Technical Report

SnapManager 7.1 for SharePoint with Data ONTAP 7-Mode: Best Practices Guide

Cheryl George, NetApp
June 2013 | TR-4192

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 Data ONTAP® 7-Mode. It also covers the best practices for the NetApp enterprise data management solution for SharePoint, which is called SnapManager® 7.1 for SharePoint.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

1.1 Purpose and Scope

This document describes the best practices and offers insight into design considerations when deploying SharePoint Foundation 2013 and SharePoint Server 2013, as well as SharePoint Foundation 2010 and SharePoint Server 2010 (all current and future service packs) on NetApp storage systems running Data ONTAP 7-Mode, with the goal of achieving effective and efficient storage deployment planning and end-to-end data protection and retention planning. 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 Foundation 2013, SharePoint Server 2013, as well as SharePoint Foundation 2010 and SharePoint Server 2010 (all current and future service packs). 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 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 Data ONTAP 7-Mode operating system
- NetApp SnapDrive® data management software
- NetApp SnapManager for SQL Server® (SMSQL)
- NetApp SnapManager for SharePoint
- Microsoft SharePoint Server architecture and administration
- Microsoft SQL Server 2012, 2008 R2, and 2008

To make sure of configuration compatibility across the NetApp stack, refer to NetApp Interoperability Matrix Tool.

For the best practices for SnapManager 7.1 for SharePoint using clustered Data ONTAP storage systems, refer to TR-4193 SnapManager 7.1 for SharePoint with Data Clustered Data ONTAP - Best Practices Guide.

2 Storage Efficiency and Manageability

Storage efficiency is the ability to store and manage data in a way that consumes the least amount of space with little or no impact on the overall performance of the system. Storage efficiency goes beyond just data deduplication; it is a combination of RAID, provisioning (overall layout and utilization), mirroring, and other data protection technologies.

2.1 NetApp Technology for Storage Efficiency

The following NetApp technologies offer to implement storage efficiency and reap its 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.

Figure 1) NetApp technologies for storage efficiency.

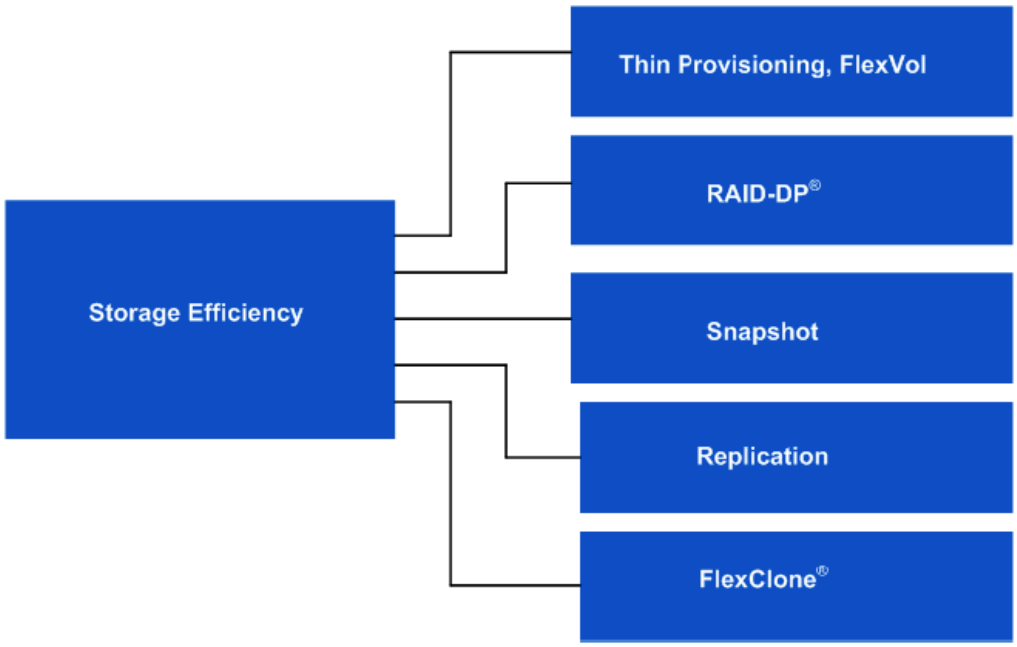

RAID-DP

RAID-DP®, the NetApp high-performance implementation of RAID 6, is double-parity RAID that adds a second parity stripe to dramatically increase data availability. With RAID-DP, aggregates and volumes can withstand up to two failed disks in a RAID group, or the more common event of one failed disk followed by an uncorrectable bit read error from the disk drive. RAID-DP makes less expensive SATA disks to store SharePoint content without worrying about data loss. It also lowers their storage acquisition costs.

RAID-DP provides considerable space savings when compared with the following:

- RAID 5 combines striping data across multiple disks with a data protection schema that generates parity bits as data is written, which means additional storage capacity is utilized.
- RAID 10, a combination of RAID 0 that stripes the data across multiple disks to maximize write speed by using multiple disks and minimizing overhead by offering no data protection capabilities.
- RAID 1, which provides one-to-one mirror copy of data by using mirror pairs of disks.

  Note: NetApp SyncMirror® can be used along with RAID-DP to provide a second layer of mirrored protection for a more robust disk protection strategy.

WAFL

Write Anywhere File Layout (WAFL®) is a block-based file system that uses inodes to describe files. It uses 4KB blocks with no fragments. Each WAFL inode contains 16 block pointers to indicate which blocks belong to the file. All the block pointers in a WAFL inode refer to blocks at the same level. Thus, inodes for files smaller than 64KB use the 16 block pointers to point to data blocks. Inodes for files larger than 64MB point to indirect blocks that point to actual file data. Inodes for larger files point to doubly indirect blocks. For very small files, data is stored in the inode itself in place of the block pointers.

Figure 2) WAFL file system.

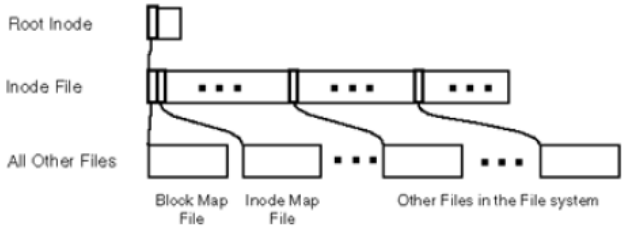

Flash Accel

Flash Accel™ is a NetApp server caching software solution. Flash Accel is fully supported software technology that can turn server-based PCI-e flash card or SSD drive into a server cache for Data ONTAP.

Best Practice
Consistent Snapshot™ copies cannot be created with write-back caching enabled without flushing the cache contents to disk first. If you do not flush the contents to disk first, you will likely have an invalid Snapshot copy or backup set. Hence, make sure write-back caching is disabled when using it with SnapManager for SharePoint and SnapManager for SQL Server.

Flash Cache

As information technology continues to grow in the business world, three fundamental resources are in short supply in many data centers: power, space, and cooling. A number of workloads no longer require as many disk drives to satisfy capacity, but the disks are still required for good performance. These additional disks require more space and more power, and the need for extra power increases the amount of cooling required. Flash Cache™, formerly known as Performance Acceleration Module (PAM), solves this dilemma by providing the ability to satisfy performance and capacity separately. Flash Cache can optimize the performance of random-read intensive workloads and caches data coming from volumes located in aggregates on the storage in the hardware module. This is much faster than retrieving it from the disk, which in turn accelerates the processing of the workload, thereby improving performance.

Flash Cache takes data that previously would have been cleared from system memory and places it in the module through the following process:

1. Data is always read from disk into memory.
2. When data must be cleared from the system memory, it is stored in the module.

Figure 3) Flash Cache process.

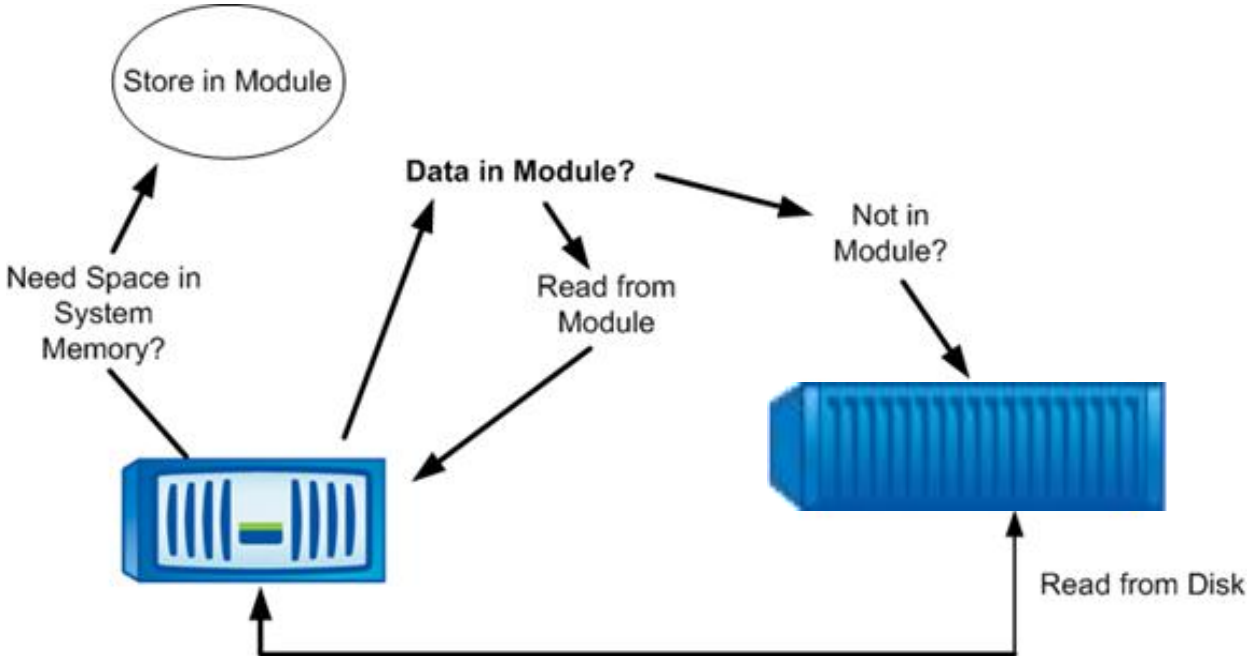

The next time the data is needed, Data ONTAP retrieves it from the module. This is much faster than retrieving it from the disk, which in turn accelerates the processing of the workload. 64-bit aggregates have a bigger address space and also take more memory for their metadata than 32-bit aggregates. This might reduce the total amount of effective data that can be cached in Flash Cache when it is used with the 64-bit aggregates present in the system.

Flash Pool

Flash Pool™ can optimize the performance of random-read and extensive write workloads and caches data coming from volumes located in aggregates on the storage.

Snapshot

NetApp Snapshot technology provides zero-cost, near-instantaneous point-in-time copies of the file system (volume) or LUN by preserving Data ONTAP WAFL consistency points. There is no performance penalty for creating Snapshot copies, because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not 100% for each backup, as with mirror copies. This means savings in storage costs for backup and restore purposes and it opens up a number of possibilities for efficient data management.

Figure 4) Example depicting how Snapshot copies work.

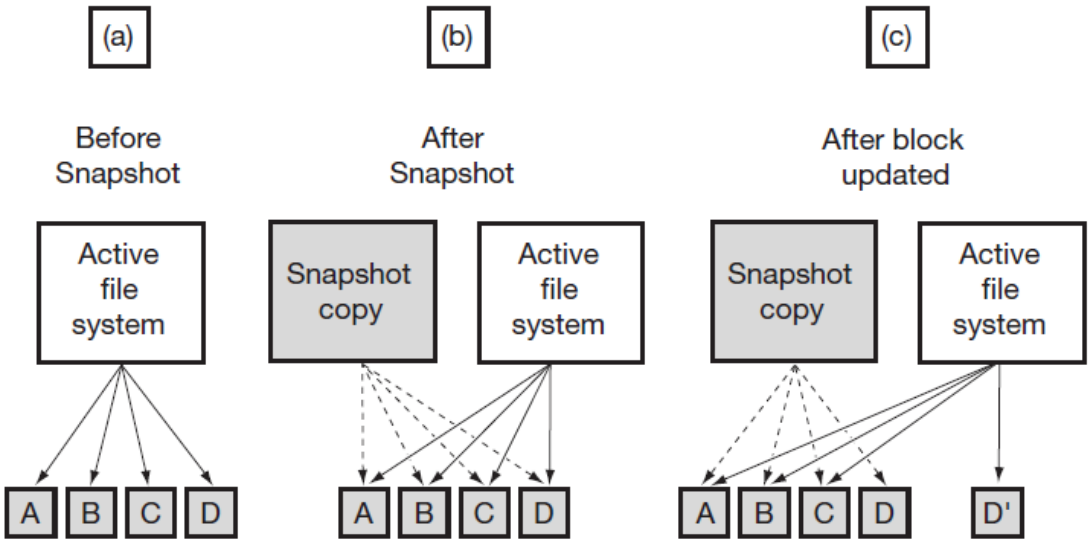

Best Practices

- Data ONTAP has a limit of 255 Snapshot copies per volume, so be careful about how many Snapshot copies are kept online. SMSQL 5.2 is integrated with NetApp SnapVault® through Protection Manager datasets. Therefore, older Snapshot copies should be archived according to their lifecycle and restore needs.
- 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 and SnapDrive for Windows®. If these schedules conflict, Snapshot copies may not be created.

Storage Thin Provisioning

In a shared storage environment, thin provisioning is a method of on-demand allocation of blocks of data, rather than the traditional method of allocating all the blocks up front, thereby optimizing utilization of available storage. Thin provisioning eliminates almost all whitespace, which helps avoid poor utilization rates. FlexVol® volumes (flexible volumes) is the enabling technology behind NetApp thin provisioning that can be thought of as the virtualization layer of Data ONTAP.

When thin provisioning is enabled, primary data and space for associated Snapshot copy are allocated on demand, to achieve the best ratio of storage efficiency for provisioning SharePoint. NetApp recommends that customers choose the thin-provisioning method to increase storage efficiency, which follows a 100% allocate-on-demand concept. When a LUN or volume is created, WAFL does not dedicate specific blocks out of the NetApp volume for the LUN or volume, or for Snapshot copies of the LUN or volume. Instead, it allocates the blocks from the NetApp aggregate when the data is actually written. This allows the administrator to provision more storage space than is actually physically present in the storage system.

Advantages of thin provisioning include:

- It requires less storage initially when buying a new storage system.
- More servers per storage system provide a greater level of consolidation.
- It allows overallocation of the application server's total disk capacity, thereby providing good return on investment (ROI).

- Monitoring storage space is very critical in thin-provisioned environments; otherwise, available storage space might become exhausted, which would cause application downtime. One way of protecting this type of scenario is to use the Snapshot autodelete and volume autogrow features.

The thin-provisioning method has the following characteristics:

- Volumes are created without space guarantee.
- The size of the volume follows the formula X + Δ, where:
  - X is the size of the primary data (sum of all user files and directories within the volume)
  - Δ is the amount of space needed to hold Snapshot copy data

Best Practices

- NetApp recommends using thin-provisioned LUNs for maximum storage efficiency. However, when enabling NetApp thin provisioning, administrators should also configure storage management policies on the volumes that contain the thin-provisioned LUNs. The use of these policies aids in providing the thin-provisioned LUNs with storage capacity as they require it. The policies include automatic sizing of a volume, automatic NetApp Snapshot copy deletion, and LUN fractional reserve.
- Thin provisioning is recommended for both SharePoint database and transaction log volumes.
- For optimal performance, NetApp strongly recommends that the thin-provisioned volumes be spanned across multiple disks. This substantially decreases the bottlenecks at the spindle level.
- 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 of the same size as or greater than the source volume.

Space Guarantee

Space guarantee enables thin provisioning. Space guarantee option can be set at the aggregate, volume, or LUN level, depending on the requirements of the application. 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 Reclamation

Space reclamation must be initiated from time to time to recover the unused space in a LUN. Storage space can be reclaimed at the storage level using the SnapDrive > Start Space Reclaimer option.

Autodelete

This volume setting allows Data ONTAP to automatically delete Snapshot copies when a threshold/trigger such as the following is met:

- **Volume.** The volume is nearly full, reported in the first line for each volume by the `df` command. Note that the volume can be full even though there might still be space in the snap_reserve areas.
- **Snap_reserve.** The snap_reserve space is nearly full.
- **Space_reserve**. The overwrite reserved space is full. This is the space determined by the LUNs with space reservations enabled and the fractional_reserve option. The reserve space will never be filled until both the volume and the snap_reserve areas are full.

Best Practices

- NetApp strongly recommends setting the trigger to volume.
- When autodelete is used, NetApp recommends that the Snapshot autodelete functionality in SnapManager for SQL Server (SMSQL) is used rather than the Data ONTAP autodelete feature. If not, SMSP will not delete the Snapshot copies based on the retention defined within the backup wizard of SMSP.
- 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 (SDW).
- NetApp recommends not using the autodelete option in NAS environments. Keeping a certain amount of space for Snapshot copies for file versioning or file restores is part of the SLAs defined for file services.
- NetApp recommends not enabling Snapshot autodelete for volumes that are currently protected by SnapManager for SharePoint. Enabling Snapshot autodelete on these volumes might disrupt SMSP and cause issues with consistency.
- 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 databases, NetApp recommends not to use 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 costly 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 to by using the -m option. If volume autosize is enabled, the default maximum size to grow to is 120% of the original volume size.

Best Practices

- NetApp recommends planning for additional buffer space when using thin provisioning for SharePoint 2013 environments. There must be enough space available in the aggregate for the autosize option to succeed.
- 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).
- Volume autosize is best practiced when deploying thinly provisioned environments.
- Space used for Snapshot copies can grow unexpectedly. Administrators can use the autosize function to make space available when reaching a certain volume threshold or when the space reserved for user data becomes low.
- 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 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 100%.

Best Practices

- Exercise caution when changing the fractional reserve value because when space is fully consumed, the write operations will fail and disrupt the SharePoint environment.
- Do not modify fractional reserve:
  - Unless there is a mechanism to monitor fractional reserve or volume and aggregate available space. SnapDrive for Windows does not provide this functionality.
  - If there are multiple LUNs in a volume and each LUN has a different rate of change, an estimation must be made of the overall volume size and the combined fractional reserve setting based on the average rate of change of all the LUNs.
- Use Snapshot autodelete and/or volume autosize when setting fractional reservation to a value less than 100%.

Table 1) Thin provisioning volume options.

Volume Option	Recommended Value	Notes
space-guarantee	none	No space reservation for volume at all. This is the key setting for thin provisioning.
fractional-reserve	0	The default is 0%.
autosize	on	Set autosize to `on`. No artificial limited volume must be monitored. The autosize function allows the user data to grow beyond the guaranteed space limit.
space-mgmt-try-first	volume_grow	Increasing the size of the volume does not destroy any data or information. Therefore, the volume size can be increased or decreased as needed. For some configurations, automatic volume growth might not be desired.

Table 2) Thin provisioning volume Snapshot options.

Volume Snapshot Option	Recommended Value	Notes
Snap_reserve	0	The value depends on the number of Snapshot copies and the change rate within the volume. Displaying only the committed usable space using SLA is the preferred way to provision storage. (However, in some situations, the Snapshot copy reserve area might be omitted.)
`autodelete`	true	Deleting Snapshot copies is not recommended in most environments so that data is not lost when SnapManager products are being used. However, if running out of space is deemed more critical than losing Snapshot copies, then change the value to `true` to enable this option. Having said that, Snapshot autodelete functionality in SnapManager for SQL Server is used rather than the Data ONTAP autodelete feature.

For best practice guidelines in thin-provisioned environments, refer to TR-3483: Thin Provisioning in a NetApp SAN or IP SAN.

Note: Snapshot copies used to create FlexClone® volumes are not deleted by the autodelete option.

NetApp FlexClone

A NetApp FlexClone volume is a writable point-in-time Snapshot copy of a parent FlexVol 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 useful in any situation where testing or development occurs, any situation where progress is made by locking in incremental improvements, and any situation where it is necessary to distribute data in changeable form without endangering the integrity of the original. A common scenario is to use FlexClone in an environment before committing a Microsoft SharePoint rollup or hotfix into production.

Figure 5) NetApp FlexClone technology.

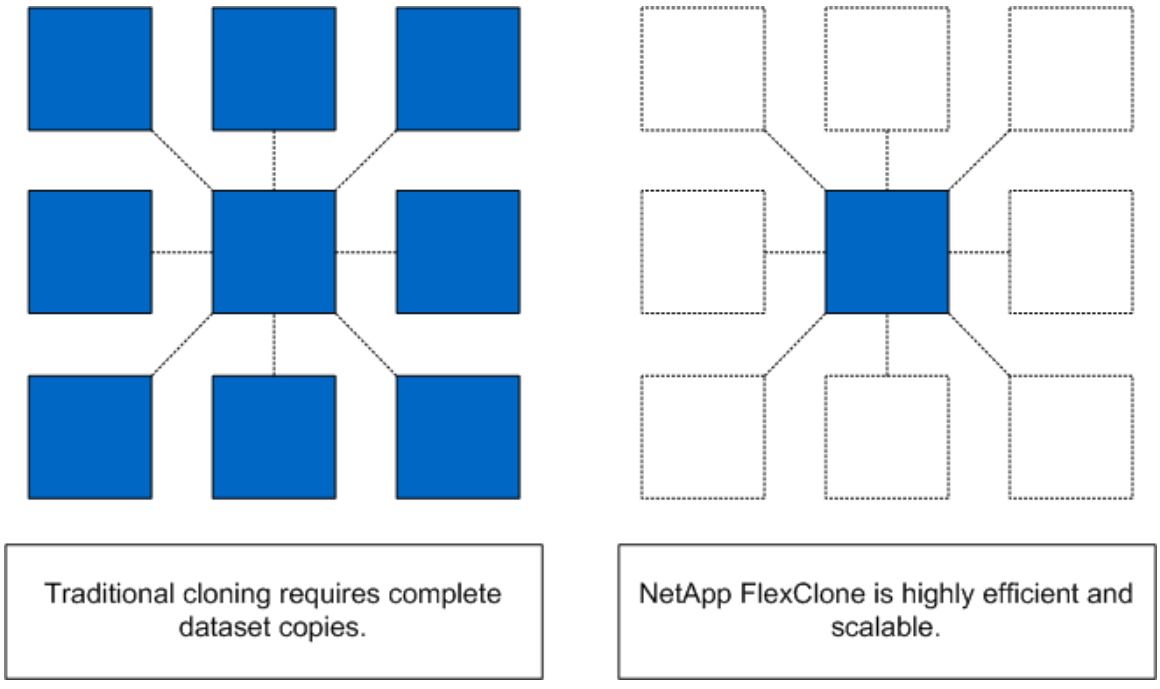

FlexClone technology can be leveraged both at the primary storage system and at the SnapMirror destinations for effective use of resources. FlexClone can also be used for disaster recovery testing without affecting the operational continuity of the SharePoint 2013 environment.

Best Practice

Use NetApp FlexClone to restore from SnapMirror or SnapVault content databases. Databases from the SnapMirror or SnapVault destination Snapshot copies can be cloned by using FlexClone and mounted on the active SQL Server node to go down to any item level and recover any items or the entire database itself. This is performed on active SQL Server instances. For clustered SQL Server instances in Windows 2008 and later, downtime of the SQL Server instance is not necessary. NetApp recommends creating FlexClone volumes by leveraging SnapDrive for Windows in SharePoint environments. This automates the creation of the FlexClone volumes and connecting the LUNs within the clone to the test and dev host.

NetApp Deduplication

NetApp deduplication is a fundamental component of NetApp core operating architecture Data ONTAP. It is a data compression technique for eliminating coarse-grained redundant data, typically to improve storage utilization. NetApp deduplication combines the benefits of granularity, performance, and resiliency with a significant advantage in the race to improve storage utilization demands. The deduplication process stores only unique blocks of data in the volume and creates additional metadata in the process.

When deduplication runs for the first time on a FlexVol volume with existing data, it scans the blocks in the volume and creates a fingerprint database, which contains a sorted list of all fingerprints for used blocks in the volume. 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.

Figure 6) NetApp deduplication.

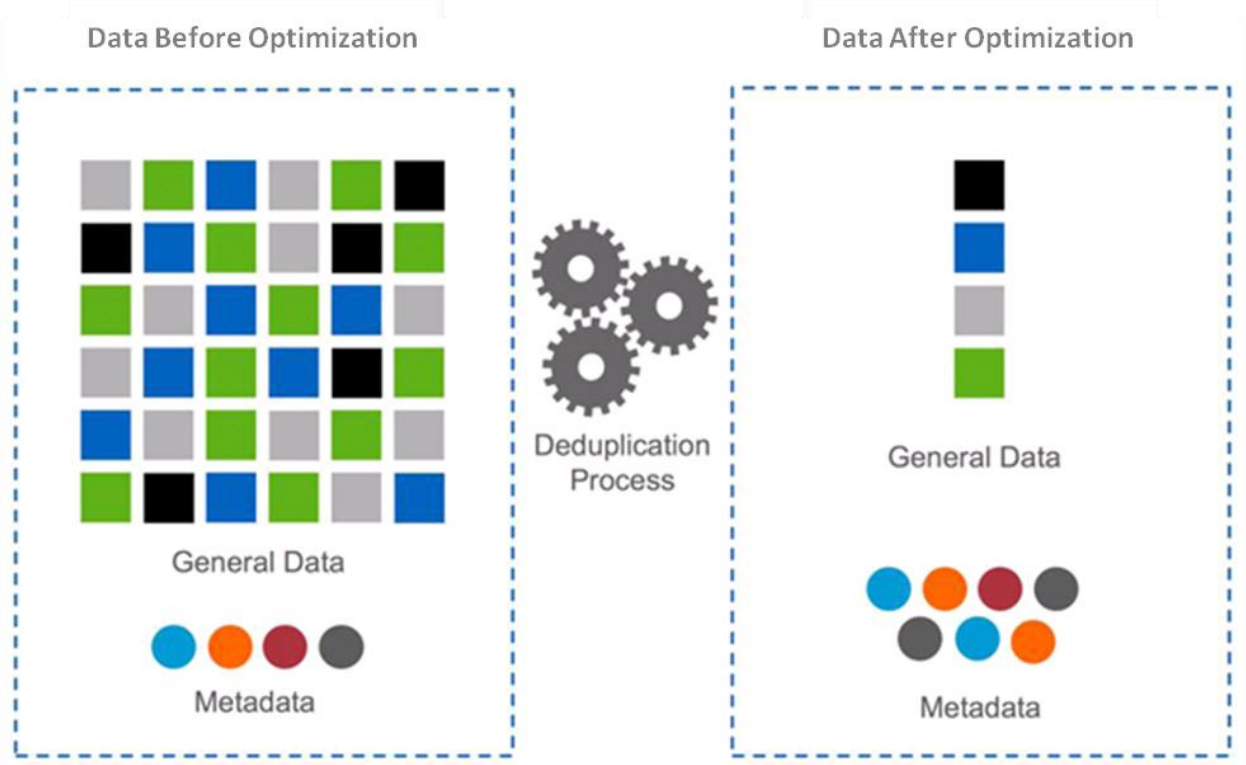

Deduplication consumes system resources and can alter the data layout on disk. Due to the application I/O pattern and the effect of deduplication on the data layout, the read and write I/O performance can vary.

Note: Deduplication is transparent to SharePoint, which does not recognize the block changes, so the SharePoint 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 proper testing should be done to determine the savings for your environment.

Note: Adequate space must be available on the FlexVol volume for the `sis on` command to complete successfully. If the `sis on` command is attempted on a FlexVol volume that already has data and is completely full, it fails because there is no space to create the required metadata.

Deduplication Metadata Overhead

Although deduplication offers substantial storage savings in many environments, there is a small amount of storage overhead associated with it that should be considered when sizing the FlexVol volume.

- **Volume deduplication overhead.** For each volume with deduplication enabled, up to 4% of the logical amount of data written to that volume is required to store volume deduplication metadata.
- **Aggregate deduplication overhead.** For each aggregate that contains any volumes with deduplication enabled, up to 3% of the logical amount of data contained in all of those volumes with deduplication enabled is required to store the aggregate deduplication metadata.

Total data = used space + saved space, as reported when using df –s (that is, the size of the data before it is deduplicated) from the Data ONTAP console. So, for 1TB of total data, the metadata overhead would be approximately 10GB to 60GB.

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. While

considering to use deduplication in SharePoint 2013 environments, be sure to consider this factor when sizing the volume layout.

Best Practices
• For volumes with deduplication enabled, volume autogrow is mandatory. • Run deduplication before creating new Snapshot copies. • Schedule deduplication only after significant new data has been written to the volume. • Configure appropriate reserve space for the Snapshot copies. • Replication of deduplicated volumes is supported by using SnapMirror. NetApp recommends not using deduplication with synchronous SnapMirror, because that could add substantial overhead on the storage subsystem and introduce performance overhead to SharePoint 2013 and SharePoint 2010 databases.

3 Virtualization

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

Best Practice
• You can choose to have all the components of the SMSP setup virtualized; make sure you have sufficient memory allocated for each VM.

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.

3.1 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. VHDs should only be created as thin fixed-type VHDs.

Best Practices
• Make sure that there is 4GB of RAM or more for virtual machines with SharePoint. • 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. • Do not use the Hyper-V Snapshot feature on virtual servers that are connected to a SharePoint products and technologies server farm because SharePoint Server uses timer jobs extensively; Snapshot latency adversely affects time-sensitive operations and can result in data corruption or data loss.

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)

3.2 VMware ESX

SMSP leverages SnapDrive for Windows to provide LUN provisioning and file system–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.

Note: NFS, VMFS, and RDM datastores are supported with Data ONTAP operating in 7-Mode.

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 a 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 Windows PowerShell™ Tool Kit (PSTK) to automate the network bubble (SRM replicated farm) and SDCLI.

Refer to Deploying VMware vCenter Site Recovery Manager 5 with NetApp FAS/V-Series Storage Systems for some keys to consider when implementing VMware vCenter™ Site Recovery Manager (SRM) version 5 in an environment using NetApp FAS storage systems.

4 Storage Planning for Microsoft SharePoint

4.1 NetApp Storage Software and Tools

Because the major part of the deployment of SharePoint is taxing on the storage, it is important to understand the underlying NetApp technology and the value proposition that it offers to our customers and partners.

- **NetApp Windows host utilities kit.** This kit contains the tools required to get support from the Microsoft iSCSI initiator. It can help configure the Windows Server to access virtual disk on a NetApp storage system 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.
- **Microsoft multipath I/O (MPIO).** Microsoft multipath I/O (MPIO) is a framework provided by Microsoft that allows storage providers to develop multipath solutions containing the hardware-specific information needed to optimize connectivity with their storage arrays, called device-specific modules (DSMs). Multipathing provides a high-availability solution for fault tolerance against a single point of failure in hardware components. Multipathing can also provide load balancing of I/O traffic, thereby improving system and application performance. MPIO is protocol-independent and can be used with Fibre Channel, Internet SCSI (iSCSI), and Serial-Attached SCSI (SAS) interfaces in Windows Server. Use NetApp DSM because it is easier to manage and path failover is faster as well.
- **SnapDrive for Windows.** SnapDrive simplifies storage and data management by using the host operating system and NetApp technologies, by hiding the complexity of steps that must be executed on both the storage system and the host system, and by removing the dependency on the storage administrator. NetApp SnapDrive for Windows helps a system administrator in:
  - LUN management:
    - Provision and manage storage directly from a server, managing LUNs on a storage system, making these LUNs available as local disks on Windows hosts.
    - Resize the storage on the fly without any disruption of the application.

- VSS integration with NetApp Snapshot technology.

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.

Best Practices

- Make sure that the NetApp storage system details are provided using the "Transport Protocol Settings" with preferred storage system IP address and appropriate protocol (HTTP or HTTPS) specified to successfully connect to the storage system.
- When specifying a UNC path to a share of a volume when creating a LUN, use IP addresses instead of host names. This is particularly important with iSCSI, as host-to-IP name resolution issues can interfere with the locating and mounting of iSCSI LUNs during the boot process.
- Create volumes and LUNs through SnapDrive for Windows instead of the OnCommand® System Manager.

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

- SnapDrive for Windows:
- SnapDrive 6.5 for Windows for Data ONTAP Operating in 7-Mode: Best Practices Guide
- SnapDrive 6. 5 for Windows for Clustered Data ONTAP: Best Practices Guide

SnapManager for SQL Server. This application is tightly integrated with Microsoft SQL Server to help streamline database storage management while simplifying storage layout planning, 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 and SnapManager for SQL Server on all nodes of the always on cluster.
- All backup and recovery processes should be created according to each environment's unique requirements. The general recommendation is to back up the SharePoint database every hour and the transaction log every 30 minutes during production time and possibly less often during off-peak hours.
- Do not schedule any SnapManager operations to overlap each other. Only one SnapManager operation can be running on the same machine at the same time.
- Do not schedule a backup to occur while a database verification is being performed, even if the verification is performed on a remote verification machine. This can result in a backup that cannot be deleted easily.
- Do not schedule verifications on SQL Server during peak usage hours. The verification process is CPU-intensive and could degrade SQL Server performance if run on SQL Server during peak usage hours.
- It is recommended that you keep a limited number of transaction logs to avoid running out of storage system disk space.

For more information, refer to:

- SnapManager for Microsoft SQL Server
- Microsoft SQL Server and NetApp SnapManager for SQL Server on NetApp Storage Best Practices Guide

SnapManager for SharePoint. This browser-based interface enables you to automate the data backup process and other administrative functions. It allows externalizing BLOB content to NetApp CIFS shares outside of SharePoint content database to improve the scalability of large SharePoint deployments.

Figure 7) Different components of SnapManager for SharePoint.

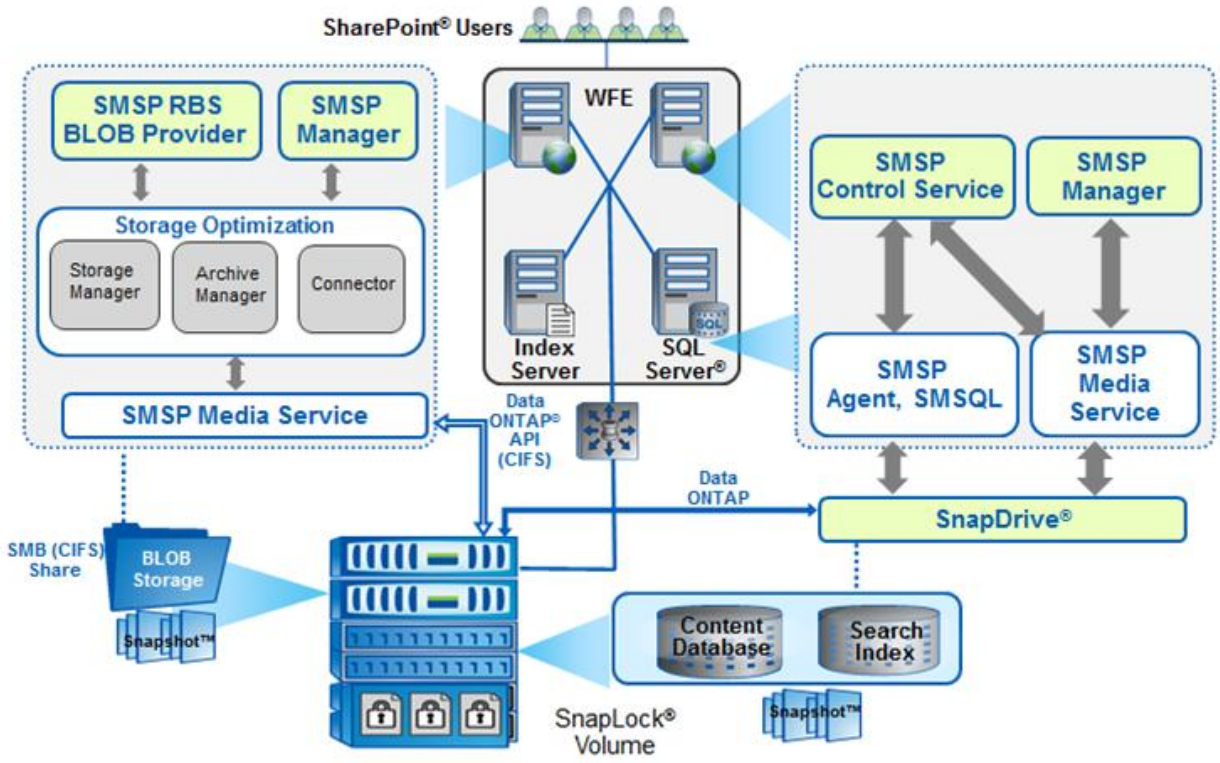

For more information, refer to SnapManager for Microsoft SharePoint.

4.2 Principles of Designing NetApp Storage for Microsoft SharePoint

Planning for Storage Layout (Aggregates, Volumes, and LUNs)

Storage must be carefully planned and laid out for the SharePoint content, considering the objectives of faster access and recoverability.

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, or item.

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.

Pooling all available disks into a single, large aggregate might maximize performance; however, it might not meet the data availability requirements. Creating separate aggregates for SharePoint databases and transaction log or SnapInfo volumes can meet the performance requirements of SharePoint Server while still providing the data availability required. However, hosting the database and the transaction logs or SnapInfo volumes in a single volume can have benefits from a storage efficiency perspective.

Best Practices

When you configure an aggregate, NetApp recommends that you use the following settings:

- **Double parity**. Select this option to benefit from RAID-DP, which is the preferred RAID level for an aggregate.
- **RAID group size**. NetApp recommends selecting the default, which is 16 in most cases.
- **Disk selection**. Automatic is selected by default and is the NetApp recommendation.
- **Disk size**. By default, any size is selected. However, NetApp recommends selecting disks of the same size when creating an aggregate.
- **Number of disks**. NetApp requires that at least three disks be assigned in order to provision a new aggregate. NetApp recommends creating the largest aggregate possible in order to benefit from the increased I/O capacity of all the spindles in the aggregate.
- NetApp recommends that you place all the SharePoint content on the same aggregate.

Volume Planning

Data ONTAP enables the creation of FlexVol volumes for managing data without the need to assign physical disks to the volumes. Instead, NetApp FlexVol volumes enjoy performance benefits from a larger pool of physical disks, called an aggregate.

This functionality results in the following additional benefits for SharePoint environments:

- A large number of volumes can be created, all with independent Snapshot copy schedules, mirroring policies, and so on.
- All volumes can be managed independently, while receiving the maximum I/O benefit of a much larger pool of disks.

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.

Figure 8) SMSP volume layout.

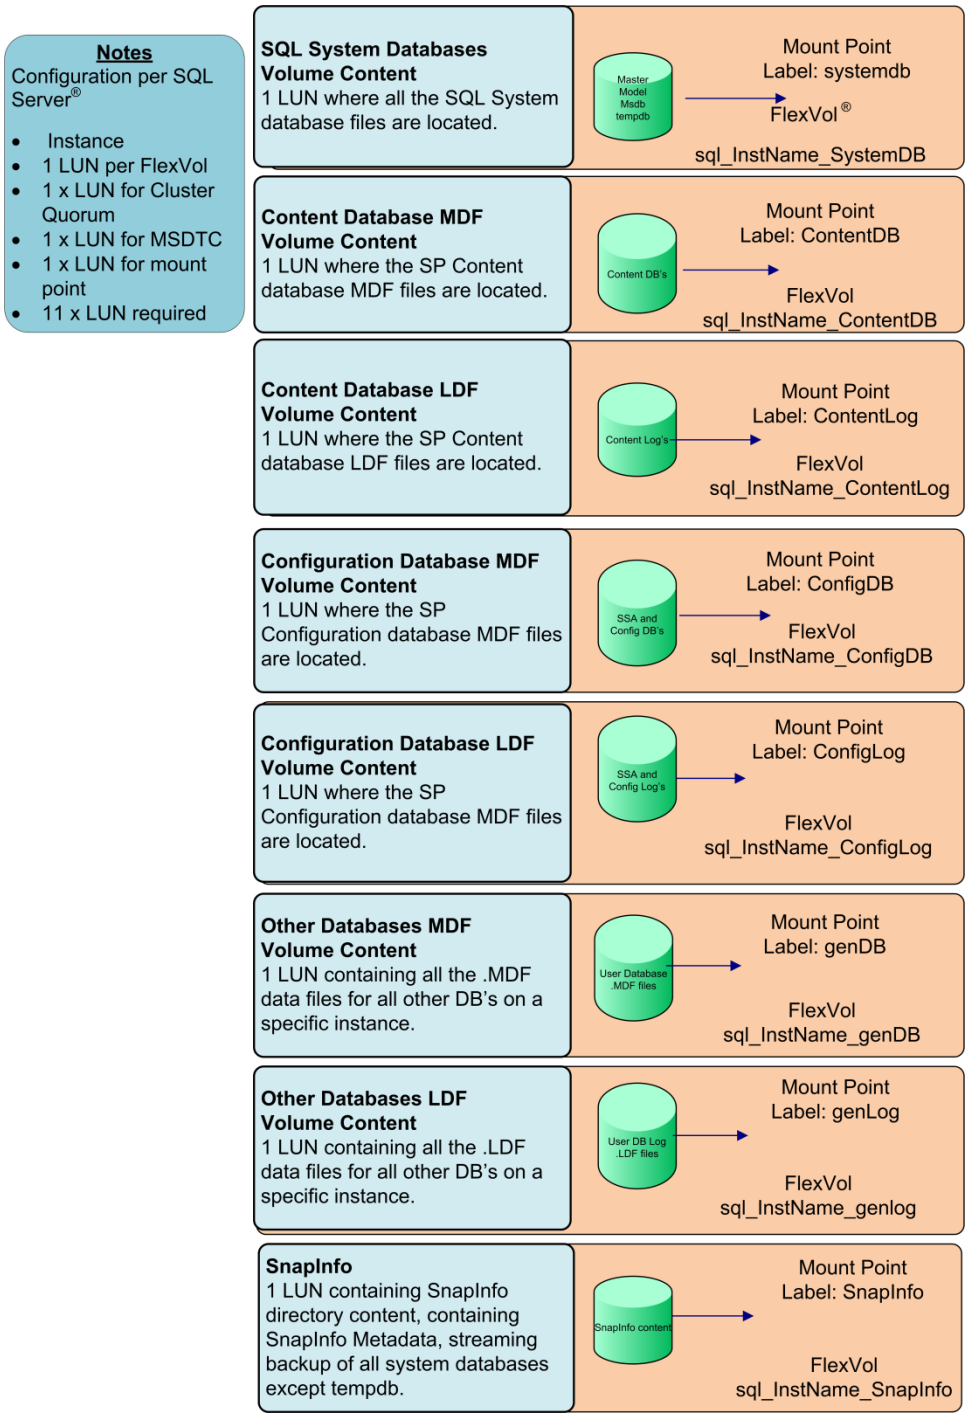

Table 3) Information on SMSP LUN.

Content	LUN	Description
SQL Server system databases	/vol/sql_Inst_Name_SystemDB/lunSQLSystem For master, model, and so on. /vol/sql_Inst_Name_SystemDB/lunTempDb For SQL temp database	Backed up using job scheduled through SMSQL directly and not SMSP.
SharePoint content databases	/vol/sql_Inst_Name_ContentDb/lunSPContentDb /vol/sql_Inst_Name_ContentLog/lunSPContentLog	These databases are backed up using SMSP. The layout of the content databases will be determined by the RTO of the databases. When you put multiple databases on the same LUN, the database restore method is a streaming restore. For larger databases, this might result in missing RTO times because streaming restores are far slower than restoring at the LUN level.
SharePoint configuration database	/vol/sql_Inst_Name_ConfigDb/lunSPCoreDbs /vol/sql_Inst_Name_ConfigLog/lunSPCoreDbLogs	These are not very read or write intensive. Hence you can also choose to: • Store the SharePoint central admin databases and service application databases. • Also host the SMSP control, SMSP stub, and archive databases. These databases can be backed up using SMSP by adding the aforementioned databases as custom databases.
Other databases	/vol/sql_Inst_Name_genDb/lunOtherDbs /vol/sql_Inst_Name_genLog/lunOtherDbLog	Databases for third-party related apps or not related to the SharePoint, but hosted on the SharePoint instance that can be backed up in SMSP using custom database option.
SnapInfo	/vol/ sql_Inst_Name_SnapInfo/lunSnapInfo	Used to store backup metadata for SMSQL.
SharePoint search index	/vol/ sql_Inst_Name_SPSearchIndex/lunSPSearchIndex	SMSP manages Snapshot copy of LUN using SDW.
SMSP 7.x index (contains job metadata, SMSP 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

Content	LUN	Description
index, WFE IIS metadata backup data)		requires an SDCLI script to take a Snapshot copy after the backup 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: SharePoint databases need to be migrated to NetApp storage system LUNs so that they are backed up using SMSP. 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.

In SharePoint 2013, databases created using SharePoint central administration website uses the model database as a template, which has specific configuration settings, such as file location, growth settings and more, it will apply these settings from the model database instead of getting created with default server configured settings.

Best Practices

- Do not create LUNs on the root storage system volume /vol/vol0.
- It is strongly recommended that you put fewer than 35 SharePoint databases on a storage volume.
- For very large SharePoint farms having multiple SQL Server instances on the back end reduces backup windows and also reduces contention for SMSQL. To attain the optimal layout for SQL Server databases with NetApp storage systems, have separate FlexVol volumes or LUNs for each SQL Server instance.
- NetApp recommends not performing a restore through SMSP when an SMSP backup is going on and vice versa.
- Each web application should be allocated a dedicated volume for all its content databases. Storing multiple databases in a single LUN or volume results in fewer mount points and faster backups due to fewer backup groups, because a fewer volumes are used for all the databases. However, make sure that the SharePoint databases are on LUNs separate from the SQL Server system database LUN.
- LUNs should be created on the faster disks, such as Fibre Channel disks or SAS disks with 15kRPM.
- Restores tend to be faster when entire databases are being restored provided a single database is hosted on a LUN or all the databases on a LUN are restored together. This is because SMSP can take advantage of LUN clone split restore for much faster restores. Individual documents are restored leveraging cloning technology.
- Do not have SharePoint content LUNs on the same storage system volume as CIFS data.
- Consolidation of multiple LUNs on a single volume can ease monitoring and administrative activities at the storage level.

4.3 Sizing for SnapManager for SharePoint

SharePoint Server Planning Considerations

Although SharePoint products 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. It is also important to have a well thought out information architecture (IA) and taxonomy that 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.

In SharePoint Server, there are certain limits that are by design and cannot be exceeded, and other limits that are set to default values that may be changed by the farm administrator. Therefore, exceeding certain limits, such as the number of site collections per web application, may only result in a fractional decrease in farm performance. However, in most cases, operating at or near an established limit is not a best practice, as acceptable performance and reliability targets are best achieved when a farm's design provides for a reasonable balance of limits 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 4 details the databases that are created as part of the SharePoint deployment, based on the product version and edition.

Table 4) SharePoint 2013 database details.

Product	Databases
SharePoint Server 2013	SharePoint system databases: • Configuration: SharePoint_ConfigCentral Administration content: SharePoint_AdminContent_<GUID> • 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> • 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>

Product	Databases
	• App management: App_Management_<GUID> • 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_AdminContent_<GUID> • 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: App_Management_<GUID> • Secure store service: Secure_Store_Service_DB_<GUID> • Usage: SharePoint_Logging • Subscription settings service: SettingsServiceDB • Business data connectivity: Bdc_Service_DB_<GUID>The Business Data

SharePoint System Databases

Configuration: SharePoint_Config

The configuration database contains data about SharePoint databases, Internet Information Services (IIS) websites, web applications, trusted solutions, web part packages, site templates, and web application and farm settings specific to SharePoint products, such as default quota settings and blocked file types.

Best Practices
• The recommended minimum size of configuration database is 2GB. Database size may increase with time, growing by approximately 40MB for each 50,000 site collections. Transaction logs for the configuration database can grow large if many items are created between transaction log checkpoints. • The default recovery model for the SharePoint configuration database is full.

Central administration content: SharePoint_AdminContent_<GUID>

The central administration content database is considered to be a configuration database. It stores all site content, including site documents or files in document libraries, list data, and web part properties, in addition to user names and rights for the central administration site collection. If Microsoft SQL Server PowerPivot for Microsoft SharePoint is installed, the central administration content database also stores the Excel® worksheets and PowerPivot data files used in the PowerPivot Management Dashboard.

Best Practices
• The recommended minimum size is 1GB. • The default recovery model for the SharePoint admin content database is full.

Content (one or more): WSS_Content

Content databases store all content for a site collection, including site documents or files in document libraries, list data, web part properties, audit logs, and sandboxed solutions, in addition to user names and rights. All the data for a specific site collection resides in one content database on only one server. A content database can be associated with more than one site collection. Content databases also contain the Microsoft Office Web Apps cache, if Office Web Apps have been deployed. Only one cache is created per web application. If multiple site collections that are stored in different content databases have Office Web Apps activated, they will all use the same cache. You can configure the size of cache, the expiration period, and the location. For more information about the size of the Office Web Apps cache, see Manage the Office Web Apps cache. Content databases also store user data for PowerPivot for SharePoint, if it has been installed in the environment.

Note: For SharePoint sizing, contact your NetApp SharePoint consulting systems engineer for assistance for sizing and layout.

Best Practice
• The default recovery model for the SharePoint content database is full.

Service Application Databases

SharePoint Search Service Application

- Search administration: Search_Service_Application_DB_<GUID>

The search administration database hosts the search service application configuration and access control list (ACL) for the crawl component.

Best Practices
• The recommendation is to allocate 10GB of space. • The default recovery model for the SharePoint search service application database is simple.

- Analytics reporting: Search_Service_Application_AnalyticsReportingStoreDB_<GUID>

The analytics reporting database stores the results for usage analysis reports and extracts information from the link database when needed.

- Crawl: Search_Service_Application_CrawlStoreDB_<GUID>

The crawl database stores the state of the crawled data and the crawl history.

Database size = sum of all content databases in the farm × 0.046.

- Link: Search_Service_Application_LinkStoreDB_<GUID>

The link database stores the information that is extracted by the content processing component and the click through information.

For SharePoint search service application, SMSP focuses on the topology of index partition; refer to the following contents to understand the best practices of SharePoint search partition index. NetApp recommends adding one index partition for every 10 million items in the search index.

Table 5 lists the examples of a medium enterprise search farm with approximately 40 million items in the search index.

If server resource is enough for full fault-tolerance and performance, distribute the index partitions of different servers to different volumes.

Table 5) Example of medium enterprise search farm.

Server Name	Index Partition			
	0	1	2	3
Host A	✓	✓		
Host B	✓	✓		
Host C			✓	✓
Host D			✓	✓

If server resource is not enough for full fault-tolerance, configure index replicas, and distribute the search index files on one server with index partitions in the same volume.

Server Name	Index Partition			
	0	1	2	3
Host A	✓	✓	✓	✓
Host B	✓	✓	✓	✓

If server resource is not enough for performance, distribute the index partitions in the same volume.

Server Name	Index Partition			
	0	1	2	3
Host A	✓	✓		
Host B			✓	✓

Note: To use volume with SnapMirror protection, select “Update SnapMirror after operation” option in SMSP when backing up search service application.

SharePoint Search Index Data Migration

Best Practices

- SMSP 7.1 does not support index migration from third-party storage to NetApp storage; hence, user is required to manually migrate SharePoint 2013 index.
- Before SharePoint search index migration:
  - Verify that the user account that is performing this procedure is an administrator for the search service application.
  - Verify that the search index is empty in the Search Administration page > System Status > Searchable items displays "0."

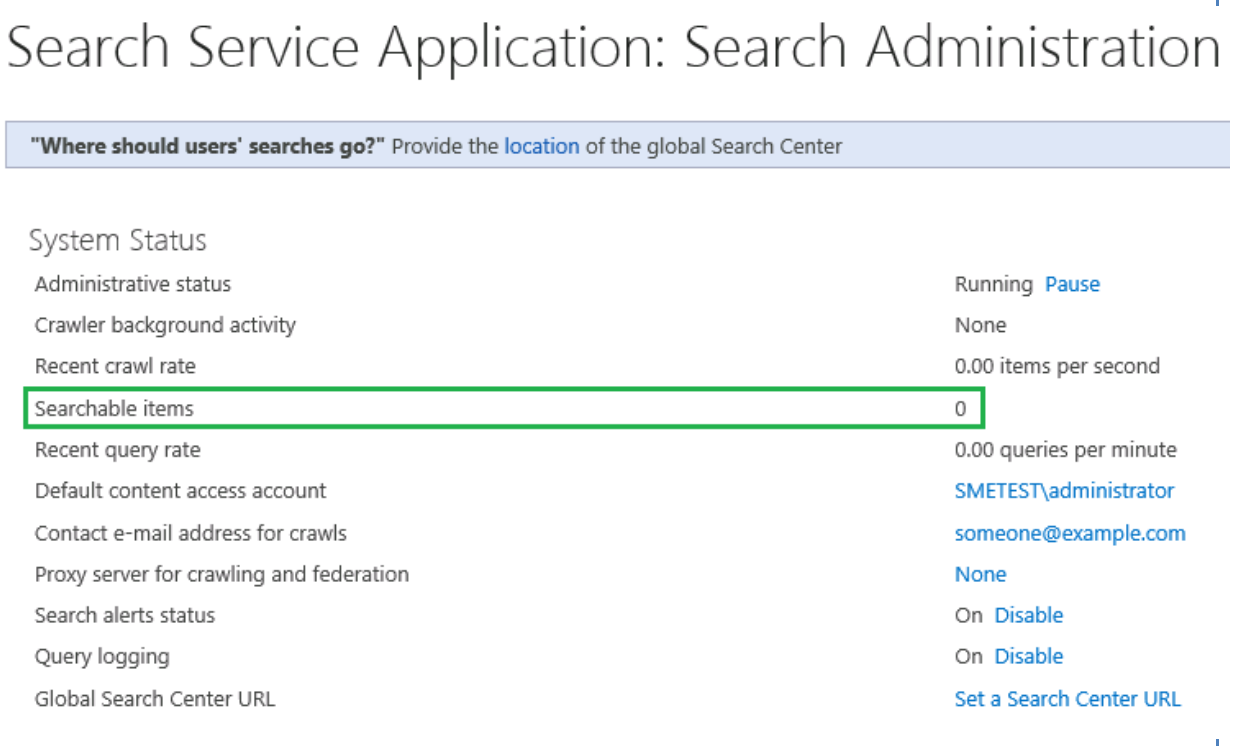

- Confirm that crawl is not running in the Search Administration > Crawling > click Content Sources > Manage Content Sources page >Status column for existing content sources should display Idle.

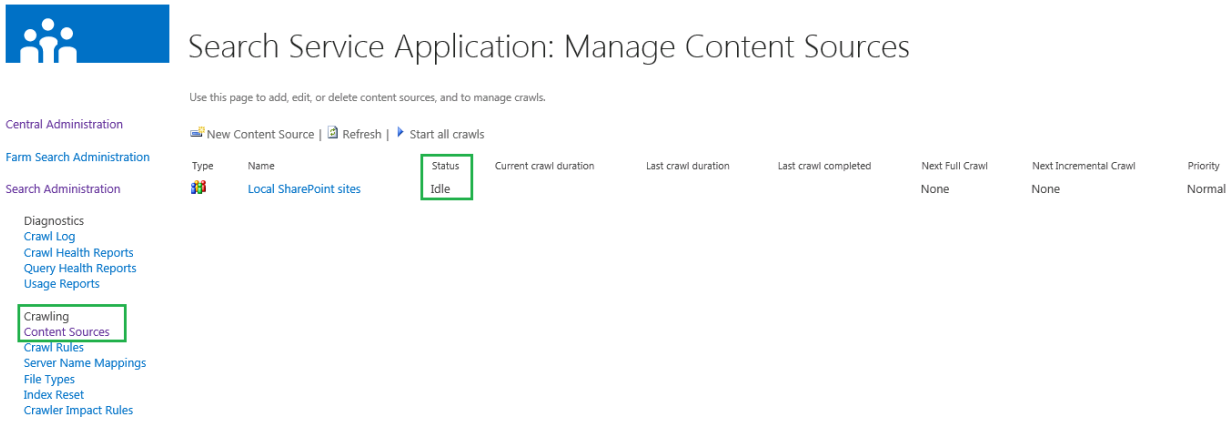

SharePoint User Profile Service Databases

- Profile: User Profile Service Application_ProfileDB_<GUID>

  The profile database stores and manages users and associated information. It also stores information about a user's social network in addition to memberships in distribution lists and sites.

  Estimate the size using the calculation: Database size = Number of users × 1MB.

- Synchronization: User Profile Service Application_SyncDB_<GUID>

  The synchronization database stores configuration and staging data for use when profile data is being synchronized with directory services such as Active Directory®.

  Estimate the size using the calculation: Database size = Number of users × 630KB.

- Social tagging: User Profile Service Application_SocialDB_<GUID>

  The social tagging database stores social tags and notes created by users, alongside their respective URLs.

- Estimate the size using the calculation: Database size = (((number of users × 0.9) x 1) + ((number of users × 0.1) × 6.3)) × 0.009MB.

- App management: App_Management_<GUID>

  The app management database is used by the app management service application. It stores the app licenses and permissions that are downloaded from the SharePoint store or app catalog.

- Secure store service: Secure_Store_Service_DB_<GUID>

  The secure store service application database stores and maps credentials, such as account names and passwords.

- Usage: SharePoint_Logging

  The usage and health data collection database is used by the usage and health data collection service application. It stores health monitoring and usage data temporarily, and can be used for reporting and diagnostics. The usage and health data collection database is the only SharePoint database that supports schema modifications.

- Subscription settings service: SettingsServiceDB

  The Microsoft SharePoint Foundation subscription settings service application database stores features and settings for hosted customers. The subscription settings service application and database are not created by the SharePoint Products Configuration Wizard; they must be created by using Windows PowerShell cmdlets.

- Business data connectivity: Bdc_Service_DB_<GUID>

  The business data connectivity service application database stores external content types and related objects.

Best Practice
• The default recovery model for the all the SharePoint service application database is simple.

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 Media Service Server

The media service is a service of SnapManager for SharePoint. The server hosting the media service manages the following items:

- Backup job metadata, backup index

- Backup IIS metadata
- Backup SharePoint 15 or 14 hives
- SharePoint 2013 shredded storage
- BLOB externalized using Storage Manager
- Archived content from Archive Manager
- Connected content from NetApp CIFS shares

The size of the backup-set indexes created by the media service depends on the level of granularity chosen when creating the backup set. 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 media service server very difficult.

Estimation of Backup Data Size

1. Each SMSP backup job is run as full backup, so one job running will generate a full data set on media service.
2. The number of backup jobs' data for one backup plan to keep on media service is determined by the backup retention policy in the storage policy. Since NetApp volume has the maximum 255 Snapshot copies for each volume, the actual number of backup job data kept on media will be much smaller than 255.
3. The SMSP backup will not save database content to media service; depending on the backup options in backup plan the following data and size can be estimated for backup job.
4. The backup job will save a catalog file on media service, which is small and typically less than 10MB.
5. 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 SharePoint content database. Based on the test results, one document item takes about 1Kb for index, and if the documents have multiple versions, each version takes about 300 bytes. For example, a 211GB content DB with about 3.1 million items, the index data size is about 3GB.
6. Shredded storage feature that "shreds" Office XML document types and stores data in content database or NetApp CIFS share.

   Note: The real-time storage manager rule is in effect with the SharePoint 2013 chunk size instead of real file size. To make chunk size bigger for SP2013 to create bigger size blob data, you can use the following Windows PowerShell script to change the setting in your farm; the default value is 25MB.

```
param([UInt32]$chunksize=25*1024*1024)

[void][System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint")
$websvc = [Microsoft.SharePoint.Administration.SPWebService]::ContentService
Write-Host "The existing FileWriteChunkSize is", $websvc.FileWriteChunkSize
if ($chunksize -gt 0x7FFF0000) { $chunksize = 0x7FFF0000 }
Write-Host "The FileWriteChunkSize will be set to", $chunksize
$websvc.FileWriteChunkSize = $chunksize
$websvc.Update()
```

7. If the backup plan selected and backs up the WFE, the WFE data will save to media service; each WFE backup (WFE and APP server) will take about 10GB storage size.
8. If FAST server is selected in the backup plan, and the FAST data is not installed on the NetApp LUN, then the FAST backup data will copy to media service instead of creating NetApp Snapshot copies. The FAST backup data size will be the FAST server data (about 2GB) plus the index data.
9. If blob backup is selected in backup plan, the index of blob is also saved to media service. Since the blob backup granular level is only at list level, so the index data only need be at list level, and the data size is relatively small in MByte level.

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 will use the size of archive data in SharePoint plus 5% (for metadata and archive manager index usage).

Media service cache location is buffer space used for granular index. If the media service is a single node, you can set it to a NetApp LUN (use the UNC path \\host\driver$\path will actually use the local disk access). If the media service is failover cluster, a file share UNC path or shared disk path needs to be used.

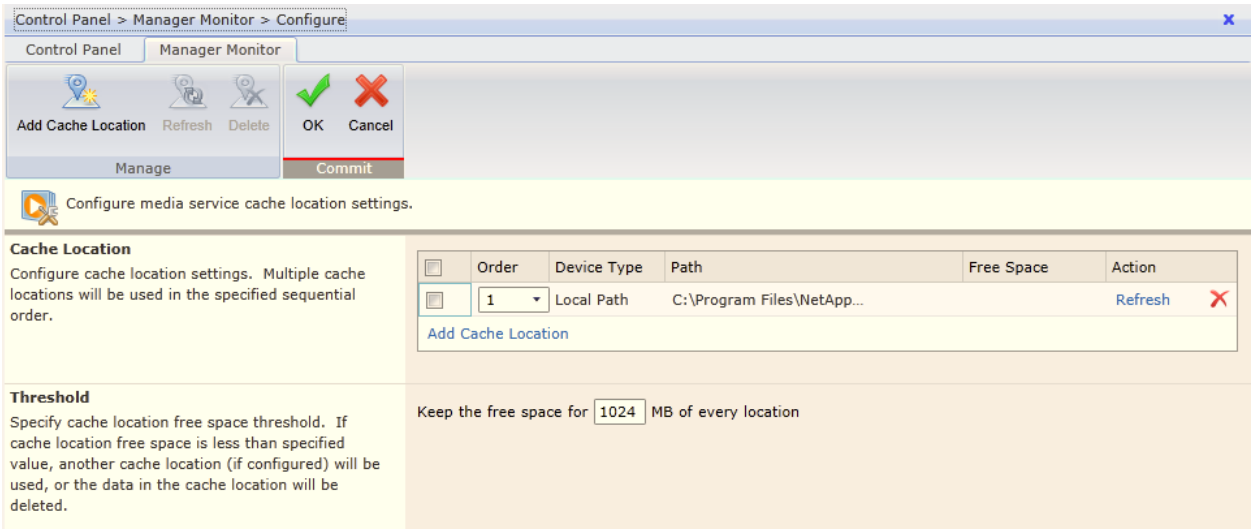

The data growth rate is small for SMSP platform since the backup data is written directly to physical device without using media service cache, and it is only for temporary data usage when creating index.

Best Practices

- Installation of media service requires 1GB free space on the system drive that install SMSP media service.
- Reserve enough space for backup granular index data space if the backup plan has schedules. The media storage size need 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 of its own. This is necessary to cope with the additional processing power needed for storage optimization tasks (storage manager externalization, connector plug in of external content from NetApp CIFS shares and aging out and archiving SharePoint content), as well as managing the backup job data (metadata and index).
- The number of backup jobs data for one backup plan to keep on media service is determined by the backup retention policy in the storage policy. Since NetApp volume has the maximum 255 Snapshot copies for each volume, the actual number of backup job data kept on media will much smaller than 255.
- 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 the Control Service Database

The control service has one control database, which contains the SMSP configuration data, and store backup plans, storage optimization module (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 control database.

Best Practices

- Always use retention policy with backup and archive plan, to make sure 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 by going to job monitor to find the backup job and its ID, then go to 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, then go back to job monitor delete the job monitor record. Deleting the line in job monitor only deletes the record in SMSP database, so you need manually delete the real data folder first.
- Since the control database is changed frequently, the log size may grow if you set the recovery model of control DB to full, unlike the default recovery model of control DB, which is simple.

Sizing for Media Service Server Datastore

Depending on the backup options in backup plan, the following data and size can be estimated for 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,
   - Index size = 2.5*(Nsc*53 + Ns*27 + L*0.45 + D*V*0.35) (Kbyte)
   - Where:
   - Nsc: Number of site collections in content database
   - Ns: Number of subsites in content database
   - L: Number of list/folders in content database
   - D: Number of items
   - V: Average number of version for documents.
   - We can estimate the index size from database size:
   - Index size = 2.5*(Nsc*53 + Ns*27 + L*0.45 + (SDB/S)*0.35) (Kbyte)
   - Where:
   - Nsc: Number of site collections in content database
   - Ns: Number of subsites in content database
   - L: Number of list/folders in content database
   - SDB: Content database size
   - S: Average document size in content database
3. If the backup plan selected backup WFE, the WFE data will save to media service, each WFE backup (WFE and APP server) will take about 10GB storage size. If user has other customization installed, the WFE backup size will also need to count the customized solutions size.

Sizing for BLOB Storage

The SMSP BLOB stub database keeps record of each blob, with each blob record using about 300 bytes.

Best Practices
• The BLOB stub database is set to simple recovery model to avoid big transaction log size. • If user needs change to full recovery model, make sure that the LUN has enough space for transaction log, and shrink the log if it is necessary. • BLOB stub database size may increase if there is massive document upload or use connector to connect NetApp CIFS shares with large number of files. • Use NetApp CIFS shares on SATA for storage optimized data for efficiency benefits.

SharePoint BLOB Data on NetApp CIFS Shares

The BLOB data, archived from the SharePoint content databases, should be placed on a NetApp CIFS shares for faster recoverability.

Volume Sizing

Volume sizing has two parts, database volume sizing and transaction log volume sizing. When sizing SharePoint volumes, NetApp recommends calculating the appropriate amount of space needed for each LUN. When creating new MDF and LDF files, presize them to minimize autogrowth events. NetApp recommends setting the file sizes appropriately and growing them manually at times of minimal system use on a planned basis. Autogrow should be used only as a safety net to prevent the files from becoming full and making the database read-only at times when unpredicted substantial growth occurs. When database files are expanded, there is an impact on performance.

Figure 9) Volume and LUN sizing decision making.

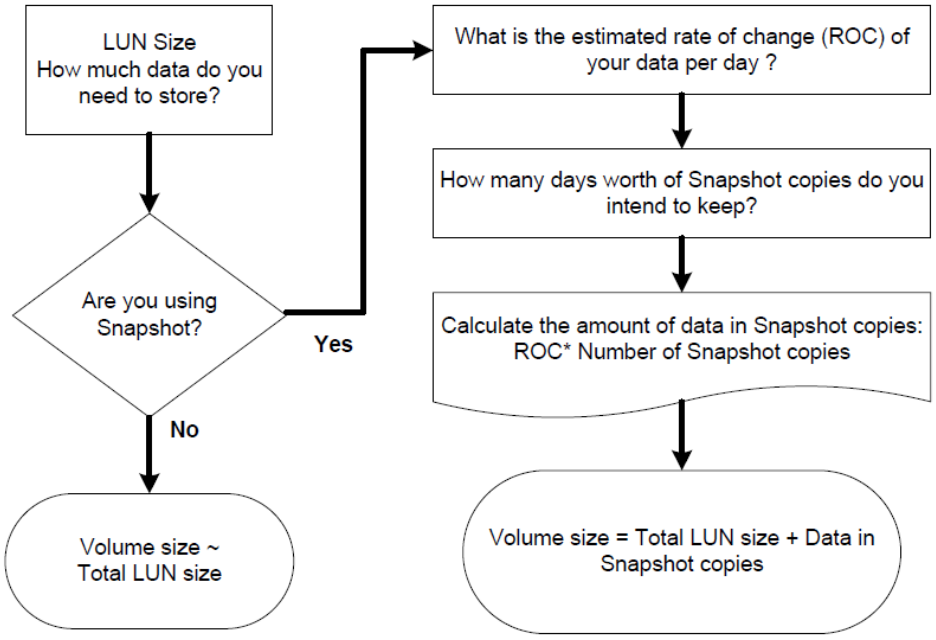

Note: Data ONTAP can support about 1,000 volumes.

General Rules

a. What is the LUN size? Do you plan to have multiple LUNs in a volume?

b. Do you want to maintain Snapshot copies?

c. If you want to maintain Snapshot copies, what is the number of Snapshot copies you want to maintain, and how long do you plan to retain them (retention period)?

d. At what rate does data in the volume change (delta)?

e. What amount of space do you need for overwrites to LUNs (fractional reserve)?

Best Practices

- NetApp strongly recommends that you use SnapDrive for Windows to configure initiator groups.
- SnapDrive for Windows also allows administrators to shrink or grow the size of LUNs. Never expand a LUN from the storage system; otherwise, the Windows partition will not be properly expanded.
- Create an immediate backup after expanding the LUN so that its new size is reflected in the Snapshot copy. Restoring a Snapshot copy made before the LUN was expanded will shrink the LUN to its former size.
- Calculate LUN size according to SharePoint sizing guides and calculate for Snapshot copy usage if Snapshot is enabled.
- If SnapVault integration is planned, the LUN should always be placed in a qtree.

4.4 Performance

For best performance with the SharePoint data access, with the optimal load on the SQL Server instances, simultaneously lowering the overall cost, follow these best practices.

Best Practices

- Leverage the SMSP storage optimization modules to externalize BLOB for faster SharePoint performance.
- SMSP archive job is fairly resource-intensive, so running multiple archive jobs simultaneously may affect the performance of the server. To avoid this condition, configure SMSP processing pool feature where archive 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 archive jobs that can be run simultaneously. The remaining archive jobs are placed in a queue.
- For SharePoint deployments where the individual content size is small enough (for example, less than 256KB), it is best to keep all of the content in SQL Server database stored on NetApp FAS disks.
- For any larger objects, it’s better to deploy RBS to keep the BLOBs in the external NetApp CIFS shares on less-expensive SATA disk storage leveraging NetApp Flash Cache technology.
- When evaluating SnapMirror performance, you should consider CPU impact, network bandwidth between source and destination NetApp storage systems, impact due to write latency caused by network devices.

To maximize performance in SharePoint deployment with RBS, follow these best practices.

Best Practices

- If the media service server is going to be virtual machines, make sure to have good memory size and CPU power—that is, a heavy host configuration. When using the storage optimization modules with SMSP, the media service server is bypassed for BLOB access.
- Disable the time synchronization for each SharePoint virtual machine. SharePoint 2013 implements timer jobs extensively, and the latency during time synchronization will cause unpredictable results in the SharePoint environment.

5 NetApp Solution for SharePoint Server

5.1 SnapManager 7.1 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). For detailed information on which SharePoint 2013 elements are supported in which modules, refer to SnapManager 7.1 for Microsoft SharePoint Platform Backup and Restore User's Guide.

This management software facilitates SharePoint administration tasks and lowers the data center administration cost.

5.2 SnapManager 7.1 for SharePoint Architecture

SnapManager for SharePoint (SMSP) 7.1 is a 64-bit application that is capable of protecting Microsoft SharePoint 2013 and 2010 databases.

It was designed with the following objectives:

- To provide centralized management of backup and recovery for multiple SharePoint farms.
- To provide data protection through the ability to back up all objects and entities of a SharePoint farm.
- To provide BLOB externalization capability using the real-time storage manager and scheduled storage manager.
- To integrate with NetApp SnapMirror and SnapVault.

Figure 10) SMSP 7.1 architecture.

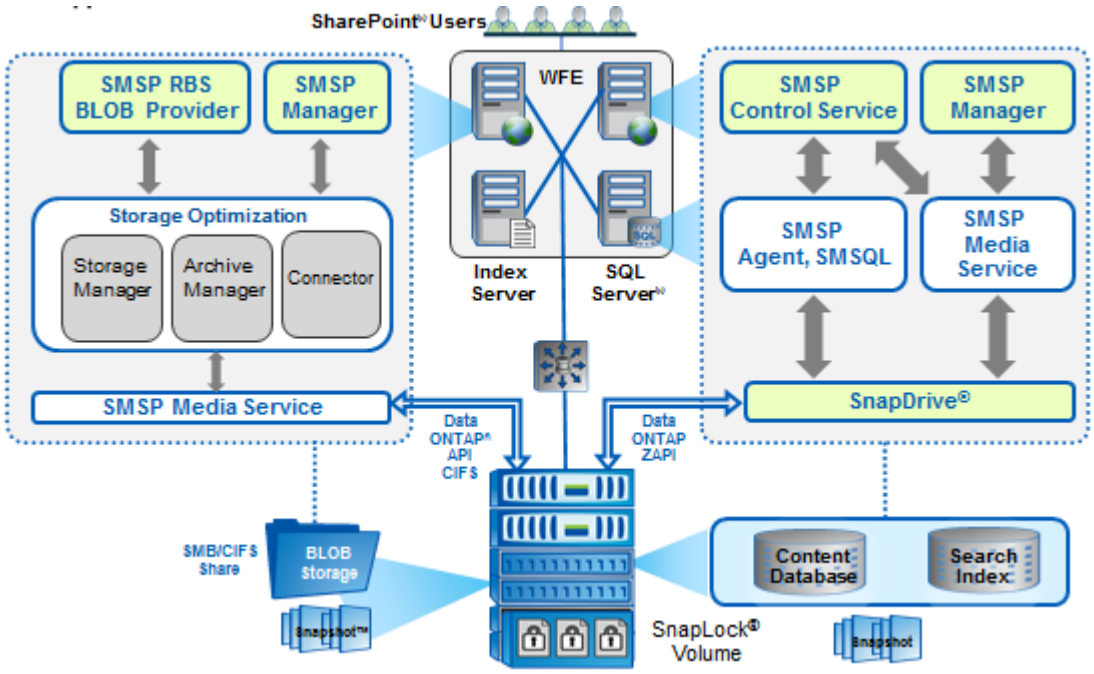

Figure 11) SMSP 7.x components overview.

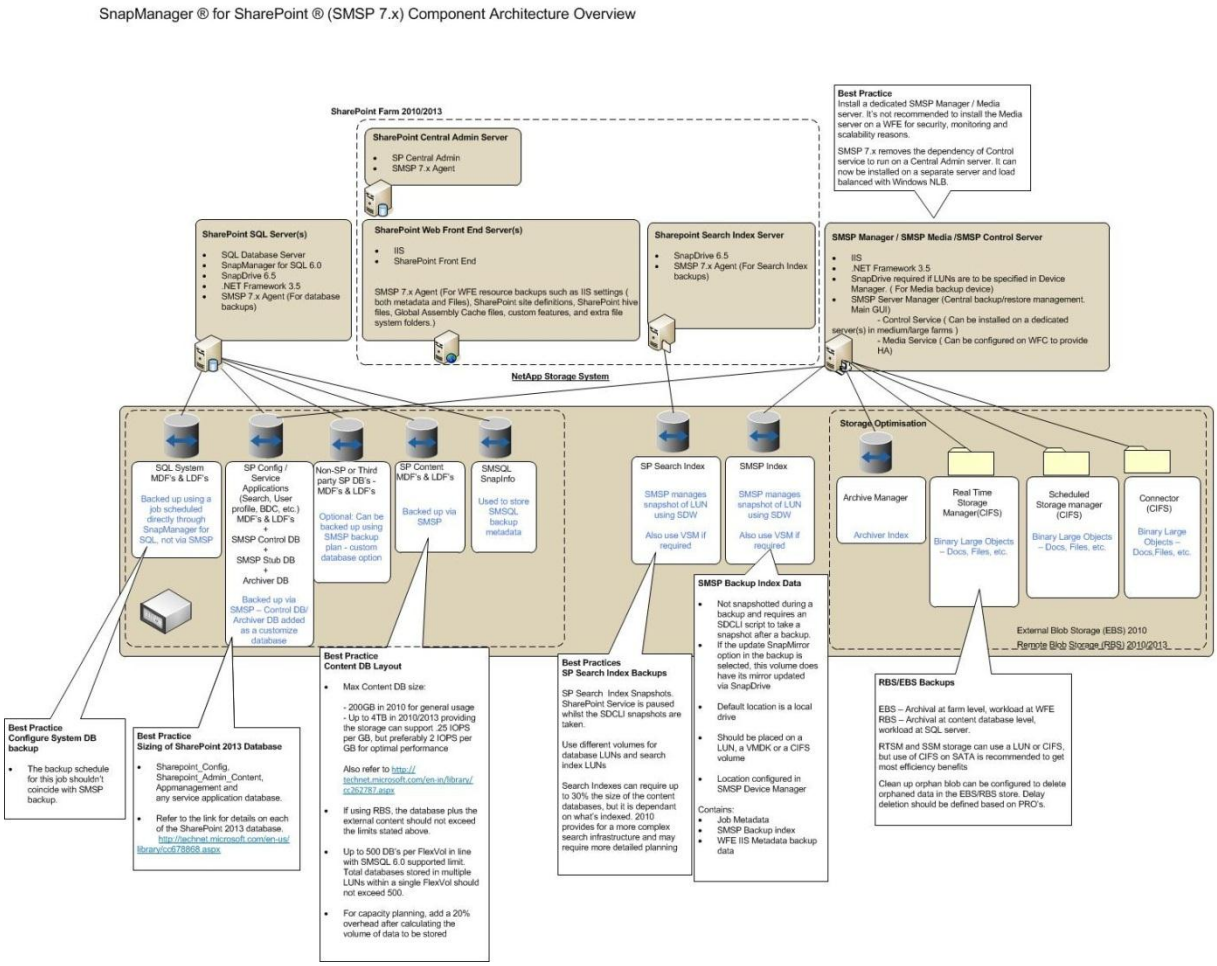

Table 6) SnapManager 7.1 for SharePoint components mapped to SharePoint farm hosts.

Server Role	SMSP 7.1 Component	Remarks
Server that is not part of the SharePoint farm	SMSP manager	Mandatory.
Media server	Media service	Mandatory. Should be installed on every media server that uses SMSP manager software.
WFE servers	SMSP agent (storage optimization – storage manager)	Optional. Enabled only for performing stub-based uploads of external documents for one or more of the web applications hosted on the WFE.
	SMSP agent (storage optimization – connector)	Optional. Enabled only for performing stub-based uploads of external documents for one or more of the web applications hosted on the

Server Role	SMSP 7.1 Component	Remarks
		WFE.
	SMSP agent (storage optimization – archive manager)	Optional. Enabled only for archiving the contents of one or more of the web applications hosted on the WFE.
Index server	SMSP agent	Optional. Installed only for backing up the SharePoint search indexes.
SQL Server host	SMSP agent	Mandatory.

5.3 Storage Optimization

Microsoft offers remote BLOB storage (RBS) as the official offloading technique for BLOB externalization, implemented by SQL Server and available in SharePoint 2013 and SharePoint 2010, based on the API supported by SQL Server 2012 and SQL Server 2008 R2. SMSP 7.1 includes storage optimization solutions to keep your SQL Server resources optimized with intelligent archiving and BLOB offloading to a NetApp CIFS share on less expensive storage. Deduplication and compression enabled on NetApp storage will work on externalized BLOBs to provide improved I/O operation. 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 only supports the SQL filestream provider for Microsoft SQL Server. SnapManager for SharePoint does not support other third-party RBS providers.

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

Best Practices

- Externalize BLOB data files larger than 256 kilobytes (KB). While this is the guidance, it also important to consider the kind of workload you have implemented. If the workload is more WORM oriented, then BLOB is a great use. However, in a high-collaboration environment (creating, editing, multiperson, multiversion), it might be best for the data to reside in SQL Server database to lessen the impact of fetching from second-tier storage.
- NetApp recommends you creating the RBS datastore on a NetApp volume 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.
- The use of RBS-enabled content databases larger than 4TB with collaboration sites is not supported. You cannot upload any document larger than 2GB to an RBS-enabled content database.
- Make sure that the volumes for the NetApp CIFS shares are big enough to hold a large amount of BLOB data.
- SMSQL can back up and restore filestream data directly when we back up/restore the database. Be aware that this will only work if local RBS filestream provider is used; the remote RBS filestream provider is not supported.

For additional information, refer to SnapManager 7.1 for Microsoft SharePoint Storage Optimization User's Guide.

5.4 Backup Guidelines

During backup, SnapManager for Microsoft SQL Server is used to perform database Snapshot backups, and SnapDrive is used to perform 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.

Guidelines on using federated backup with SMSQL:

- SMSP 7 does not support the federated backup function of SMSQL because of the following reasons:
  - In federated backup, the SMSQL agent that manages the backup can get detailed report of cmdlet; all other SMSQL servers will not send detailed job report of cmdlets, so SMSP can't get the detailed job report of SMSQL cmdlet, which is important for job report and troubleshooting.
  - The federated backup can only pass one user credential at command line; if the SMSQL servers use different credential the federated backup will not work.

Instead of using federated backup, user can install SMSP agent on all the SQL Server instances (with SMSQL installed), and add the database on SQL Server as customized databases; then SMSP will trigger backup on the SQL Server instances and get the correct job report of SMSQL cmdlet on each SQL Server instance. Since SMSP on each server can configure with different user accounts, different user credential on different SQL Server can also be used.

Best Practices

- SharePoint and FAST data should reside on a NetApp LUN for them to be backed up by SnapManager for SharePoint.
- 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 DR SLAs demand more frequent backups.
- SQL Server database backup and BLOB data backup schedule frequencies do not have to match, due to differences in storage natures (size, item number, on different LUNs with limited Snapshot capacity, and so on); they can be scheduled separately and freely, as long as there is a BLOB storage backup happening not too long after database backup.
- Performing a backup does not affect the state of the farm/web application/service application/site collection. However, it does require resources and might slightly affect farm performance when the backup is running. You can avoid performance issues by backing up the farm during hours when farm use is lowest, such as outside office hours.
- Backing up the farm configuration will not back up the information needed to restore service applications. If you want to restore a service application, you must perform a configuration and content backup of the farm.
- SQL Server database backup and BLOB storage backup serve different purposes in SMSP. Database backup is for both disaster recovery (DR) and non-DR (item-level recovery). BLOB storage backup is only for DR.
- Make sure that you plan the layout of SharePoint databases on the FlexVol volumes (flexible volumes) and LUNs since the amount of data backed up inversely relates to the duration of the subsequent backup window.
- Storage of BLOB data for SMSP storage manager and connector MUST reside on NetApp storage since SMSP checks this when you leverage a physical device on NetApp storage system CIFS share. SMSP requires that the SharePoint databases reside on the NetApp LUN to leverage SMSQL and NetApp Snapshot technology.
- You cannot back up during a restore and vice versa. Hence, use multiple SQL Server instances in large environments to avoid SMSQL contention.
- SnapManager for SharePoint supports VMDK disks, which interact with SMSP identically to LUNs, except that SnapVault integration is not supported on VMDK disks.
- In a mirrored setup, if SQL Server authentication is the SharePoint content database's 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.

For additional information, refer to the SnapManager 7.1 for Microsoft SharePoint Platform Backup and Restore User's Guide. In this guide, section "SharePoint Components Supported for Backup" has more details about the SharePoint 2013 components supported by SMSP 7.1 backup feature.

5.5 Restore Guidelines

During restore, SnapManager for Microsoft SQL Server is used to perform restore from database Snapshot backups.

Restore from Alternate Location

Restore from SnapVault

The SnapVault backup data (or the remote backup data) is used only when the local backup has already been deleted. Restore from SnapVault requires manual intervention; that is, clone database from

SnapVault secondary to primary SQL Server with the name required by SMSP that needs to be verified to be able to perform granular content restore.

Best Practice
• NetApp SnapVault restore will overwrite the primary qtree/volume; hence you need to make sure that you copy/back up the CIFS share first, after SnapVault restore, then copy back to merge data in the folder together.

Note: Only database backup to SnapVault is supported. Restoring the SharePoint search index and BLOB from SnapVault is not supported.

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 SMSP can operate on the destination farm. SMSP does not require the destination content database to be on NetApp LUN, which is a requirement of SMSQL and NetApp Snapshot technology. 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. 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/IP and/or iSCSI.

Best Practices

- NetApp recommends verifying the backups before performing any restore operations. For verification on the SnapVault and SnapMirror destinations, offload SQL Server verification and storage to the SQL Server instance in the destination location.
- You need restore BLOB data first, then restore content and stub databases, also disable the garbage collection (blob retention) before the database restore finishes.
- Performing a restore does not affect the state of the farm/web application/service application/site collection. However, it does require resources and might slightly affect farm performance when the backup is running. You can avoid performance issues by backing up the farm during hours when farm use is lowest, such as outside office hours.
- To restore customizations successfully, NetApp recommends that you deploy the .wsp file for both trusted and sandboxed solutions to the destination.
- BLOB restore is at the site collection level and Connector BLOB restore is at the list level. No item-level BLOB restore is possible.
- You cannot restore the configuration database of a server farm or the content database of the central administration application individually. But restoring a farm, including its complete configuration database and the content database of the central administration application, is not a flexible option because with that option server names and topology information for the restoration target farm must be identical to the corresponding data for the source farm. NetApp recommends that after SharePoint Foundation has been installed to a target farm, farm administrators restore the configuration settings and then restore web applications and other content as needed. If you do this, the configuration settings function similar to “farm template” that does not presuppose any particular farm topology.
- You cannot back up during a restore, and vice versa. Hence, use multiple SQL Server instances in large environments to avoid SMSQL contention.
- Restore the farm components prior to restoring front-end resources.
- Make sure that the source node and the destination node are the same version of 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.

For additional information, refer to SnapManager 7.1 for Microsoft SharePoint Platform Backup and Restore User's Guide.

5.6 High Availability

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

Control Service High Availability

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. High availability of control service configured using Microsoft Windows cluster failover configuration means if any are offline; there is still access to SMSP 7.1 and the ability to manage SharePoint farm and storage. First control service installed is the master, which can be changed. Also because the ControlDB for the control service is now in SQL Server, clustering and log shipping apply for HA.

Media Service High Availability

If you are using SnapManager 7.1 for SharePoint to manage your SharePoint farm, then media service plays a very important role. It is critical to provide high availability for the media service. High availability of media service can be configured using Microsoft Windows cluster failover configuration; it requires all LUN/CIFS physical devices have the same drive letter and mount point on both 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.

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.

6 NetApp SnapVault

SnapVault data sets backup for backups of SharePoint databases and search indexes, storage manager BLOB data, archive data volume and connector BLOB data, and verification of SnapVault targets. SnapVault protects data on a primary system by maintaining a number of read-only versions of that data on a secondary system. First, a complete copy of the dataset is pulled across the network to the SnapVault secondary system. This initial, or baseline, transfer might take some time to complete because it duplicates the entire source dataset on the secondary system, much like a level-zero backup to tape. Each subsequent backup transfers only the data blocks that changed since the previous backup. When the initial full backup is performed, the SnapVault secondary system stores the data in a WAFL file system and creates a Snapshot image of the volume for the data to be backed up. A Snapshot copy is a read-only, point-in-time version of a dataset. SnapVault creates a new Snapshot copy with every transfer, allowing retention of a large number of copies according to a schedule configured by the backup administrator. Each copy consumes an amount of disk space proportional to the difference between it and the previous copy.

Note: Use Data ONTAP command SnapVault restore to restore using the secondary Snapshot copy. However, there are some existing issues. In the following case, we use the file stream restore method to restore from the destination Snapshot copy.

Best Practices

- Make sure that you set up the Protection Manager datasets using SMSQL.
- User needs to set up the SnapVault and policy in NetApp Protection Manager and relationship between source NetApp storage system and the secondary storage. SMSP 7.1 will trigger SnapVault update each time a backup is run, provided the option to update SnapVault is selected in the backup plan.
- The SnapVault update can be configured for different data, the SharePoint database backup data, the blob backup data, or the backup data on SMSP storage policy device.
- To restore database from SnapVault:
  - You can restore from SnapVault directly as long as SMSQL still has the right configuration and Protection Manager policy.
  - Choose to "Restore from Alternate Storage Location" and the SMSP GUI will display the Snapshot name and the user manually restore from SMSQL GUI with Protection Manager.
- To restore blob from SnapVault, user needs to find the volume and Snapshot from job report and manually restore with Protection Manager.

7 SharePoint Disaster Recovery with SMSP

There are numerous ways to augment data availability in the event of hardware, software, or even site failures. The SMSP DR is more for warm and cold scenarios. The hot scenario is considered to be high availability (HA), and SMSP 7.x does not support HA. In warm DR, you can do in-place replace of farm by restore database only. Cold DR basically means you rebuild the whole farm, then use the SMSP farm rebuild wizard. The RPO requirements defined decide the type of data synchronization (synchronous, asynchronous, and semi-synchronous).

7.1 NetApp SnapMirror

NetApp SyncMirror maintains two copies of the SharePoint data online so that the data is available and 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 Sync 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.

Benefits of Using SyncMirror

- Synchronous replication makes sure that a copy of the data (SharePoint databases, search index, storage manager BLOB, archive data, and connector BLOB) is always up-to-date and available in the local data center.
- Transparent operation enables SharePoint to access the synchronously replicated data with no programming or system changes.
- Integration with active-active controller configuration offers complete hardware redundancy with automatic failover for the highest level of data availability.

For information about the supported and unsupported configurations for synchronous and semi-synchronous SnapMirror modes, refer to TR 3326: SnapMirror Sync and SnapMirror Semi-Sync.

NetApp MetroCluster with SyncMirror

Both stretch MetroCluster™ and storage controllers in active-active configurations are supported up to 500 meters apart with 2Gb/sec connectivity, or 270 meters with 4Gb/sec speeds. This generally means that all controllers are in the same data center. SyncMirror is required as part of stretch MetroCluster to make sure that an identical copy of the data exists in case the original data center is lost.

Note: To extend active-active configurations across data centers up to 160 kilometers apart, consider fabric MetroCluster along with SyncMirror.

The benefits of MetroCluster and SyncMirror for high availability are as follows:

- Fabric MetroCluster and SyncMirror provide the highest level of storage resiliency across a local region.

MetroCluster and SyncMirror provide the highest levels of regional storage resiliency for continuous data availability in a particular geography.

Best Practices

- Since MetroCluster is used for mirroring data for NetApp storage controller in distance, MetroCluster is more used for DR scenario. User can configure the data on storage controller to be mirrored to another NetApp storage controller at another location; the other location will also have the same SharePoint farm topology (host name and so on) as source side in “cold.” When the source side is down, the DR site farm can start up and use the data on the storage controller on DR side. Assuming the DNS at DR side can point to the new storage controller, then all SMSP device can be used; otherwise the SMSP storage system profile needs to update the storage system IP when the MetroCluster failover occurs.
- User needs to manually set up the SnapMirror relationship between source NetApp storage system and the SnapMirror partner. SMSP 7.1 will trigger SnapMirror update each time run a backup if the option to update SnapMirror is selected. The SnapMirror update can be configured for different data, the SharePoint database backup data, the blob backup data, or the backup data on SMSP storage policy device. To minimize write impact of the primary workloads, do not exceed 2 milliseconds of round trip time (RTT) for synchronous SnapMirror.
- To thin provision an aggregate on the SnapMirror source, create FlexVol volumes with a guarantee of none or file so that the volume size is not limited by the aggregate size. The total size of the FlexVol volumes can be larger than that of the containing aggregate.
- Use caution when thin provisioning an aggregate at the destination system, because SnapMirror fails when the volume runs out of space. Starting with Data ONTAP 7.3, it is possible to set guarantees on the SnapMirror destination volume so that the SnapMirror updates never fail on that volume. The default behavior is that the volume guarantees are turned off. The following example demonstrates space usage with and without volume guarantees on the SnapMirror destination volume, For example: For a 1TB SnapMirror source volume that is 75% full, the SnapMirror destination volume (or replica) needs 750GB with the guarantee disabled and the full 1TB with the guarantee enabled.
- When the primary NetApp storage system data is damaged, user needs to manually break the SnapMirror volume and make the secondary SnapMirror volume as primary.
- To restore database from SnapMirror, the “Restore from Alternate Storage Location” is selected, and the SMSP GUI will display the Snapshot name and database name on the SnapMirror volume, which will be used for restore.
- To restore blob from SnapMirror, user needs to find the volume and Snapshot copy from job report and manually restore. To make sure that storage manager and connector can be used after disaster recovery, you must add their stub databases as customized databases and include those stub 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 and the storage system profiles/physical devices/logical devices are configured using the correct configurations/paths before performing the disaster recovery.

References

The following references were used in this report:

- TR-4193 SnapManager 7.1 for SharePoint with Data Clustered Data ONTAP - Best Practices Guide
- 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
- 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 7.1 for Microsoft SharePoint Installation Guide
- SnapManager 7.1 for Microsoft SharePoint Storage Optimization User's Guide
- SnapManager 7.1 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
- Protection Manager (example)
- SharePoint Community

Version History

Version	Date	Document Version History
Version 1.0	June 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