good memory size and CPU power—that is, a heavy host configuration.

To maximize performance in SharePoint deployment with RBS, consider the following best practice.

Best Practice

When externalizing BLOBs to NetApp CIFS shares on less-expensive SATA disk storage as a result of using the SMSP storage optimization modules, add NetApp Flash Cache technology to improve controller performance for random read workloads.

7 NetApp Solution for SharePoint Server 2013

7.1 NetApp Storage Software and Tools

Microsoft SharePoint Server is a transaction-intensive application that can put heavy loads on storage systems. It is important to understand the NetApp technology and the value proposition that it offers to our customers and partners for a highly performing SharePoint deployment.

- **NetApp Host Utilities Kit.** Installation of this kit sets timeout and other operating system–specific values to their recommended defaults and includes utilities for examining LUNs provided by NetApp storage through the Fibre Channel, iSCSI, or FCoE protocol. It also helps to align the master boot record for the Microsoft VHD file layout, preventing it from getting out of alignment with the underlying NetApp LUN. This is very important for optimal I/O performance. For complete details for a given NetApp tested and supported configuration, refer to Interoperability Matrix Tool. You can download this software from the support site. This kit is not necessary on Windows Server 2012 when Data ONTAP DSM is installed.
- **Microsoft Windows and native multipath I/O (MPIO**). Microsoft MPIO is a Microsoft framework for developing multipath solutions that contain hardware-specific information required to enhance connectivity for storage arrays through iSCSI, Fibre Channel, or SAS. To operate as intended, clustered Data ONTAP requires MPIO to access volumes/LUNs and asymmetrical logical unit access (ALUA) to determine the state of a given path. When using the iSCSI protocol, it is necessary to configure multipath support on iSCSI devices in the MPIO Properties administrative application. Navigate to the Discover Multi-Paths tab, select the "Add support for iSCSI devices" checkbox, and click Add. It is also necessary to create multiple sessions from the host initiators to the target iSCSI LIFs on the clustered Data ONTAP system. This can be accomplished using the native iSCSI initiator: Select the "Enable multi-path" checkbox when logging on to a target. Sessions can also be managed by using the NetApp SnapDrive iSCSI management pane. This is the preferred method, because SnapDrive remembers which target logical interfaces already have an established session and preselects an unused target portal.
- **Data ONTAP DSM.** NetApp Data ONTAP DSMs for Windows MPIO help NetApp storage systems to integrate with Microsoft MPIO on Windows Server and provide high availability to applications by using path-failover methods. Use the Data ONTAP DSM over the native MPIO implementation by Windows 2008 and Windows 2012 to achieve optimal performance and advanced path decisions in the NetApp DSMs and also provide for auto-configuration, heuristics for specific storage arrays, statistical analysis, and integrated management. During installation, the Data ONTAP DSM sets a number of Windows registry values to optimize performance and provide correct behavior during the failover scenarios. Data ONTAP uses asymmetric logical unit access (ALUA) to identify optimized paths. There can be multiple active paths and multiple passive paths. You can have a maximum of 32 paths to a LUN. This maximum applies to any mix of FC and iSCSI paths. If all the active paths fail, the DSM automatically switches to the passive paths, maintaining the host's access to its storage.

  Consult the Interoperability Matrix Tool for current information on supported configurations.

  For further details, refer to Data ONTAP DSM 4.0.1 Installation and Administration Guide.
- **SnapDrive for Windows.** Key SnapDrive for Windows functionality includes SAN storage provisioning on the host, consistent data Snapshot copies, and rapid application data recovery from Snapshot copies. SnapDrive complements the native file system and volume manager technology, and it integrates seamlessly with the clustering technology supported by the host operating system to provide high availability of the service to its users.

  The unit of management when using NetApp SnapDrive with clustered Data ONTAP is at the level of individual SVMs, not at the node or cluster level. As virtual storage servers, SVMs are meant to provide a secure multi-tenancy environment in which specific volumes, network interfaces, target ports, and management users are logically grouped. Therefore, a host connected to multiple SVMs that are part of the same cluster, accessing the same physical disks and network interfaces, would not have visibility into the fact that they are logical entities distributed across physically separate, but interconnected nodes.

Best Practices

- Make sure the virtual server name is added in the domain naming system (DNS), which needs to be resolved to the management LIF.
- Make sure that the SVM details are provided using the Transport Protocol settings with IP address of the management LIF and management user for managing clustered Data ONTAP LUNs and SMB shares.
- Make sure the CIFS server name and the SVM names are not identical.
- When specifying a UNC path to a share of a volume when creating a LUN/SMB share, use IP addresses instead of host names. This is particularly important with iSCSI, because host-to-IP name resolution issues can interfere with the locating and mounting of iSCSI LUNs during the boot process.
- Depending on the level of security required, the administrator can decide whether to provide clustered Data ONTAP server credentials on the host side to enable Snapshot copy scheduling. Otherwise, the clustered Data ONTAP administrator (cluster administrator) should reserve adequate space for Snapshot copies while creating the SVM and aggregates for volume provisioning. This is necessary because the SVM administrator does not have sufficient privileges to add storage space to the aggregate.

For more information on SnapDrive for Windows, refer to the following guides:

- SnapDrive 7.0 for Windows
- TR-4000: SnapDrive 7.0 for Windows for Clustered Data ONTAP 8.2—Best Practices Guide for SAN Environment

- **SnapManager for SQL Server.** This application is tightly integrated with Microsoft SQL Server to help streamline database storage management while simplifying backup and restore operations for SQL Server databases. SMSP leverages SMSQL to back up and restore SharePoint databases.

Best Practices

- It is highly recommended that you install SnapDrive for Windows and SnapManager for SQL Server on all nodes of the always-on availability group.
- Do not schedule verifications on SQL Server of the SharePoint farm during peak usage hours. The verification process is CPU-intensive and will degrade SQL Server performance if run on SQL Server during peak usage hours.

For more information, refer to SnapManager for Microsoft SQL Server and Microsoft SQL Server and NetApp SnapManager for SQL Server on NetApp Storage Best Practices Guide.

8 SnapManager 8.0 for SharePoint

8.1 SnapManager 8.0 for SharePoint Overview

SnapManager for SharePoint is an enterprise-strength backup, recovery, and data management solution for SharePoint Foundation 2013 and SharePoint Server 2013, as well as SharePoint Foundation 2010 and SharePoint Server 2010 (all current and future service packs). Refer to SnapManager 8.0 for Microsoft SharePoint Platform Backup and Restore User's Guide.

The combination of NetApp storage solutions and Microsoft SharePoint enables the creation of enterprise-level database storage designs that can meet today’s most demanding application requirements. This release of SnapManager 8.0 for SharePoint (SMSP 8.0) aligns with Data ONTAP 8.2. The new features are:

- **Native SnapVault integration.** SMSP provides direct integration with SnapManager 7.0 for SQL Server (SMSQL 7.0) for recovering from vaulted backups for all restore workflows, including database-level restores and granular restores of BLOB data on external CIFS shares. SnapVault support in SMSQL 7.0 requires database (SAN/NAS) backups to be copied using SnapVault on secondary storage (archived backup). Unlike in 7-Mode, DataFabric® Manager (DFM) is not required for SnapVault datasets on clustered Data ONTAP. SnapDrive 7.0 for Windows (SDW 7.0) will provide native support for SnapVault for clustered Data ONTAP.

- **SharePoint cloning.** SMSP allows cloning of SharePoint objects such as web applications, sites, subsites, and libraries from a SharePoint topology tree view like the one that exists in the backup builder.
- **Granular functional RBAC in SMSP.** SharePoint administrators should be able to use the SMSP Account Manager module in control panel of SMSP for granular delegation of SMSP functions to users in the context of specific SMSP modules.
- **SMB 3.0 support.** You can choose to have SharePoint content databases reside either on NetApp LUN or on SMB 3.0 share with database file path that can be identified as UNC path.

Table 3) SnapManager 8.0 for SharePoint components mapped to SharePoint farm hosts.

Server Role	SMSP 8.0 Component	Remarks
Server that is not part of the SharePoint farm	SMSP manager	Mandatory.
WFE servers	SMSP agent and storage optimization modules such as Storage Manager, Connector, and Archive Manager	Optional. Enabled only for performing stub-based uploads of external documents for one or more of the web applications hosted on the WFE.
		Optional. Enabled only for archiving the contents of one or more of the web applications hosted on the WFE.
SharePoint index server	SMSP agent	Optional. Installed only for backing up the SharePoint search indexes.
SQL Server host	SMSP agent	Mandatory.

Note: Installing SMSP agent on a SQL Server that runs Windows Server Core is not supported by SMSP and SMSQL.

Best Practices

- It is recommended that you use the default settings unless a known conflict with an existing port persists. While customizing the ports, make sure that the ports you are using are available.
- If there are multiple SnapManager for SharePoint services installed on the same server, make sure all of the required ports are enabled on that server.
- Make sure that the antivirus on the host is not configured to block the ports being used by SMSP.

8.2 Backup Guidelines

During backup, SnapDrive is used to perform database Snapshot backups and SharePoint index Snapshot backups. The backup data of the other SharePoint components is sent to the configured storage policy and stored together with the backup job metadata and index.

Prerequisites

- Because the external BLOB data in the NetApp CIFS share is immutable, there is no need to perform frequent backups. NetApp recommends taking these backups daily or weekly, unless your disaster recovery SLAs demand more frequent backups.
- To make sure that Storage Manager and Connector can be used after disaster recovery, you must add their respective databases (stub and archive databases) as customized databases and include these databases in the full farm backup. In addition, you must use the same SnapManager for SharePoint control database when reinstalling the new SnapManager for SharePoint Manager; also, make sure all BLOB data is ready at the DR side before performing the disaster recovery.
- Storage of BLOB data for SMSP storage manager and connector MUST reside on NetApp storage because SMSP checks this when you leverage a physical device on NetApp storage system SMB share. SMSP requires that the SharePoint databases reside on the NetApp LUN to leverage SMSQL and NetApp Snapshot technology.
- SMSP requires the NetApp controller user account be able to log in to the Data ONTAP storage system and be able to perform the following operations:
  - Query/list SVM, CIFS shares, volumes
  - Create Snapshot copies for CIFS volumes
  - SnapMirror and SnapVault operations (query, update, and so on)
  - Expose Snapshot copy as CIFS share

Example for Backup Schedule

The customer wants the schedule as follows:

1. Hourly; between 08:00 and 18:00 (weekdays) with 14 days retention
2. Daily; with a retention of 12 weeks (every weekday)
3. Weekly; on Sunday with a retention of 24 weeks
4. Monthly; on Sunday with a retention of 52 weeks

You can define only one storage policy, and that is either daily, weekly, or standard.

You can choose to create one backup job with the following schedules:

1. Hourly; between 08:00 and 18:00, 7 days local retention and 14 days SnapVault retention = full farm backup (no item level) local management group standard
2. Daily on Sunday with “Standard” selected for SnapManager for SQL Server (SMSQL) management group and SnapVault job of daily retention of 12 weeks = item-level backup with index
3. Weekly on Sunday with “Standard” selected for SnapManager for SQL Server (SMSQL) management group and SnapVault job of weekly retention of 24 weeks item-level backup with index
4. Monthly on Sunday with “Standard” selected for SnapManager for SQL Server (SMSQL) management group and SnapVault job of monthly retention of 52 weeks = item-level backup with index

This way, you can keep only the seven days’ worth of Snapshot copies on local storage.

Best Practices

- Make sure the SMSP backup network share location is accessible by all nodes of the availability group.
- Keep the SharePoint databases (sync or async) in separate AGs based on their supported configurations as mentioned in Supported high availability and disaster recovery options for SharePoint databases (SharePoint 2013). Also, new databases added to the SharePoint farm need to be explicitly added to the availability group because this does not happen automatically.
- In SMSP, when backing up a SharePoint content database using a SQL Server FILESTREAM provider that is part of the availability group, make sure FILESTREAM is enabled on every server instance that hosts the availability replica for the availability group.
- In a mirrored setup, if SQL Server authentication is the SharePoint content database authentication, make sure the SMSP agent account has sufficient permissions to log into the destination SQL Server instance. Otherwise, the mirroring databases cannot be backed up when being used as failover databases.
- If you enable RBS and externalize the BLOB to NetApp CIFS share, the DB verification time at backup should be much shorter. The granular index generating time at backup should be the same because the number of entries has not changed. The SharePoint search crawl indexing time might be better because RBS access should be faster than SQL Server on large BLOB content.
- Make sure the recovery model for control database, archive database, and stub database has changed from simple to full.
- In the case of stub database migration, set the site collections to read-only so that no new BLOB record is created during this migration.

For additional information, refer to the SnapManager 8.0 for Microsoft SharePoint Platform Backup and Restore User's Guide.

8.3 Restore Guidelines

During restore, SnapManager for Microsoft SQL Server is used to perform restore from database Snapshot backups. In SMSP 8.0, item-level restore of BLOB from SnapVault is possible.

Best Practices

- SMSP supports the verification of backups on SnapMirror destinations and SnapVault secondary locations, thus offloading the read I/O from the production database servicing users.
- Restore BLOB data first, then restore content and stub databases. Also, disable the garbage collection (BLOB retention) before the database restore finishes.
- An out of place restore or restore to alternate location can be done to any SharePoint server provided it has the SMSP agent installed.
- To restore customizations successfully, NetApp recommends that you deploy the .wsp file for both the trusted and sandboxed solutions to the destination.
- Make sure the SharePoint file system resources (SharePoint hives) are restored prior to restoring the farm components.
- Make sure that the source node and the destination node are the same version and patch level for SharePoint. You can neither restore backed-up SharePoint 2010 data to SharePoint 2013 nor restore backed-up SharePoint 2013 data to SharePoint 2010. If the site within SharePoint 2013 is a SharePoint 2010 mode site, the content can only be restored to a SharePoint 2013 site that is in SharePoint 2010 mode.

Out of Place Restore

SMSP granular restore helps you recover backed up data to an alternate SharePoint location (out of place), either in the same farm or across different farms. The only requirement is to make sure that the SMSP agent is installed on the destination farm so that the SMSP can operate on the destination farm.

To restore to a location that does not exist in SharePoint, use the blank box at the bottom of each level in the SharePoint tree. To create a new site collection, the full URL is needed, but for other levels (from site down to folder), only the name needs to be specified. You should always select a container for the content that is either on the same level or one level higher than what is being restored. For example, a site should always be restored either to a site or to a site collection, and a list or library should always be restored either to a site or to another list or library.

Figure 2) Restore to alternate location screenshot.

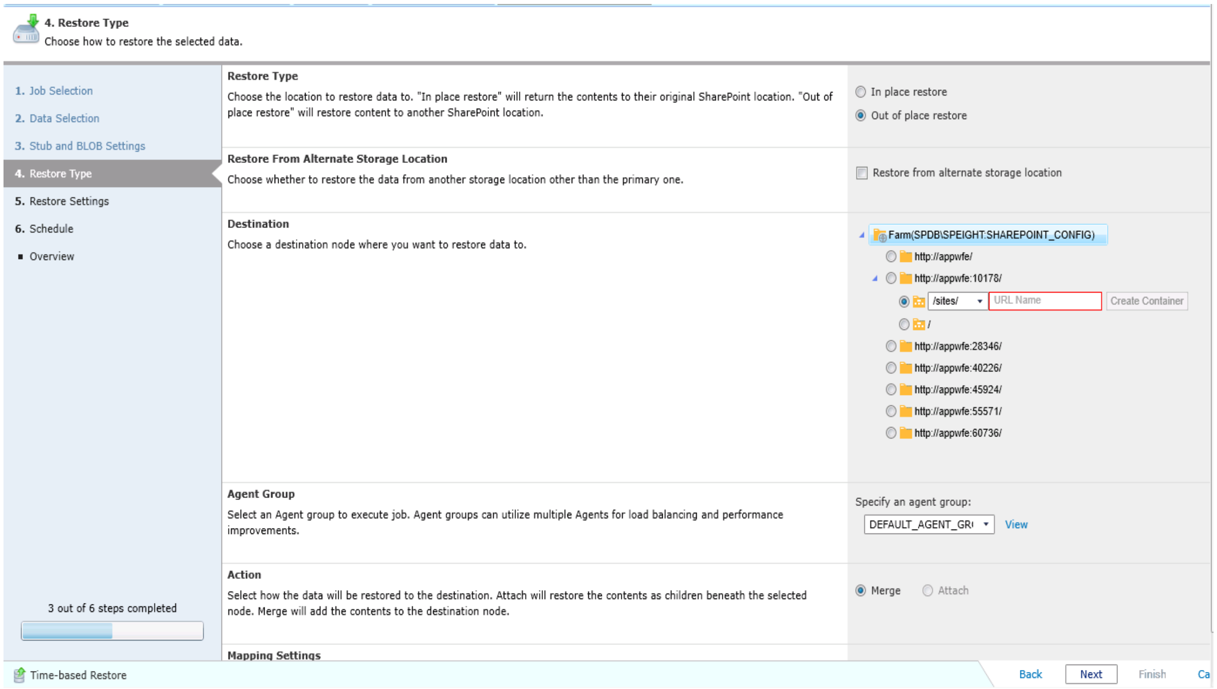

4. Restore Type

Choose how to restore the selected data.

1. Job Selection
2. Data Selection
3. Stub and BLOB Settings
4. Restore Type
5. Restore Settings
6. Schedule
- Overview

3 out of 6 steps completed

Restore Type

Choose the location to restore data to. "In place restore" will return the contents to their original SharePoint location. "Out of place restore" will restore content to another SharePoint location.

○ In place restore

● Out of place restore

Restore From Alternate Storage Location

Choose whether to restore the data from another storage location other than the primary one.

□ Restore from alternate storage location

Destination

Choose a destination node where you want to restore data to.

Farm(SPDB\SPEIGHT:SHAREPOINT_CONFIG)
- ○ http://appwfe/
- ○ http://appwfe:10178/
  - ● /sites/ URL Name Create Container
  - ○ /
- ○ http://appwfe:28346/
- ○ http://appwfe:40226/
- ○ http://appwfe:45924/
- ○ http://appwfe:55571/
- ○ http://appwfe:60736/

Agent Group

Select an Agent group to execute job. Agent groups can utilize multiple Agents for load balancing and performance improvements.

Specify an agent group: DEFAULT_AGENT_GR(View

Action

Select how the data will be restored to the destination. Attach will restore the contents as children beneath the selected node. Merge will add the contents to the destination node.

● Merge ○ Attach

Mapping Settings

To accomplish an out of place restore for content from SnapVault and/or SnapMirror destinations, the network must allow the source to connect to the destination over FC and/or iSCSI.

Best Practice

For SharePoint content, SMSP can do granular out of place restore to a different farm at a different granular level (site collection to item/item version level). However, SMSP does not support out of place restores of SharePoint services and components because SMSQL-based DB backup is based on local Snapshot, which might not be available to a SQL Server agent on other farm. For the same reason, SMSP does not support out of place database restore to a different SQL Server instance as well, even though this is possible through SMSQL.

For additional information, refer to [SnapManager 8.0 for Microsoft SharePoint Platform Backup and Restore User's Guide.]

8.4 Storage Optimization

Microsoft offers RBS as the official offloading technique for BLOB externalization, implemented by SQL Server, and is available in SharePoint 2013 and SharePoint 2010, based on the API supported by SQL Server 2012 and SQL Server 2008 R2. SMSP 8.0 includes storage optimization solutions to keep your

SQL Server resources optimized with intelligent archiving and BLOB offloading to a NetApp SMB share on tiered storage. Deduplication and compression enabled on NetApp storage will work on externalized BLOBs to provide improved I/O operation, deduplication, and/or compression. However, RBS does not increase the storage limits of content databases. The supported limits will still hold true for SP2013 databases. RBS has the following limitations:

- BLOB externalization granularity restore is limited to the site collection level, and for Connector, it is list level.
- RBS is not supported with SQL Server 2005.

  Note: RBS has to be run on the local computer that is running SQL Server 2012, SQL Server 2008 R2, or SQL Server 2008 R2 Express. SharePoint 2013 requires you to use the version of RBS that is included with the SQL Server remote BLOB store installation package from the feature pack for SQL Server 2012/2008 R2.

  Note: SnapManager for SharePoint supports only the SQL Server FILESTREAM provider for Microsoft SQL Server, provided this content resides on NetApp storage. Although be aware that this will only work if local RBS FILESTREAM provider is used; the remote RBS FILESTREAM provider will not be supported. SnapManager for SharePoint does not support other third-party RBS providers.

  Note: The connector uses RBS provider to represent file as BLOBs in SharePoint. Hence you cannot connect NetApp SMB shares on premises to SharePoint online or Office 365 because O365 does not allow RBS.

Best Practices

- Keep BLOB shares separate for each farm or SQL server instance for backup data retention management ease.
- Place the SMSP index on a separate LUN/SMB share from that of the BLOB share.
- Make sure to select which SMSP database, such as stub database, control database, and archive database, is explicitly be included in backup.
- The stub database that maintains pointers to BLOB content on the NetApp SMB share needs to be hosted on a SQL Server instance closest to SharePoint WFE, and this SQL Server instance must be available to all WFEs that run RBS provider and the SMSP manager.
- Use separate stub database per web application:
  - To be able to divide the web applications into multiple backup plans.
  - Also when BLOB backup is part of backup plan, the restore will not overwrite the stub data from a different backup plan.
- NetApp recommends that you create the RBS datastore on a NetApp LUN that does not contain the operating system, paging files, database data, log files, or the tempdb file.
- Although RBS can be used to store BLOB data externally, do not access or change those BLOBs manually outside of normal SharePoint operations.
- Make sure that the volumes for the NetApp SMB shares are big enough to hold a large amount of BLOB data.
- If you currently use SQL Server FILESTREAM and want to move to using SMSP and RBS, use the SMSP data import wizard to convert the FILESTREAM RBS BLOB to SMSP RBS BLOB. After BLOB converting, the FILESTREAM BLOB becomes orphaned; you will have to run the RBS garbage collection task outside of SMSP.
- If you choose to use SQL authentication when creating the archiver database, make sure the user specified also has db_creator and security admin database roles to SQL Server.

For additional information, refer to SnapManager 8.0 for Microsoft SharePoint Storage Optimization User's Guide.

8.5 High Availability

The best solution for high availability requires careful planning in terms of deciding whether to create fault-tolerant server hardware, create virtualization infrastructure, or increase the redundancy of roles for the SharePoint farm.

Control Service High Availability

SnapManager for SharePoint control service high availability can be achieved by installing SnapManager for SharePoint control service on multiple servers using the same control service database. HA is automatically performed by the Windows operating system within the same Windows network load-balanced cluster.

- You can choose to associate the SnapManager for SharePoint control database with a specific failover SQL Server instance that is used in conjunction with SQL Server database mirroring.
- Because the control DB for the control service is now in SQL Server, clustering and log shipping apply for HA.
- Make sure you register the agents and media service to the control service. In case this control service becomes unavailable, reregister the agents and media service to another control service in order to continue to access SnapManager for SharePoint.
- First control service installed is the master, which can be changed.
- Also make sure to configure a report location in Job Monitor before you can use the Log Manager and Job Monitor with SnapManager for SharePoint control service high availability. Otherwise, each server where control service is installed will retain its own log only for the jobs carried out by the control service installed on the server.

Make sure the following requirements are met:

- Enter the hostname or IP address of each individual server when installing SMSP control service on the corresponding server.
- Use the public IP address when installing other SMSP service(s).
- Use the public IP address when accessing SMSP.
- When using SQL authentication, make sure the specified account has DB owner permission of the existing SMSP control database or DB creator of the newly created control database.

Media Service High Availability

If you are using SnapManager 8.0 for SharePoint to manage your SharePoint farm, then media service plays a very important role. High availability of media service can be configured using Microsoft Windows cluster failover configuration for load-balanced access to the data storage locations; it requires all LUN/SMB physical devices have the same drive letter and mount point on all nodes. In cluster administrator, set all SMSP manager services as cluster generic services. Set control service or media service as a dependent on the shared drives. Use the media service server cluster name and IP address for any interaction with it.

Media service is only used to store the generated index and run postprocessing after the index has been generated and transferred by SMSP agent on SQL Server. When there are many SQL Server agents running index generation in parallel, configure the backup plans to use different storage policies that leverage multiple media servers to make sure the media servers used within a storage policy are polled sequentially.

All servers that belong to a server farm, including database servers, must physically reside in the same data center. Redundancy and failover between closely located data centers that are configured as a single farm ("stretched farm”) are not supported in SharePoint 2013. Refer to Hardware and Software Requirements for SharePoint 2013.

9 NetApp SnapVault

SnapVault is a NetApp disk-to-disk backup solution that is built into NetApp Data ONTAP. Enabling SnapVault on your NetApp system is as simple as installing a license key; no additional hardware or software must be installed. SnapVault allows you to replicate your data to a secondary volume and to retain the data for a longer period of time than you might on your primary volume.

SnapVault for clustered Data ONTAP is introduced in Data ONTAP 8.2.

Note: Both primary and secondary storage systems must be running clustered Data ONTAP 8.2 or later.

SnapManager 7.0 for SQL Server (SMSQL 7.0) provides native SnapVault support with clustered Data ONTAP 8.2. One important architectural change is that SnapVault in clustered Data ONTAP replicates at the volume level as opposed to the qtree level, as in 7-Mode SnapVault. This means that the source of a SnapVault relationship must be a volume, and that volume must replicate to its own volume on the SnapVault secondary.

Prerequisites

- Make sure valid FlexClone and CIFS licenses are installed on the SnapVault system when restoring from a SnapVault storage system.
- By default, the CIFS server is set to the same as the SVM name. Make sure the DNS name for the SVM and CIFS server is set up correctly. It should be different so it can be resolved correctly recovering BLOB data residing on NetApp SMB shares.
- There are four entries that need to be added to the DNS server or the SQL Server `etc\hosts` file:
  - Source SVM name
  - SnapVault destination SVM name
  - CIFS server name on source
  - CIFS server name on destination
- Make sure you add the SnapVault SVM management LIF IP address in the SDW 7.0 Transport Protocol settings on the primary storage system and vice versa.
- After making the preceding necessary changes, restart the SnapDrive service and the SnapDrive management service.
- If the secondary CIFS server is not in the same domain as the primary CIFS server, make sure a two-way trust relationship between the two domains exists.

Restore from SnapVault

The SnapVault backup data (or the remote backup data) is used only when the local backup has been deleted based on retention settings. The process of cloning database from SnapVault secondary to primary SQL Server with the name required by SMSP needs to be verified to be able to perform granular content restore.

Best Practice

Make sure the database and SnapInfo LUNs/SMB shares are on separate volumes to avoid SnapVault retention policy from overwriting Snapshot copies.

For additional information, refer to:

- SnapVault Best Practices Guide Clustered Data ONTAP
- SnapManager 7.0 for SQL Server (SMSQL) Installation and Administration Guide

10 SharePoint Disaster Recovery with SMSP

The organization's business requirements expressed using recovery time objective (RTO) and recovery point objective (RPO) are derived by determining the downtime cost to the organization if a disaster occurs and help build the SharePoint 2013 disaster recovery strategy. The best practice is to clearly identify and quantify your organization's RTO and RPO before developing the recovery strategy.

10.1 NetApp SnapMirror

NetApp SnapMirror maintains two copies of the SharePoint data online so that the data is available and is up to date at all times, even in the event of hardware outages, including a very unlikely triple disk failure. NetApp SnapMirror technology performs block-level mirroring of the SharePoint data volumes to the SnapMirror destination for data availability and to meet stringent RTO and RPO requirements. If a disaster occurs at a source site, mission-critical SharePoint data can be accessed from its mirror on the NetApp storage deployed at a remote facility for uninterrupted data availability. This approach can be tailored to meet your information availability requirements by providing a fast and flexible enterprise solution for mirroring data over LAN, WAN, and FC networks.

NetApp SnapMirror enables you to achieve the highest level of data availability with the NetApp active-active controller configuration. The client receives an acknowledgement only after each write operation is written to both primary and secondary storage systems. Therefore, the round trip time should be added to the latency of the application write operations.

Volume SnapMirror works at the physical level; therefore, any data that is compressed and deduplicated on the source retains the savings during the transfer and on the destination. This also reduces the network utilization between the source and destination by sending compressed/deduplicated data over the wire rather than the larger uncompressed/duplicate versions of data. Because the data remains compressed/deduplicated after the transfer, no additional load is imposed on the destination system by compression or deduplication.

Best Practice
NetApp recommends having adequate bandwidth over a WAN for the initial transfer.

For more information, refer to TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP 8.2.

11 Virtualization

Businesses of all sizes are virtualized and perform server consolidation across their application infrastructure to lower cost, improve scalability, and improve service-level agreements. SharePoint 2010 and 2013 as an application supports virtualization, so we can similarly virtualize the SnapManager for SharePoint components.

Best Practice
The SMSP manager service (control, media) servers can be virtualized. SharePoint WFE application server (AS) and SQL Server on which SMSP agent is installed can be virtualized. Make sure you have sufficient memory allocated for each VM, as defined for system requirements in the section “Preparing to Install SnapManager for SharePoint” of the SnapManager 8.0 for SharePoint Installation Guide.

During the planning of virtualization, it is necessary to evaluate and decide between the virtualization technology and the differentiating factors of multiple vendors, specifically Microsoft Hyper-V® or the VMware® ESX® virtualization stack.

11.1 Microsoft Hyper-V

SnapManager for SharePoint supports the Hyper-V feature introduced in Windows Server 2008 R2 and Windows Server 2012 through SnapDrive for Windows and enables users to provision LUNs to VMs and pass-through disks on a Hyper-V virtual machine without shutting down the virtual machine.

Best Practices

- Make sure that there is 4GB of RAM or more for virtual machines for the different SharePoint Server roles.
- To reduce disk contention, store system files on aggregates dedicated to storing virtual machine data. Keep the SharePoint content on a separate aggregate. This will make sure SharePoint I/O is separate from that of virtual machines.
- VHDs should only be created as thin fixed-type VHDs.
- SharePoint Server uses timer jobs extensively; hence use SnapManager for Hyper-V (SMHV) leveraging remote VSS for optimized backup of VMs, instead of the Hyper-V Snapshot feature because Snapshot copy latency adversely affects time-sensitive operations and can result in data corruption or data loss.
- NetApp recommends limiting the use of pass-through disks in Hyper-V except where considered necessary. This is because a limitation of pass-through disks is that Hyper-V Snapshot copies are not supported.

For best practices specific to Hyper-V, refer to TR-3702: NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

For additional information, refer to the following Microsoft TechNet links:

- Use best practice configurations for the SharePoint 2013 virtual machines and Hyper-V environment
- Best practices for virtualization (SharePoint Server 2010)
- Virtualization planning for on-premise or hosted technologies (SharePoint Server 2010)

11.2 VMware ESX

SMSP leverages SnapDrive for Windows to provide LUN provisioning and application-consistent backups and recovery leveraging NetApp storage array Snapshot copies for VMs hosted in a VMware vSphere® environment. The NetApp Virtual Storage Console (VSC), which is a server-side plug-in, needs to be installed on the vCenter™ system. Make sure the ports used by SMSP Manager (control and media services) are open on the guest OS VM. If the user plans to have the SMSP manager service VM OS disk on NetApp storage, the user is recommended to follow NetApp Storage Best Practices for VMware vSphere. Normally, the VMFS and NFS are used for OS VMDK storage for VMware HA, and the RDM/VMDK is only for data storage (content DB, search data, and media storage).

Best Practices

- Always use SnapManager for SharePoint to create consistent Snapshot copies of datastores.
- Use the NetApp VSC plug-in to create and manage datastores to host SharePoint data.
- It is a good practice to have fewer, but larger datastore volumes so that the time taken to mount a large number of such volumes decreases during the recovery. This might also translate to fewer protection groups on your setup.
- Have only FC-attached datastores/iSCSI-attached datastores in the same ESX or ESXi™ host or in different hosts in the same cluster. Do not mix them.
- Use the Data ONTAP PowerShell Toolkit (PSTK) to automate the test bubble (SRM replicated farm) and SDCLI.
- The RBS provider needs to be installed on all SharePoint WFEs to provide correct BLOB rendering of content to the end user.
- BLOB data needs to reside on the NetApp CIFS volume directly and not on VMDK on VMFS/NFS datastore.
- When using VMware HA, make sure the net share path used for SMSP job report location is accessible from the failover machine as well.

When using SQL Server 2012 availability groups, review Microsoft Clustering on VMware vSphere: Guidelines for Supported Configurations (1037959) for supportability.

12 Storage QoS and Nondisruptive Operations

12.1 Storage QoS

Data ONTAP 8.2 introduces storage quality of service (QoS), which can help you manage the risks that accompany meeting performance objectives for expensive SharePoint workloads. An entire SVM, or a group of volumes or LUNs/SMB shares within an SVM can be dynamically assigned to a policy group, which specifies a throughput limit defined in terms of IOPS or MB/sec. This can be used to reactively or proactively throttle workloads and prevent them from affecting the performance of other system workloads. QoS policy groups can also be used by service providers to prevent tenants from affecting each other, as well as to avoid performance degradation of the existing tenants when a new tenant is deployed on the shared infrastructure. Storage QoS is supported on clusters that have up to eight nodes.

Best Practices

- Enable QoS at the volume/LUN level with SharePoint content databases on SAN, depending on the storage layout to define isolation boundaries for collaboration and document publishing workloads, which are typically resource intensive with office client traffic and coauthoring.
- Enable QoS at the volume level when using SMB shares to define isolation boundaries for collaboration and document publishing workloads, which are typically resource intensive with office client traffic and coauthoring.
- Make sure your QoS policy is with least priority when applied to a cloned database.

For information about how to use storage QoS, see the Clustered Data ONTAP System Administration Guide for Cluster Administrators.

12.2 Nondisruptive Operations

Nondisruptive operation (NDO) is the continuous data availability to client and host requests architected with NetApp clustered Data ONTAP. The encompassing operations enable 24/7 activities, facilitated by a feature set built into the clustered Data ONTAP operating system.

A NetApp clustered Data ONTAP system is able to nondisruptively fail network connections over or relocate aggregates, while a SharePoint server is executing live read and write I/O to the volume that hosts SharePoint content databases. However, there are intricacies of the Server Message Block (SMB) protocol that prevent continuous data availability because the lock state information is not persistent. Such interruptions to data availability are direct results of the stateful protocol, and they have the same impact regardless of the underlying storage infrastructure.

You should be able to move a volume with SharePoint content databases and BLOB content to a different aggregate and fail back on demand, with no interruption to SharePoint workload data access.

Best Practices

To nondisruptively move SharePoint databases residing on SAN/NAS, consider the following when performing DataMotion™ migration:

- Should be performed to a supported version of Data ONTAP.
- Do not create or delete any LUNs inside volumes associated with the SVM after starting the migration process and before cutover completes. NetApp recommends performing all of the SVM resource and dataset-related change operations before the initial baseline transfer.
- NetApp recommends you refrain from using any backup and restore commands during the cutover or rollback phases.

- For more information on DataMotion for volumes, refer to TR-3975: DataMotion for Volumes Overview for Clustered Data ONTAP.
- For more information on best practices for DataMotion for volumes, refer to TR-4075: DataMotion for Volumes Best Practices and Optimization for Systems Operating in Cluster-Mode.
- For more information on NDO and SMB file shares, refer to TR-4100: Nondisruptive Operations and SMB File Shares for Clustered Data ONTAP.

References

The following references were used in this report:

- Capabilities and features in SharePoint 2013
- SharePoint 2010 Capabilities
- Hardware and software requirements for SharePoint 2013
- Technical diagrams for SharePoint 2013
- Plan for SharePoint 2013
- Overview of Shredded Storage in SharePoint 2013
- Software boundaries and limits for SharePoint 2013
- Capacity planning for SharePoint Server 2013
- Plan service deployment in SharePoint 2013
- Architecture design for SharePoint 2013 IT pros
- Database types and descriptions (SharePoint 2013)
- What's new in SharePoint 2013 upgrade
- Overview of the upgrade process to SharePoint 2013
- Verify database upgrades in SharePoint 2013.
- Introduction to Shredded Storage in SharePoint 2013
- "System error 2148073478," "extended error," or "Invalid Signature" error on SMB connections in Windows 8 or Windows Server 2012
- Restore web applications in SharePoint 2013

- Plan for high availability and disaster recovery for SharePoint 2013
- SnapDrive for Windows
- SnapManager for Microsoft SQL Server
- SnapManager for Microsoft SharePoint
- Data ONTAP Installation and Administration Guide
- TR-3483: Thin Provisioning in a NetApp SAN or IP SAN.
- SnapManager 8.0 for Microsoft SharePoint Installation Guide
- SnapManager 8.0 for Microsoft SharePoint Storage Optimization User's Guide
- SnapManager 8.0 for Microsoft SharePoint Platform Backup and Restore User's Guide
- Restore farms in SharePoint 2013
- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization
- Use best practice configurations for the SharePoint 2013 virtual machines and Hyper-V environment
- Deploying VMware vCenter Site Recovery Manager 5 with NetApp FAS V-Series S
- TR-3326: SnapMirror Sync and SnapMirror Semi-Sync
- SharePoint Community

Version History

Version	Date	Document Version History
Version 1.0	November 2013	Initial release

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

NetApp®

www.netapp.com