Technical Report

Microsoft SharePoint and SnapManager 8.2 for SharePoint Best Practices Guide

Cheryl George, NetApp

June 2015 | TR-4431

Abstract

This document discusses the planning considerations and best practices for deploying Microsoft SharePoint 2013 and Microsoft SharePoint 2010 on NetApp® storage systems. It also covers the best practices for the NetApp enterprise data management solution for SharePoint, which is called NetApp SnapManager® 8.2 for SharePoint.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

Microsoft SharePoint is a collaboration platform used by enterprise businesses for their intranet/extranet websites as a method to centralize access to share and organize information, people, and projects. Microsoft SharePoint implementations have become more complex and require a greater degree of reliability. The underlying physical storage architecture that supports SharePoint farms is expected to scale to meet growing capacity needs with suitable data protection capabilities for high-availability (HA) and disaster recovery purposes.

1.1 Purpose and Scope

This document provides guidance for achieving effective and efficient storage planning and end-to-end data protection with design considerations and best practices when deploying SharePoint Server 2013 and SharePoint Server 2010 on NetApp storage systems, with SnapManager 8.2 for SharePoint (SMSP 8.2) as the data protection solution. 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. End-to-end implementation is out of the 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 high-performing, available, and easy-to-manage SharePoint environment. The information also enables them to meet stringent service-level agreements (SLAs). Readers should have working knowledge of the following:

- NetApp clustered Data ONTAP® operating system
- NetApp SnapDrive® for Windows (SDW) data management software
- NetApp SnapManager for SQL Server (SMSQL)
- NetApp SnapManager for SharePoint (SMSP)
- Microsoft SharePoint Server 2013 and 2010 architecture and administration
- Microsoft SQL Server 2014, 2012, 2008 R2, or 2008

For configuration compatibility across the NetApp stack, refer to the NetApp Interoperability Matrix Tool.

1.2 Intended Audience

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

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

2 Introduction to SharePoint Server 2013

SharePoint is an n-tier service-oriented architecture that consists of a SharePoint farm with web and service applications. SharePoint Server 2013 is the latest release to provide features for content management, social computing, enterprise search, business intelligence, identity management, mobile computing, and branding through customization. For more details, refer to Technical diagrams for SharePoint 2013.

2.1 Physical Architecture

The Microsoft SharePoint Server 2013 environment consists of multiple servers, each of which is assigned a specific role, as displayed in Figure 1.

Figure 1) SharePoint physical architecture.

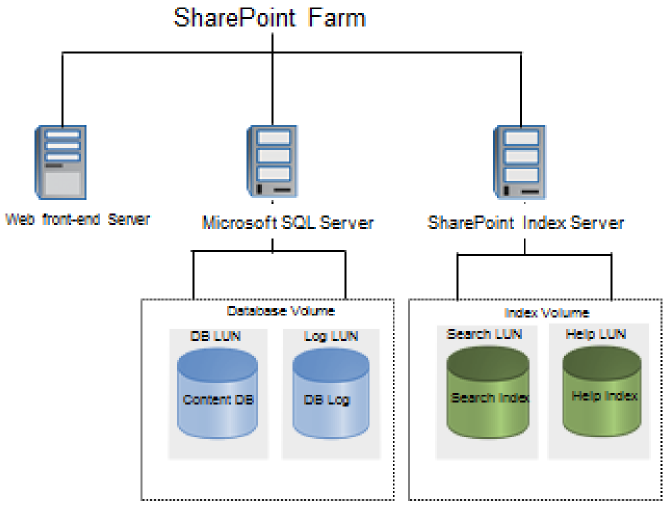

SharePoint Farm Topologies

The SharePoint farm topology can vary from a single server or standalone farm topology to large farms that generally include more SharePoint services and more web applications and typically handle more server requests. For more details on SharePoint topologies, see these Microsoft articles:

- Topologies for SharePoint Server 2010
- Traditional topologies for SharePoint 2013

2.2 Logical Architecture

Figure 2 shows the logical architecture of the SharePoint farm configuration. For more details, refer to the article Plan logical architectures for SharePoint 2013.

Figure 2) SharePoint logical architecture.

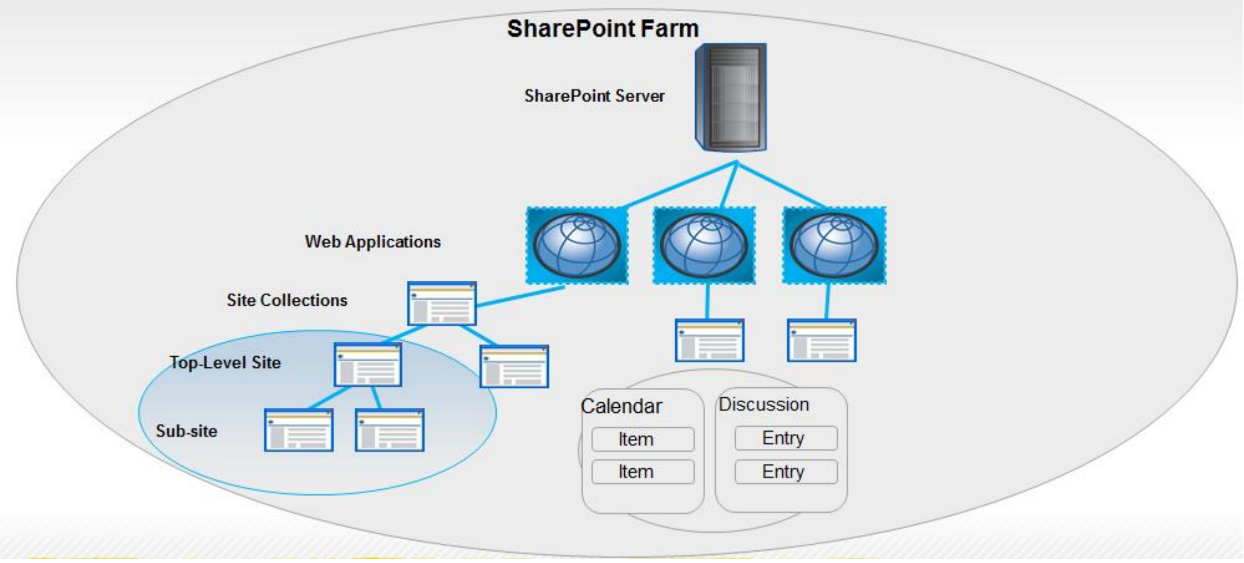

For complete details on logical and physical architecture for SharePoint farms, refer to Architecture Design for SharePoint 2013 IT Pros.

3 Database and Storage

The combination of NetApp storage solutions and Microsoft SharePoint enables the creation of enterprise-level database storage designs that can meet the requirements of the most demanding applications. To optimize both technologies, an appropriate layout of SharePoint databases is necessary for performance, faster access, recoverability, and management of the SharePoint infrastructure. A well-designed storage layout for a SharePoint farm database supports successful initial deployment that allows smooth growth over time and does not affect performance and management of the SharePoint infrastructure.

3.1 Aggregate

Aggregates are the primary storage containers for NetApp storage and contain one or more RAID groups consisting of data disks and parity disks. With large-sized serial ATA (SATA) disks, an increased spindle count with increased disk count can help maximize performance and maintain high storage efficiency due to large size thresholds. NetApp recommends having a separate aggregate to host the SharePoint application with more RAID groups and spindles to optimize and improve storage performance. It becomes easier for administrators to manage for two reasons:

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

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

- Number of SharePoint farms with high-availability requirements for various SharePoint farm content
- Size of the SharePoint content
- Service applications deployed
- Features installed

Best Practices
• NetApp recommends having at least 10% free space available in an aggregate hosting SharePoint data for optimal storage performance. • Make sure that both LUNs for data files of a particular database reside in the same controller node\aggregate because a node failover of such a failover event in a rolling upgrade will cause the database to go offline. Choose to use both controller node CPUs in the HA pair for better performance while storing the NetApp SnapMirror® copy of this database for disaster recovery purposes or an Availability Group (AG) replica.

For additional information, refer to TR-3929: Reallocate Best Practices Guide.

3.2 Volume

NetApp FlexVol® volumes are created and reside inside aggregates. Volume layout is critical in creating and sustaining a highly available SharePoint environment through careful consideration of workload I/O characteristics and backup and recovery requirements.

Figure 3 shows a typical volume layout for a SharePoint environment.

Figure 3) SharePoint volume layout.

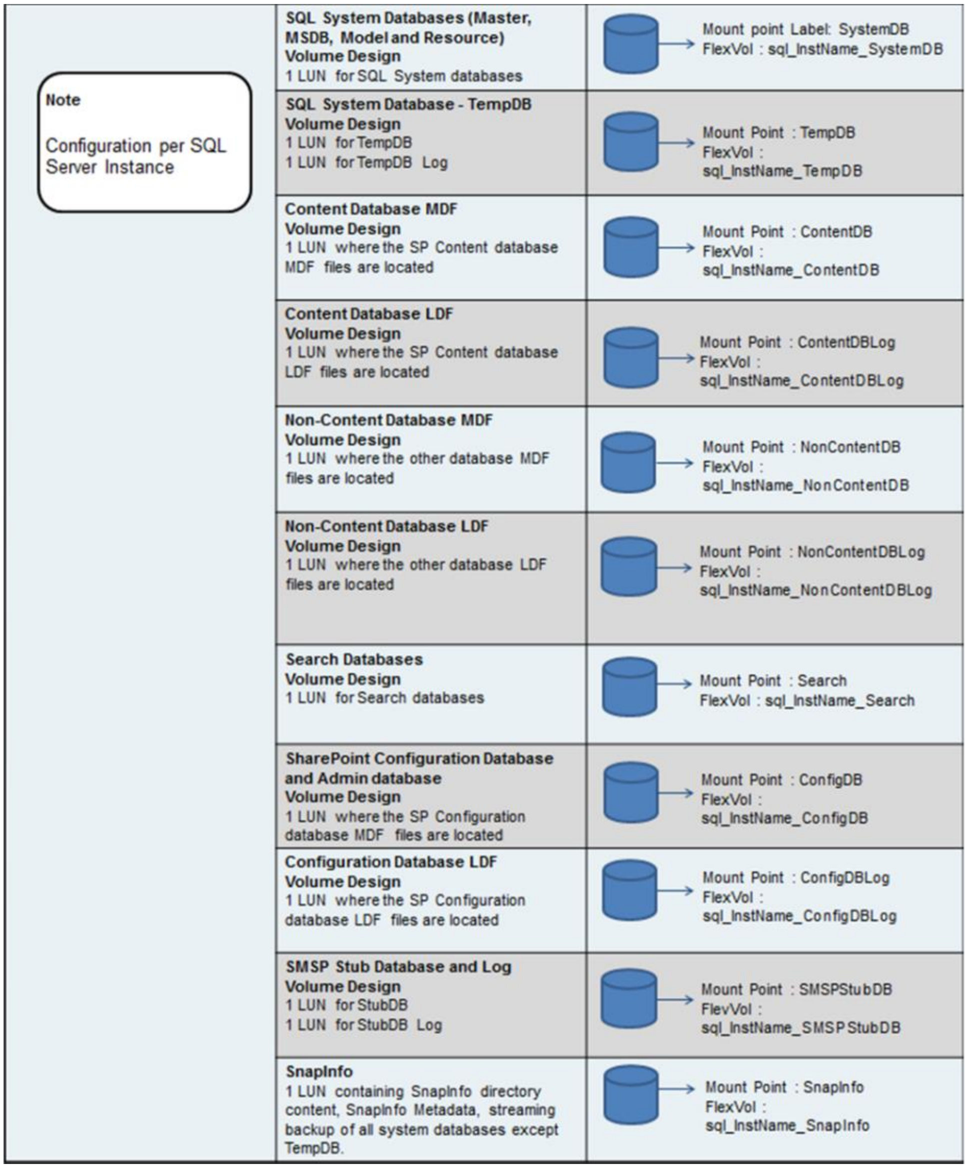

Note Configuration per SQL Server Instance	**SQL System Databases (Master, MSDB, Model and Resource)** **Volume Design** 1 LUN for SQL System databases	Mount point Label: SystemDB FlexVol : sql_InstName_SystemDB
	SQL System Database - TempDB **Volume Design** 1 LUN for TempDB 1 LUN for TempDB Log	Mount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

Best Practices

- Use separate FlexVol volumes to store Windows OS and SharePoint binaries.
- Place the SQL Server system databases on a dedicated volume or virtual machine disk (VMDK). This is needed because colocating system databases with SharePoint databases prevents NetApp Snapshot® technology backup of the user databases, with backups getting streamed into the SnapInfo LUN.
- `tempdb` is a system database used by SQL Server as a temporary workspace, especially for write I/O–intensive operations; for example, during `DBCC CHECKDB` operations. Place this database on a dedicated volume with a separate set of spindles. In large environments in which volume count is a challenge, after careful planning you can consolidate `tempdb` into fewer volumes and store it in the same volume as other system databases. Data protection for `tempdb` is not required because this database is recreated every time SQL Server is restarted.
- Place SharePoint data files (.mdf) on separate volumes from those of transaction logs to isolate the random read-write I/O from the sequential write I/O of the log files, thereby significantly improving SQL Server performance.
- For large SharePoint content databases, consider using multiple data files for improved performance.
- Allocate a dedicated volume with a separate set of spindles for the SharePoint search databases.
- Make sure that the SharePoint databases and the BLOB data reside on separate volumes.
- For SMSP Storage Manager the I/O performance has no impact when using multiple farm-web applications sharing one volume or separate volumes for each farm-web application. However, from a backup/restore point of view, the binary large object (BLOB) storage volume should not be mixed between farms to make backup data retention management easier. Also, if users plan to create multiple backup plans to group web applications together, NetApp recommends using one dedicated volume for the web applications in the backup plan, so BLOB backup/restore can be easier to manage with retention.
- NetApp recommends using one index partition for every 10 million items in the search index.
- If server resources are sufficient, place index partitions for different servers in different volumes to make sure that search performs well. If server resources are insufficient, place index partitions for different servers on the same volume. However, configure index replicas for fault tolerance.
- Avoid sharing volumes and datastores between different Windows host machines.
- Disable opportunistic locking (oplocks) on volumes hosting CIFS shares containing SharePoint BLOB data to avoid corruption from caching.
- Configure volume autosize policy whenever appropriate to help avoid out-of-space conditions.

3.3 LUNs

NetApp storage can be presented to Windows hosts as logical units called LUNs that appear as local hard disks to the server. NetApp Fibre Channel (FC) or iSCSI protocol LUNs can be created using SnapDrive for Windows.

Table 1 lists information on SMSP and SharePoint LUN layout.

Table 1) Information on SMSP and SharePoint LUN layout.

Content	LUN	Description
SQL Server system databases	/vol/sql_Inst_Name_SystemDB/lunSQLSystemDB For master, model, and so on. /vol/sql_Inst_Name_TempDB/lunTem	Place the SQL Server system databases on a dedicated volume separate from the volume hosting the user databases. These databases are backed up using jobs

Content	LUN	Description
	pDB /vol/sql_Inst_Name_TempDB/lunTempDBLog For optimal performance, separate the TempDB data and log files into separate LUNs within the TempDB volume.	scheduled through SMSQL directly, and not SMSP. TempDB should not be included in a backup because the data it contains is temporary. Place tempdb on a LUN/SMB share that is in a storage system volume in which Snapshot copies are not created; otherwise, large amounts of valuable Snapshot copy space could be consumed.
SharePoint content databases	/vol/sql_Inst_Name_ContentDb/lunSPContentDB /vol/sql_Inst_Name_ContentDBLog/lunSPContentDBLog	These databases are backed up using SMSP. The layout of the content databases is determined by the RTO of the databases. When you place multiple databases on the same LUN, the restore of individual databases is performed through the SnapDrive sub-LUN restore feature.
SharePoint configuration database	/vol/sql_Inst_Name_ConfigDB/lunSPCoreDBs /vol/sql_Inst_Name_ConfigDBLog/lunSPCoreDBLogs	These are not very read/write intensive. Therefore, you can also choose to: • Store the SharePoint central admin databases and service application databases. • Host the SMSP control and archive databases. These databases can be backed up using SMSP by adding them as custom databases.
SMSP stub database	/vol/sql_Inst_Name_StubDb/lunSMSPStubDBs /vol/sql_Inst_Name_StubDb/lunSMSPStubDBLogs	The SMSP stub database is highly read-write intensive in a collaboration environment. Place the stub database and log on separate LUNs in its own volume. This allows you to host all of the stub databases created per web application within the SharePoint farm.
SnapInfo	/vol/sql_Inst_Name_SnapInfo/lunSnapInfo	This LUN is mapped to the SQL Server used in the SharePoint farm that stores backup metadata for SMSQL. Make sure that the databases residing on LUN/SMB shares within a volume are separate from that used by the SnapInfo volume to avoid stream-based backup. Instead, leverage the NetApp Snapshot technology.
Other databases	/vol/sql_Inst_Name_genDb/lunOtherDbs /vol/sql_Inst_Name_genLog/lunOtherDbLog	Databases for third-party-related apps are not related to SharePoint but are hosted on the SharePoint instance that can be backed up in SMSP using a custom database option.
SharePoint search index	/vol/sql_Inst_Name_SPSearchIndex/lunSPSearchIndex	This LUN is mapped to the SharePoint AS on which the search service application is provisioned and whose Snapshot copy is taken by SMSP using SDW.
SMSP 8.2 index (contains job metadata, SMSP	/vol/sql_Inst_Name_SMSPIndex/lunSMSPIndex	This LUN, VMDK, or CIFS volume is mapped to the storage configuration in the SMSP control panel and whose Snapshot copy is

Content	LUN	Description
backup index, web front end (WFE) IIS metadata backup data)		created by SDW.

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

Best Practices

- Verify that the SnapInfo LUN is not shared by any other type of data such as Windows OS and SharePoint binaries that could potentially corrupt the backup Snapshot copies.
- Verify that the SharePoint databases and SnapInfo LUNs are on separate volumes to avoid the retention policy from overwriting Snapshot copies, especially when used with NetApp SnapVault® technology.
- For clustered instances of SQL Server (FCI) of the SharePoint farm:
  - The SnapInfo LUN must be a cluster disk resource in the same cluster group as the SQL Server instance being backed up by SMSP.
  - Place SharePoint databases onto shared LUNs that are physical disk cluster resources assigned to the cluster group associated with the SQL Server instance.
- Verify that the storage virtual machine (SVM, formerly called Vserver) name is resolvable to the respective management LIF IP address, either by using the Domain Name System (DNS) or adding an entry into the Windows Server `etc\hosts` file. This enables SDW to create and display LUNs/SMB shares as expected and SMSQL to list them correctly.
- Disable automatic Snapshot copy scheduling configured through SDW.
- In order to enable SMSP data protection, the SharePoint content needs to reside on NetApp LUNs.
  - The SharePoint databases created use the model database as a template with specific configuration settings, such as file location, growth settings, and more, instead of getting created with default server–configured settings. By default, these databases are placed on the same volume/LUN as the SQL Server system databases. Use the SMSP database migrator tool to migrate these SharePoint databases to their respective LUNs on NetApp storage.

Note: The migrate database tool does not support the migration of the databases in the AlwaysOn AG and the SQL mirroring databases.

  - Similarly, use the SMSP index migrator tool to migrate the SharePoint search index to respective LUNs on NetApp storage.

Note: Do this work during nonbusiness hours, because SharePoint services are stopped during this migration process.

For complete details, refer to SnapManager 8.2 for Microsoft SharePoint Platform Backup and Restore User's Guide.

3.4 SMB Shares

With clustered Data ONTAP 8.2, support for the SMB 3.0 NAS protocol was introduced (a feature of Windows Server 2012). The SMB 3.0 protocol provides file-based access to SharePoint databases on NetApp CIFS shares.

Best Practices

- If you choose to place the SharePoint database on SMB shares, make sure that all the database files (.MDF and .LDF) of a SharePoint database reside on SMB shares, instead of placing them across LUNs and SMB shares.
- Configure the SDW transport protocol setting, where the SVM management LIF must be connected (by providing the SVM IP address, user name, and password) to view all of the SMB shares on its CIFS server, which then becomes visible to SMSQL.
- For SnapManager to be able to recognize the database file path as a valid file path hosted on NetApp storage, you must use the CIFS server name on the storage system in the SMB share path instead of the IP address of the management LIF or other data LIF. The path format is `\\<CIFS server name>\<share name>`. If the database uses the IP address in the share name, manually detach and attach the database by using the SMB share path with the CIFS server name in its share name.
- Avoid antivirus scanning on the SMB/CIFS shares where SharePoint BLOB is stored to avoid failed transactions resulting from scan delays.
- Make sure Windows host caching is disabled on the SMB/CIFS share where SharePoint data is located to avoid corruption due to caching.

When SharePoint data is hosted on NetApp storage, it is important to get a good understanding of the various storage efficiency methodologies used to store and manage this data in a way that consumes the least amount of storage space, with little or no effect on overall system performance. Storage efficiency goes beyond data deduplication; it is a combination of RAID, provisioning (overall layout and utilization), mirroring, and other data protection technologies. To understand further, refer to the section “Storage Efficiency and Manageability” in the Microsoft SharePoint and SnapManager 8.1 for SharePoint with Clustered Data ONTAP: Best Practices Guide

4 Sizing for SnapManager for SharePoint

4.1 SharePoint Server 2013 Planning Considerations

SharePoint farms vary in complexity and size; therefore, 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 decide the number of disks and type of disks depending on required I/O. Some of the many factors that need to be considered when planning a SharePoint environment to size it correctly include workload type, I/O operations per second (IOPS), requests per second (RPS), latency, read/write ratios, and working set size.

It is also important to have a well-thought-out information architecture (IA) and taxonomy, which go a long way in helping SharePoint to be more discoverable, logical, and manageable. When you have a 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. A range of capacity management and usage considerations affect how to determine the most appropriate hardware and configuration options.

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 configured and managed centrally that many web applications can consume.

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

The document Database types and descriptions (SharePoint 2013) details the databases created as part of the SharePoint deployment, based on the product version and edition. Each of the following databases has different requirements for:

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

When planning a SharePoint solution it is important to understand that all SharePoint databases are not created equal. Each of these factors affects the decisions made on types of disks to use; for example, placing TEMPDB on SSD or NetApp Flash Pool™ intelligent caching for performance versus placing BLOB content on lower-tier storage (SATA).

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

Best Practices

- Make sure that the `AUTO_CREATE_STATISTICS` option is off, because it is not supported for SharePoint and the required settings are automatically provided by SharePoint Server during provisioning and upgrade.
- Set the maximum degree of parallelism (MAXDOP) option to 1, in which a single SQL Server process serves each request, thereby confirming optimal query plans.
- Consider the following databases as Flash Pool candidates: TEMPDB, search, and usage.
- Configure the autogrowth value to a fixed number of megabytes versus percentage for safety reasons. This is to reduce the frequency with which SQL Server used by the SharePoint farm increases the size of a data file, because this blocking operation involves filling new space with empty pages. In addition, proactively monitor and manage the growth of the data and log files. For further details, refer to Considerations for the "autogrow" and "autoshrink" settings in SQL Server.
- Document management sites have a database priority for faster disks, as follows:
  - TEMPDB (mdf and ldf)
  - Content database (ldf)
  - Search databases
  - Content database (mdf)
- Read-oriented publishing portal sites have a database priority for faster disks, as follows:
  - TEMPDB (mdf and ldf)
  - Content database (mdf)
  - Search databases
  - Content database (ldf)

For additional information refer to:

- Database types and descriptions (SharePoint Foundation 2010)
- SharePoint Infrastructure planning and design process

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

4.2 Sizing the Control Database

The SMSP Manager hosting control service connects to one control database that contains the SMSP configuration data and backup plans, storage optimization (Storage Manager, Archive Manager, and connector) rules, and job records. The data growth rate on the control database is relatively small, and, with retention on jobs, the job record can be automatically pruned from the control database.

Best Practices

- Always define retention rules for backup and archive data in the storage policy to confirm that the job record is pruned from the control database.
- In case you are manually deleting the backup job, make sure to delete the job and backup data.
- NetApp highly recommends configuring a job-pruning policy if you run backups frequently, to make sure that the control database is not overloaded with job data.
- Change the recovery model for the SMSP control database to full because it is changed frequently, causing the log size to grow.

4.3 Estimating Backup Data Size

Depending on the backup options in your backup plan, the following form a major part of the backup data index stored on the storage policy device:

- The catalog file for the backup job saved on the backup device LUN is small, typically less than 10MB. The backup retention policy defined in the storage policy decides the number of backup jobs for the backup plan.
- If the backup plan includes backing up SharePoint Server (WFE/APP) backup data, this data is streamed to the backup device LUN; each backup includes the SharePoint hive, global assembly cache, web parts, IIS metadata, site definition, and custom solutions. This size can vary depending on how many custom SharePoint solutions are deployed from third-party ISV developers or internal development efforts; it is typically less than 10GB.
- The SharePoint 2010 FAST Search server data. If FAST server is selected in the backup plan and the FAST program files data is not installed on the NetApp LUN, the backup of FAST program data copies to the storage policy device. The backup data size contains the FAST program files folder and the configuration files folder.
- The number of items in the SharePoint content database. For example, one document item takes about 1Kb for the index; if the documents have multiple versions, each version takes approximately 300 bytes.
- If BLOB backup is selected in the backup plan, the index of BLOB is also saved to the backup device LUN.
- The granular index level for content database selected. In a normal SharePoint web application, as the level of granularity becomes finer, the number of objects that must be indexed increases and therefore the size of the index increases. Normally, it is difficult to get a count of the number of objects at each level of granularity, which makes sizing the index very complex and difficult.

According to Storage and SQL Server capacity planning and configuration (SharePoint Server 2013), the formula to estimate the content database size is:

```
Database size = ((D × V) × S) + (10 KB × (L + (V × D)))
```

Here:

- D: Calculate the expected number of documents.

- S: Estimate the average size of the documents that you will store.
- L: Estimate the number of list items in the environment.
- V: Determine the approximate number of versions.

The above formula can be reorganized as:

```
D = (SDB - (10 x L))/ (V x S+ 10 x V) ≈ SDB/(V x S)
```

Here the average document size is much bigger than 10KB.

Based on this, the SMSP-generated granular index data size (in item version level) for one database in one backup job can be estimated as:

```
Index size(KByte) = 2.5*(Nsc*53 + Ns*27 + L*0.45 + D*V*0.35)
```

Here:

- N_{sc}: Number of site collections in content DB
- N_{s}: Number of subsites in content DB
- L: Number of list/folders in content DB
- D: Number of items
- V: Average number of versions for documents

If you don't have the total number of items, you can also use the database size SDB (KB) to estimate the size of the index in SMSP as:

```
Index size(KByte) = 1.5*(Nsc*53 + Ns*27 + L*0.45 + (SDB/S)*0.35)
```

The preceding calculation is an estimation of one content database in one backup job.

If you have many content databases and know the total data size **Dts (GB)** of all content databases, you can also use it to estimate the total index size for one SMSP backup job as follows:

```
IndexSize(MB)= 1.5*( 53*Nsc+ 27*Ns + 0.45*(Dts*1024*1024/S)/L + 0.35*(Dts*1024*1024/S) )/1024
```

If the backup job is run with granular index enabled, the index data will be stored on the media service. The size of the index data is related to the number of items in the content database.

Based on the test results:

```
Estimate index size (MB) = 1.5*(53*Nsc + 27*Ns + 0.45*(D*1,024*1,024/S)/L +
0.35*(D*1,024*1,024/S))/1,024
```

Here:

- N_{sc}: Number of site collections in content DB
- N_{s}: Number of subsites in content DB
- L: Number of lists/folders in content DB
- D: Total data size
- S: Average document size

The SMSP backup does not save the database content itself to the backup device LUN.

4.4 Sizing the Stub Database

The WFE servers are responsible for processing all RBS processes using the SMSP RBS provider installed on the different WFE servers. The SMSP RBS provider creates records in the stub database that correlate the content on NetApp SMB (CIFS) with the content contained within the SharePoint content database. The SharePoint content database contains only the RBS auxiliary table containing the BLOB

ID; the stub DB contains the information of the RBS BLOB storage and how the BLOB ID is mapped to the real BLOB storage location. The SMSP stub database keeps a record of each BLOB, with each BLOB record using about 300 bytes. Stub database size increases with the increase in documents uploaded or synchronized with the connector library. The stub database size is small (for example, for a content database with a maximum of 60 million document BLOBs, the stub database size is less than 20GB).

Best Practices
• Reserve enough space for backup granular index data to hold all data of various backup plans based on retention. • NetApp recommends placing the media service on a dedicated physical host or virtual machine to cope with the additional processing power needed for managing the backup job data (metadata and index). • For largely distributed deployments, NetApp recommends deploying media service 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 because of hardware failure. • NetApp does not recommend installing the media service on WFE for security, monitoring, and scalability purposes. • The media service cache is used for granular index generating buffer space. In the storage policy, use a LUN if the media service is a single node. When using media service high availability, use a CIFS device.

4.5 Estimation of Archive Data Size

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

Best Practice
Verify that SMSP archive databases are included in the backup and added as custom databases.

5 Performance

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

Best Practices

- Leverage the SMSP storage optimization modules to externalize BLOB data onto less expensive NetApp CIFS shares to help increase SQL Server performance by offloading write-intensive operations.
- An SMSP synchronization job is fairly resource intensive; therefore, running multiple synchronization jobs simultaneously might affect the performance of the server on which the control service is installed. To avoid this condition, configure the SMSP processing pool in which synchronization jobs that are added to the processing pool become threads. The number of jobs you allow in the processing pool is the maximum number of synchronization jobs that can be run simultaneously; the remaining jobs will be put into a queue.
- If the media service is virtualized, make sure that you have sufficient memory size and CPU power.
- When externalizing BLOBs to NetApp CIFS shares, add NetApp Flash Cache technology to improve controller performance for random-read workloads.
- There are certain “by design” SharePoint limits that cannot be exceeded and some whose default values might be changed by the farm administrator. Make sure that you operate within established limits because acceptable performance and reliability targets are best achieved when a SharePoint farm’s design provides a reasonable balance of limit values. This also aids the manageability of the SharePoint farm.

  For a comprehensive list of limits, refer to Software boundaries and limits for SharePoint 2013.

6 NetApp Solution for Microsoft SharePoint 2013

When planning the backup and restore of a SharePoint farm, the following objectives must be clearly defined according to customer SLAs:

- **Recovery point objective (RPO)**

  To what point in time must the data be recovered?
- **Recovery time objective (RTO)**

  How long will it take to get the database back online and rolled forward or backward to the RPO?

6.1 SnapManager 8.2 for SharePoint Overview

SMSP is an enterprise-strength backup, recovery, and data management solution for SharePoint Server. The combination of NetApp storage solutions and Microsoft SharePoint enables the creation of enterprise-level database storage designs that can meet today’s most demanding application requirements.

For more details refer to SnapManager 8.2 for Microsoft SharePoint.

Table 2 lists the SnapManager 8.2 for SharePoint components mapped to SharePoint farm hosts.

Table 2) SnapManager 8.2 for SharePoint components mapped to SharePoint farm hosts.

Server Role	SMSP 8.2 Component	Remarks
Server that is not part of the SharePoint farm	SMSP manager	Mandatory.
WFE and AS	SMSP agent	Mandatory.

Server Role	SMSP 8.2 Component	Remarks
WFE and AS	SMSP storage optimization modules such as: 1. Storage Manager 2. Connector 3. Archive Manager The reason for the storage optimization modules to be installed on WFE and AS is so that SMSP can scale in the complexity of the environment, which also allows the execution of jobs and processes to be split accordingly. Hence, when there is noncompliance, NetApp SnapManager for SharePoint Manager (SMSP) generates error messages and job failure might occur.	1. Storage Manager is optional. It is enabled for performing stub-based uploads of documents for one or more of the web applications hosted on the WFE. 2. Connector is optional. It is enabled to collaborate on NetApp CIFS and NFS shares and cloud storage resources directly in SharePoint without ingesting content. 3. Archive Manager is optional. It is enabled to archive SharePoint content to a tiered storage system for long-term retention.
SharePoint index server	SMSP agent	Mandatory. Installed for backing up the SharePoint search indexes.
SQL Server host	SMSP agent	Mandatory.

Figure 4 represents the various SnapManager 8.2 for SharePoint components that are mapped to SharePoint hosts in the farm.

Figure 4) SnapManager for SharePoint (SMSP) data protection solution.

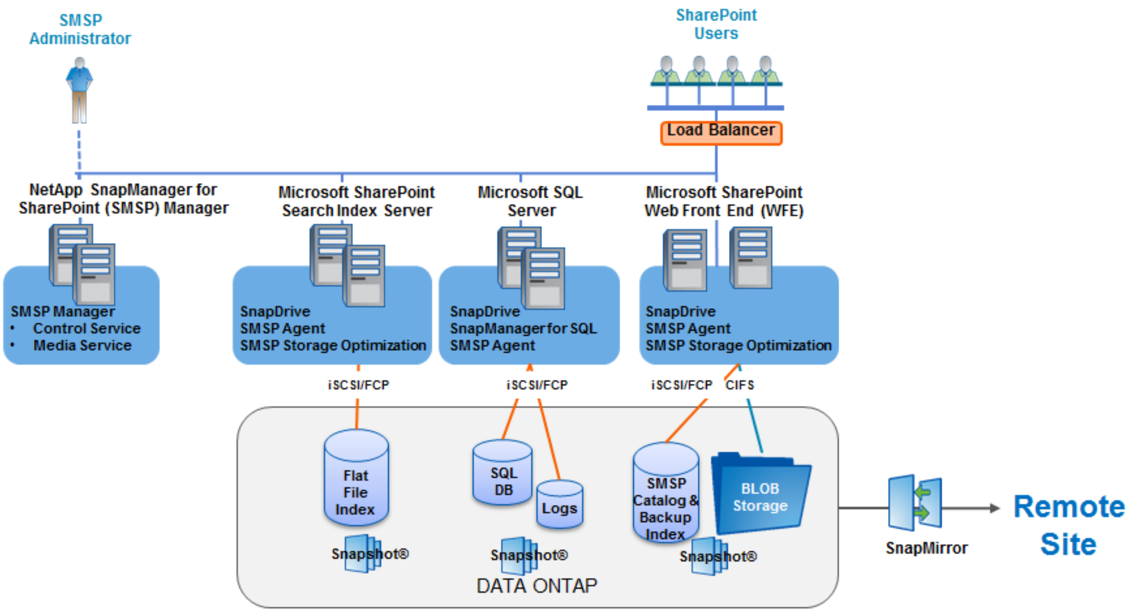

Note: Installing an SMSP agent on a SQL Server that runs Windows Server core is not supported by SMSP and SMSQL.

Best Practices

- When using customized ports for SMSP, confirm that those ports are available and not blocked by antivirus technology. If multiple SMSP services are installed on the same server, make sure that the required ports are enabled on that server.
- Confirm that the SVM name is added in the DNS that needs to resolve to the management LIF.
- Synchronize the system clock on the host running SnapManager with the clock on the storage system to confirm that SDW functions correctly.
- Use the SMSP Health Analyzer to verify that the necessary prerequisites for system, permissions, and others to use SMSP have been met.
- The SMSP Health Analyzer scans the SharePoint farm according to rules selected in the health analyzer profiles to report on any issues that might affect SMSP modules. Therefore, verify that the user account with which you run the Health Analyzer belongs to the following groups:
  - SMSP administrators group
  - SharePoint farm administrators group
  - Local administrators group on each server in the SharePoint farm
- Make sure that the SMSP manager is accessible by all SMSP agent servers.
- Make sure that the SMSP agent account is added to the SharePoint Farm administrators group and that full control is given on the web application to enumerate all SharePoint content through the SMSP manager to allow a backup.
- When specifying a UNC path while creating an SMB share, use IP addresses instead of host names. This is particularly important with iSCSI because host-to-IP name resolution issues can interfere with locating and mounting the iSCSI LUNs during the boot process.
- Confirm that you have SDW and SMSQL installed on all nodes of the SQL Server failover cluster instance (FCI). When using SQL AG, make sure that both are installed on the server replicas selected for backup.
- The Snapshot copy verification process is CPU intensive and degrades SQL Server performance. Therefore, configure a SQL Server instance that is not used by the SharePoint farm to run these database verification operations and schedule them to run during peak usage hours.

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

6.2 Backup Guidelines

The criticality matrix shown in Figure 5 provides a conceptual strategy on how to plan and schedule backups of various kinds of SharePoint data. The best situation is to be able to back up the whole farm every day. However, as the SharePoint farm grows you might need to schedule separate backups based on the change frequency of the data. For example, Search Help Index and WFE servers do not need to be backed up every day, so it is reasonable to schedule those backups weekly (or even monthly). Using another example, MySites data might not be as critical as other content, so you can back up that data less frequently (assuming you followed the best practice and gave MySites a set of separate content databases).

Figure 5) Backup planning guidance provided by the criticality matrix.

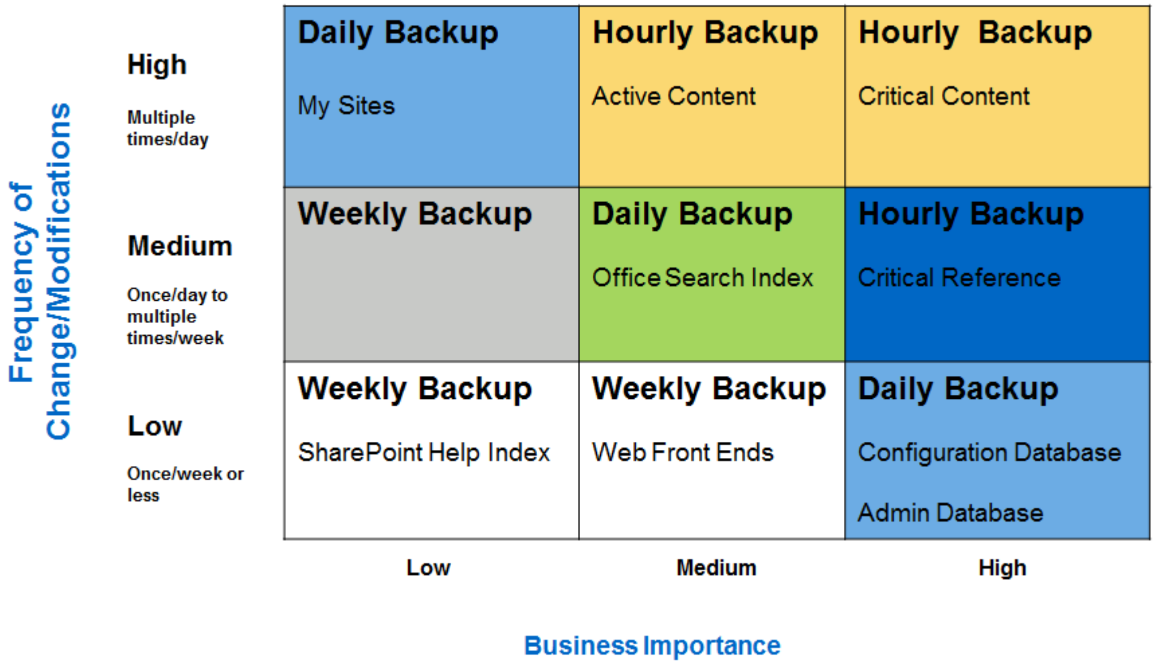

Frequency of Change/Modifications	Low	Medium	High
High Multiple times/day	**Daily Backup** My Sites	**Hourly Backup** Active Content	**Hourly Backup** Critical Content
Medium Once/day to multiple times/week	**Weekly Backup**	**Daily Backup** Office Search Index	**Hourly Backup** Critical Reference
Low Once/week or less	**Weekly Backup** SharePoint Help Index	**Weekly Backup** Web Front Ends	**Daily Backup** Configuration Database Admin Database

Business Importance

Note: This criticality matrix is only an example. Businesses need to first organize their content appropriately and then categorize based on their SLAs, RTO, RPO, and so on.

During SharePoint farm backup, SMSP works as a Volume Shadow-Copy Service (VSS) requestor to start a VSS session and uses the SharePoint Foundation VSS reference writer to query the VSS components (databases and search index) that need to be backed up. The SPF-VSS reference VSS writer in turn simply references SQL VSS writer and search VSS writer. SMSQL leverages them to perform database Snapshot backup copies of SharePoint databases and SDW for search index Snapshot copy backup. The backup data is then sent to the configured storage policy and stored with the backup job metadata and index.

The SMSP backup data typically contains the following types of data:

- SharePoint database backup created using SMSQL
- SharePoint search index data
- Externalized BLOB data
- Storage policy device metadata for backup jobs

End-Process Management in SMSP 8.2

When trying to stop jobs in previous releases of SMSP by configuring the file "`C:\Program Files\AvePoint\DocAve6\Manager\Control\Config\controlcastle.config` on the control service host, users had to wait 30 minutes for an unresponsive agent to set a job to "failed." Also, the act of stopping a job does not remove the backup data already created (or the Snapshot copies). Hence, in SMSP 8.2, if you are aware that the job has failed or the agent is offline, you can stop the job on demand and reset its status from the job monitor without having to wait the full 30 minutes, as shown in Figure 6. This feature allows execution of subsequent pending jobs. When the SMSP agent connects to the SMSP manager again, it kills the respective process on the agent server. If the SMSP agent does not reconnect with the SMSP manager after stopping the job, the user has to reset the process manually using the task manager on the agent server.

Figure 6) SMSP 8.2 force stop.

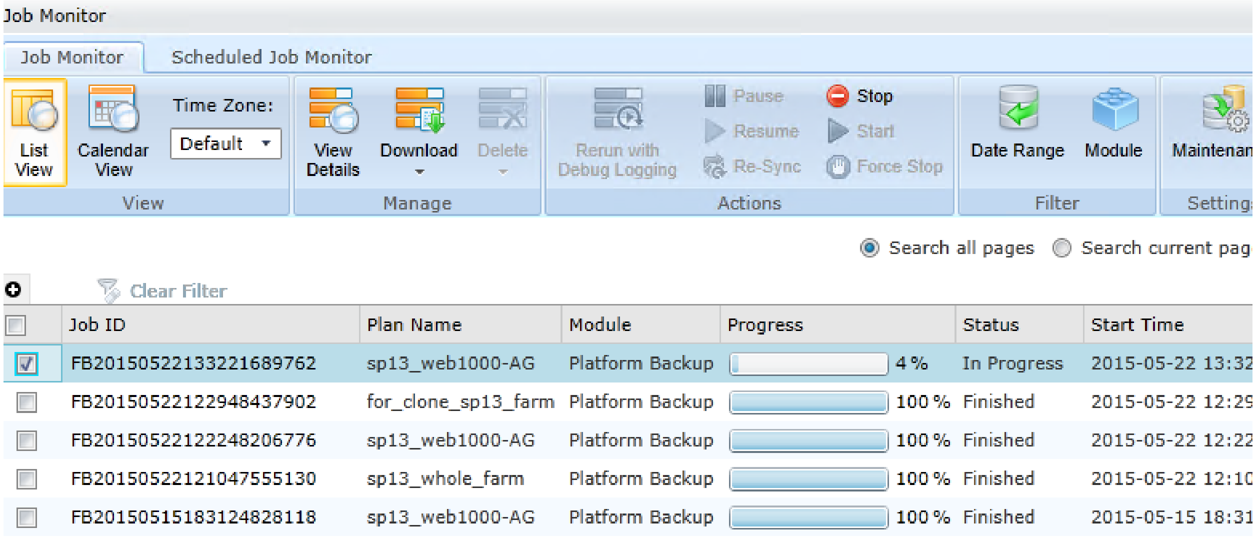

□	Job ID	Plan Name	Module	Progress	Status	Start Time
☑	FB20150522133221689762	sp13_web1000-AG	Platform Backup	4 %	In Progress	2015-05-22 13:32
□	FB20150522122948437902	for_clone_sp13_farm	Platform Backup	100 %	Finished	2015-05-22 12:29
□	FB20150522122248206776	sp13_web1000-AG	Platform Backup	100 %	Finished	2015-05-22 12:22
□	FB20150522121047555130	sp13_whole_farm	Platform Backup	100 %	Finished	2015-05-22 12:10
□	FB20150515183124828118	sp13_web1000-AG	Platform Backup	100 %	Finished	2015-05-15 18:31

Prerequisites

- SMSP requires that the SharePoint databases reside on the NetApp LUN and BLOBs created by SMSP Storage Optimization modules on NetApp CIFS shares to leverage NetApp Snapshot technology. Make sure to also add the stub and archive databases in SMSP full farm backups as custom databases. This enables a complete full farm recovery from SnapMirror Snapshot copies during a disaster.
- Use the same SMSP control database when reinstalling the new SnapManager for SharePoint Manager in the disaster recovery site.
- SMSP requires that the NetApp controller user account be able to log in to the Data ONTAP storage system and perform the following operations:
  - Query/list SVM, CIFS shares, volumes.
  - Create Snapshot copies for CIFS volumes.
  - Perform SnapMirror and SnapVault operations (query, update, and so on).
  - Expose Snapshot copy as CIFS share.

Store Backup Data and Index to Nonroot Folders

In previous SMSP releases, you could not specify a start folder for a LUN device, and SMSP always created its data folder on the root of the LUN device. Customer IT policy might require that the application data folder be placed under a folder and be assigned to a certain department/group with different ACLs. With SMSP 8.2, you can alter the default location to select a specific folder in which to store SMSP backup data using the “Folder Name” parameter for the LUN device, as shown in Figure 7. You can do this rather than use the root location of the LUN when LUN backup devices are selected in the physical device wizard.

Figure 7) Specify folder name for LUN physical device.

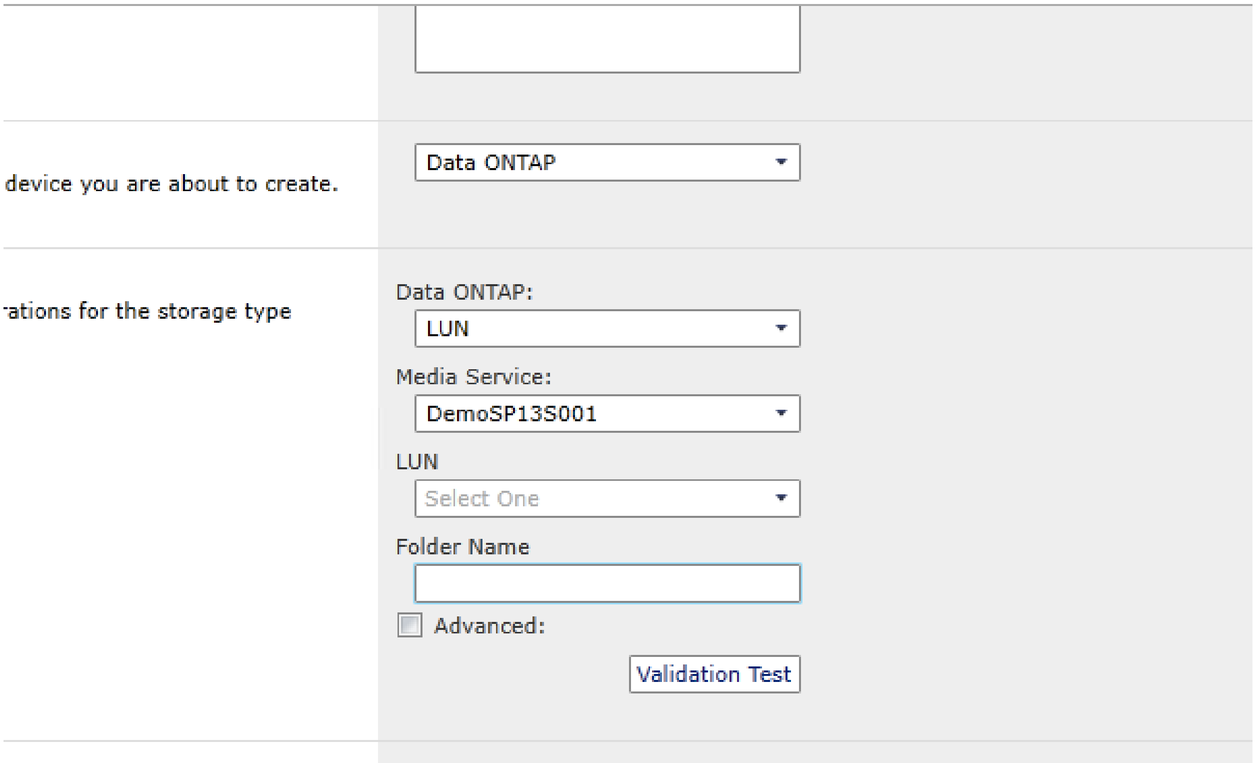

device you are about to create.

Data ONTAP

ations for the storage type

Data ONTAP: LUN

Media Service: DemoSP13S001

LUN: Select One

Folder Name:

□ Advanced:

SMSP Backup Using AlwaysOn Availability Group Replica

Previously, when SMSP found the database in a SQL AG and the AG listener was used for a SharePoint farm, SMSP used SQL query to find the backup “preferred replica” server and then ran the `export-config` command to locate the database to verify if it resided on a NetApp LUN. SMSP backup of SQL AG executed the SMSQL backup job on the “preferred backup” replica or the primary replica and did not support SMSQL’s ability to back up AlwaysOn databases in different ways. When a secondary replica went down, that replica indicated its state as not synchronized and that the other replica, including the primary replica, still worked well. If the primary replica goes down and the AG has automatic failover configured, failover will happen and another replica will become the primary replica. This will need the SQL administrator’s involvement to manually fail over or fix the primary replica.

When creating an SMSP 8.2 backup containing an AG with SMSQL installed, as shown in Figure 8, users will be able to select:

- Back up only on the primary replica
- Back up on the secondary replica
- Back up on all replicas

Figure 8) SMSP backup replica settings.

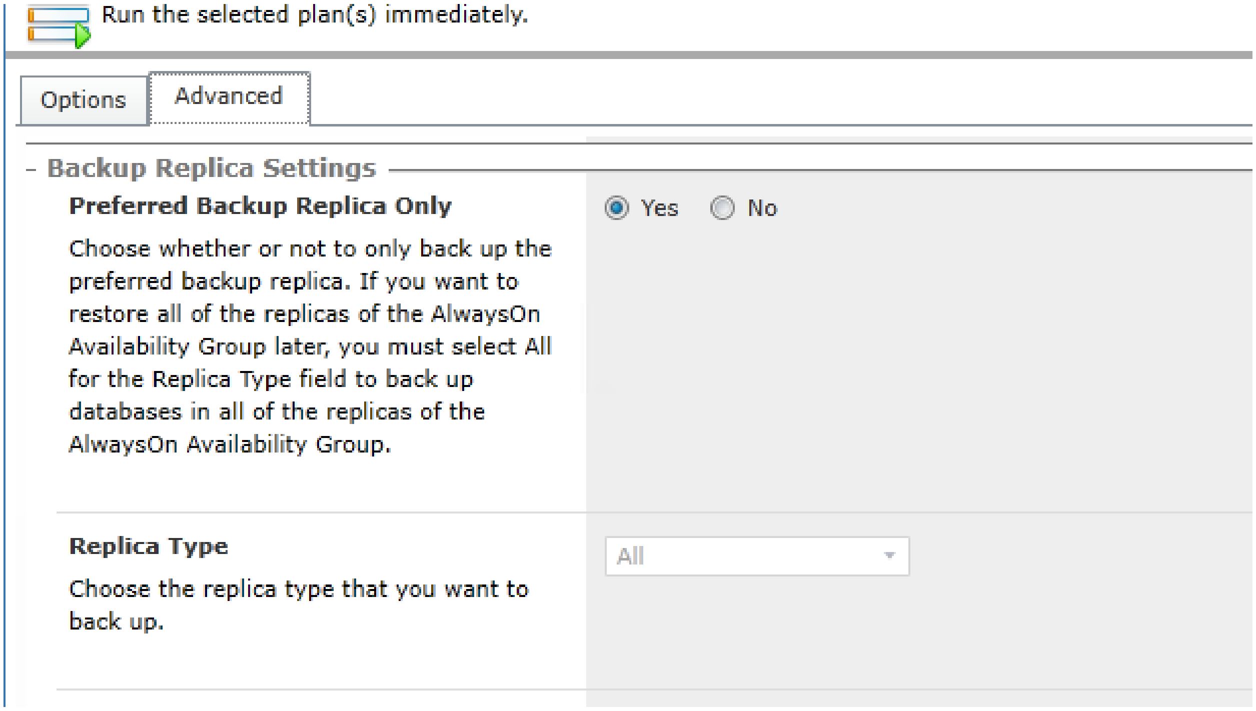

Run the selected plan(s) immediately.

- Backup Replica Settings

Preferred Backup Replica Only

Choose whether or not to only back up the preferred backup replica. If you want to restore all of the replicas of the AlwaysOn Availability Group later, you must select All for the Replica Type field to back up databases in all of the replicas of the AlwaysOn Availability Group.

● Yes ○ No

Replica Type

Choose the replica type that you want to back up.

All

Note: Since there might be multiple AGs and replicas (up to nine in SQL Server 2014), SMSP does not allow selection of specific replica settings for individual AGs in the backup plan. SMSP replicates how SMSQL selects secondary replica for backup; that is, it either selects all secondary replicas or selects the secondary replicas based on the backup priority range.

Verification Group

In SMSP 8.1, all database verification was sequential. SMSP 8.2 supports the capability of SMSQL 7.2 to register multiple verify servers and use it as a group to run verify jobs to perform parallel database verification on the same verification server or server group, as shown in Figure 9. This setting enables the verification of backups created on different machines in parallel on multiple verification servers, thereby reducing verification time.

Figure 9) Verification server group.

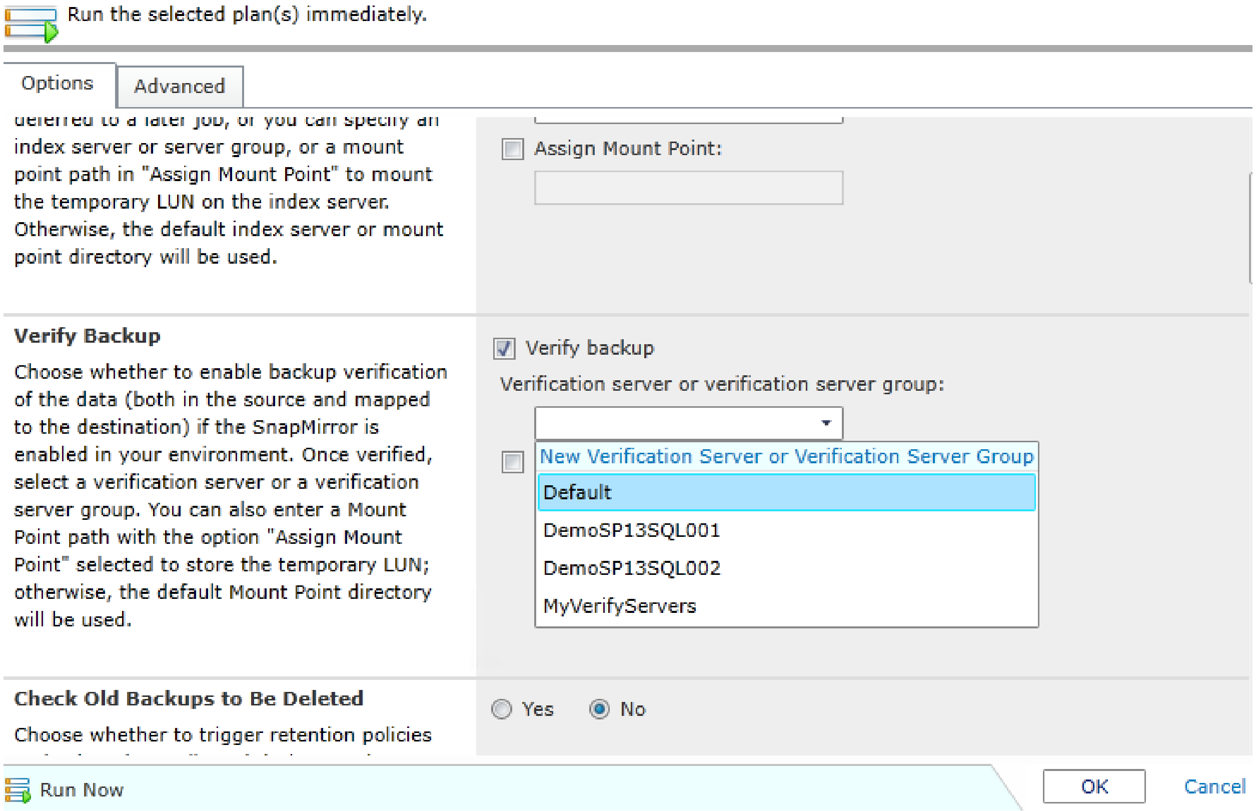

Run the selected plan(s) immediately.

deferred to a later job, or you can specify an index server or server group, or a mount point path in "Assign Mount Point" to mount the temporary LUN on the index server. Otherwise, the default index server or mount point directory will be used.

□ Assign Mount Point:

Verify Backup

Choose whether to enable backup verification of the data (both in the source and mapped to the destination) if the SnapMirror is enabled in your environment. Once verified, select a verification server or a verification server group. You can also enter a Mount Point path with the option "Assign Mount Point" selected to store the temporary LUN; otherwise, the default Mount Point directory will be used.

☑ Verify backup

Verification server or verification server group:

□

Check Old Backups to Be Deleted

Choose whether to trigger retention policies

○ Yes ● No

Best Practices

- If you have SharePoint Servers in the DMZ, create a custom agent group that excludes the SMSP agents installed on these servers. This confirms that the backup operation does not connect to these servers, avoiding blocked access because of firewall restrictions that cause timeouts during backup.
- When backing up a SharePoint content database with the BLOB provider configured, confirm that the respective provider is enabled on each SQL Server instance of FCI or AG with necessary permissions.
- Leverage SMSP verification (at the end of a backup job or deferred verification) to verify that consistent databases are created during backup.
- Make sure that you periodically use SnapMirror to copy the following volumes for disaster recovery purposes:
  - SMSP backup Snapshot copies containing the SharePoint content databases and search index data
  - NetApp CIFS shares containing externalized BLOB data
  - The SMSP control database used by SMSP manager
  - The stub DB used by the SMSP agent for BLOB access
- In a mirrored setup, if SQL Server authentication is the SharePoint content database authentication, make sure that 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.
- Confirm that the recovery model for the control database, archive database, and stub database is changed from simple to full.
- In the case of stub database migration, set the site collections to read only so that no new BLOB record is created during this migration.
- NetApp recommends not using too many separate CIFS volumes to run multiple jobs using different profiles to externalize site collection BLOB to different CIFS shares. The backup job will take longer to find the respective volumes and created backup Snapshot copies with Data ONTAP operating system cmdlets.
- Make sure that the AG replica is configured with “readable secondary”=yes. Otherwise, SMSP will not be able to get the correct status of the database to complete the backup operation. Also, because the user might configure all secondary replicas with equal backup priority or use “any replica” as a backup preference, make sure to configure all secondary replicas as “readable secondary”=yes.

For additional information, refer to the SnapManager 8.2 for Microsoft SharePoint Platform Backup and Restore User's Guide.

6.3 Restore Guidelines

SMSP provides the flexibility to restore an entire SharePoint farm, site collection, sites, subsites, individual documents, and document versions as needed, all within minutes.

Full Farm Clone Restore Operation

Previous versions of SMSP could perform a clone restore of a web application, including content databases and BLOB data, to a new web application for the purpose of creating test/developments or disaster recovery environments. The new SMSP 8.2 Farm Clone Wizard, as shown in Figure 10, leverages full farm backup data (including WFE backup) to create a point-in-time temporary farm, created leveraging split-clone (if using SMSQL 7.2 in both farms) or standard FlexClone technology. After cloning, both farms will be active and online, and registered to the same SMSP manager with externalized BLOB data working as expected. However, no plans, profiles, or jobs will be carried from the production farm to the temporary farm because these must be recreated in the temporary farm. It will appear as a new farm with a different name in the backup wizard interface. The new farm created by Farm Clone does not need the same topology as the original farm, so it can have a different number of servers and a different host name.

Figure 10) SMSP 8.2 Farm Clone.

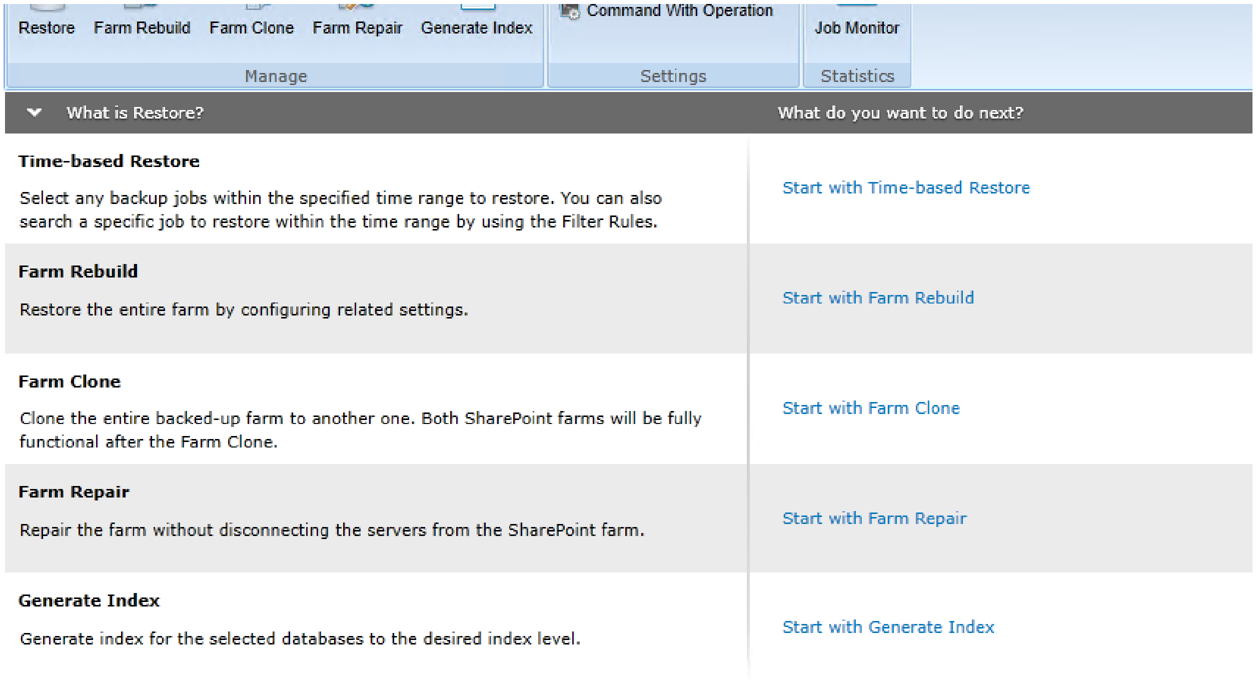

Time-based Restore

Select any backup jobs within the specified time range to restore. You can also search a specific job to restore within the time range by using the Filter Rules.

Start with Time-based Restore

Farm Rebuild

Restore the entire farm by configuring related settings.

Start with Farm Rebuild

Farm Clone

Clone the entire backed-up farm to another one. Both SharePoint farms will be fully functional after the Farm Clone.

Start with Farm Clone

Farm Repair

Repair the farm without disconnecting the servers from the SharePoint farm.

Start with Farm Repair

Generate Index

Generate index for the selected databases to the desired index level.

Start with Generate Index

Figure 11 shows the option to rebuild a temporary farm from an SMSP full farm backup.

Figure 11) Full Farm Clone to temporary farm.

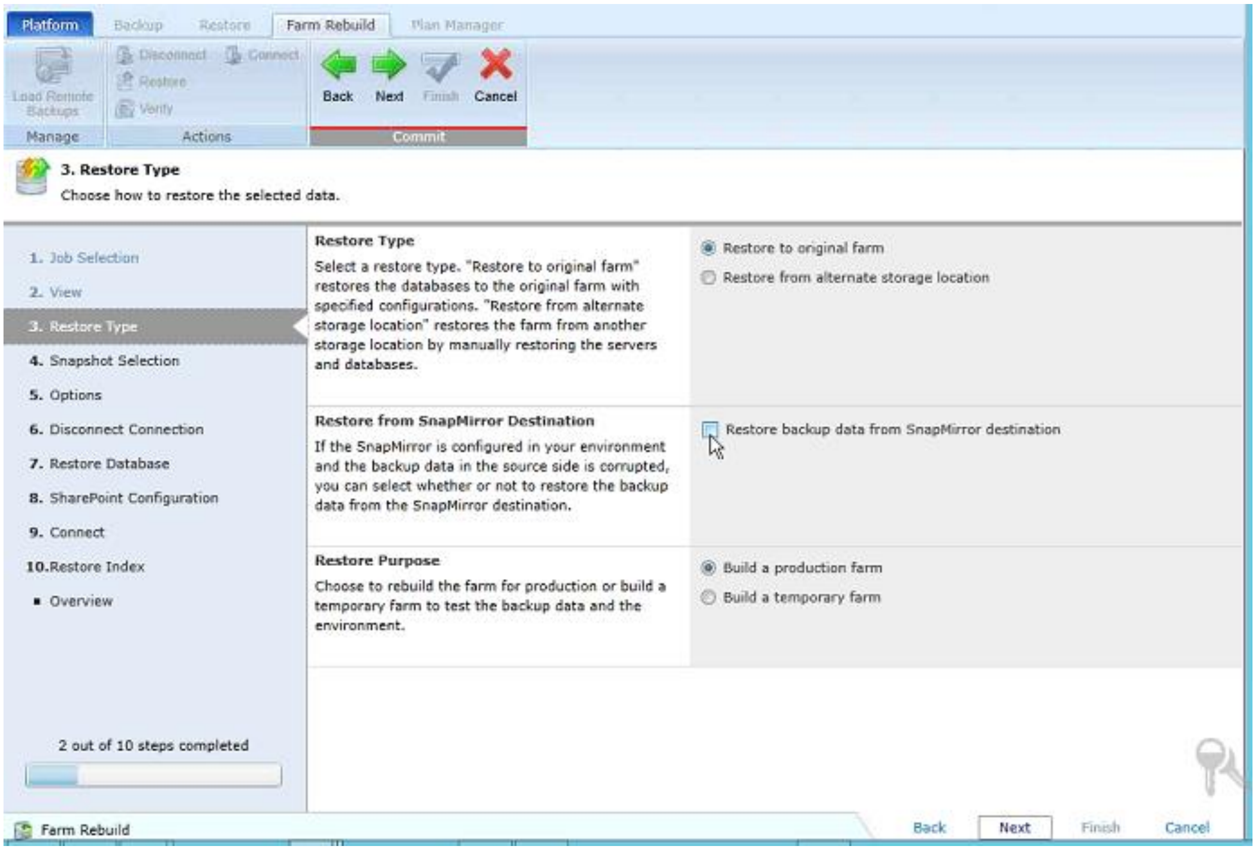

3. Restore Type

Choose how to restore the selected data.

1. Job Selection
2. View
3. Restore Type
4. Snapshot Selection
5. Options
6. Disconnect Connection
7. Restore Database
8. SharePoint Configuration
9. Connect
10. Restore Index

- Overview

2 out of 10 steps completed

Restore Type

Select a restore type. "Restore to original farm" restores the databases to the original farm with specified configurations. "Restore from alternate storage location" restores the farm from another storage location by manually restoring the servers and databases.

● Restore to original farm
○ Restore from alternate storage location

Restore from SnapMirror Destination

If the SnapMirror is configured in your environment and the backup data in the source side is corrupted, you can select whether or not to restore the backup data from the SnapMirror destination.

□ Restore backup data from SnapMirror destination

Restore Purpose

Choose to rebuild the farm for production or build a temporary farm to test the backup data and the environment.

● Build a production farm
○ Build a temporary farm

An SMSP 8.1 farm rebuild from SnapMirror requires you to manually break the current SnapMirror relationship and manually mount the volume/LUN at the SnapMirror disaster recovery destination site prior to the rebuild operations. However, with SMSP 8.2, the wizard saves the SnapMirror relationship information, automatically breaks the SnapMirror relationship, and mounts the volume/LUN, as shown in Figure 12.

Note: This does not support the cloning of a virtual machine (VM) within a farm because both farms will be active after the clone has completed.

This temporary farm topology is derived from production topology with the number of servers, server names, and device names changed to accommodate smaller QA or development farms.

The temporary farm includes all of the following as a result of the full Farm Clone operation:

- The configuration database is included with the farm ID metadata modified in the database to make sure that there is no duplication with the original active farm.
- All web applications and content databases.
- All BLOB data is created by the storage manager with stub databases with new devices and device IDs to make sure that there is no overwrite of the original BLOBs created by the production farm.

  Note: Connector data cloning is currently not supported.
- Service applications such as Search Service Application, User Profile Service, Managed Metadata Service, and so on.
- Additional third-party services that require customizations.

This requires that SharePoint is already deployed, installed, and configured in the temporary farm, with all changes made to settings, configurations, and customizations deployed manually.

Note: The SharePoint version and patches must match those of the production farm, which is confirmed by a validation test in the full Farm Clone Wizard.

Figure 12) SQL Server mapping for full Farm Clone operation.

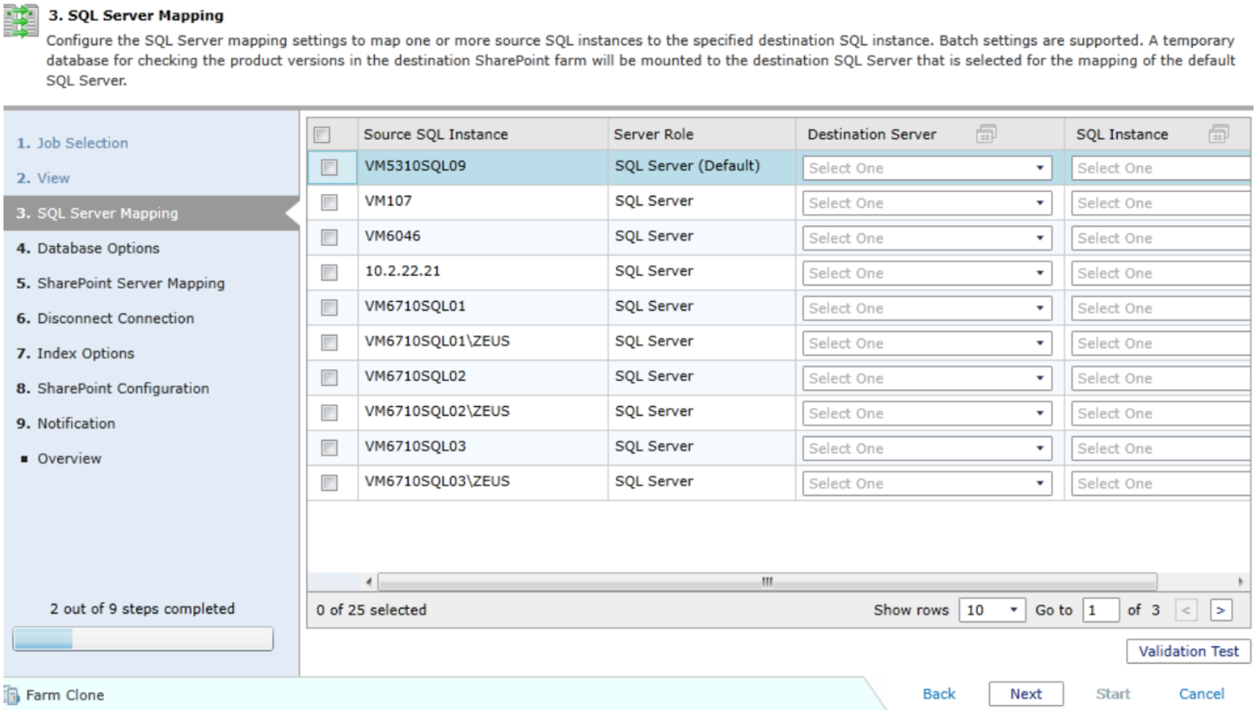

3. SQL Server Mapping

Configure the SQL Server mapping settings to map one or more source SQL instances to the specified destination SQL instance. Batch settings are supported. A temporary database for checking the product versions in the destination SharePoint farm will be mounted to the destination SQL Server that is selected for the mapping of the default SQL Server.

□	Source SQL Instance	Server Role	Destination Server	SQL Instance
□	VM5310SQL09	SQL Server (Default)	Select One	Select One
□	VM107	SQL Server	Select One	Select One
□	VM6046	SQL Server	Select One	Select One
□	10.2.22.21	SQL Server	Select One	Select One
□	VM6710SQL01	SQL Server	Select One	Select One
□	VM6710SQL01\ZEUS	SQL Server	Select One	Select One
□	VM6710SQL02	SQL Server	Select One	Select One
□	VM6710SQL02\ZEUS	SQL Server	Select One	Select One
□	VM6710SQL03	SQL Server	Select One	Select One
□	VM6710SQL03\ZEUS	SQL Server	Select One	Select One

0 of 25 selected

Show rows 10 Go to 1 of 3

2 out of 9 steps completed

Figure 13) SharePoint Server mapping for full Farm Clone operation.

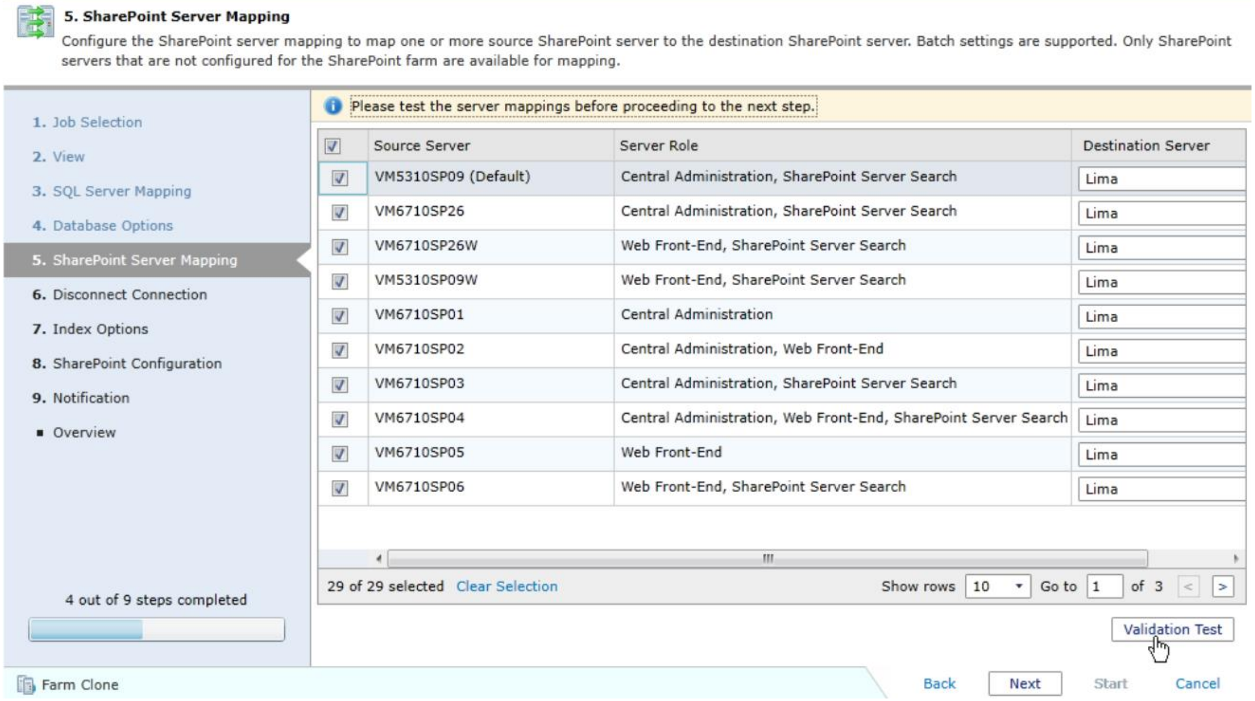

5. SharePoint Server Mapping

Configure the SharePoint server mapping to map one or more source SharePoint server to the destination SharePoint server. Batch settings are supported. Only SharePoint servers that are not configured for the SharePoint farm are available for mapping.

Please test the server mappings before proceeding to the next step.

☑	Source Server	Server Role	Destination Server
☑	VM5310SP09 (Default)	Central Administration, SharePoint Server Search	Lima
☑	VM6710SP26	Central Administration, SharePoint Server Search	Lima
☑	VM6710SP26W	Web Front-End, SharePoint Server Search	Lima
☑	VM5310SP09W	Web Front-End, SharePoint Server Search	Lima
☑	VM6710SP01	Central Administration	Lima
☑	VM6710SP02	Central Administration, Web Front-End	Lima
☑	VM6710SP03	Central Administration, SharePoint Server Search	Lima
☑	VM6710SP04	Central Administration, Web Front-End, SharePoint Server Search	Lima
☑	VM6710SP05	Web Front-End	Lima
☑	VM6710SP06	Web Front-End, SharePoint Server Search	Lima

29 of 29 selected Clear Selection

Show rows 10 Go to 1 of 3

4 out of 9 steps completed

Figure 14) Index server mapping for full Farm Clone operation.

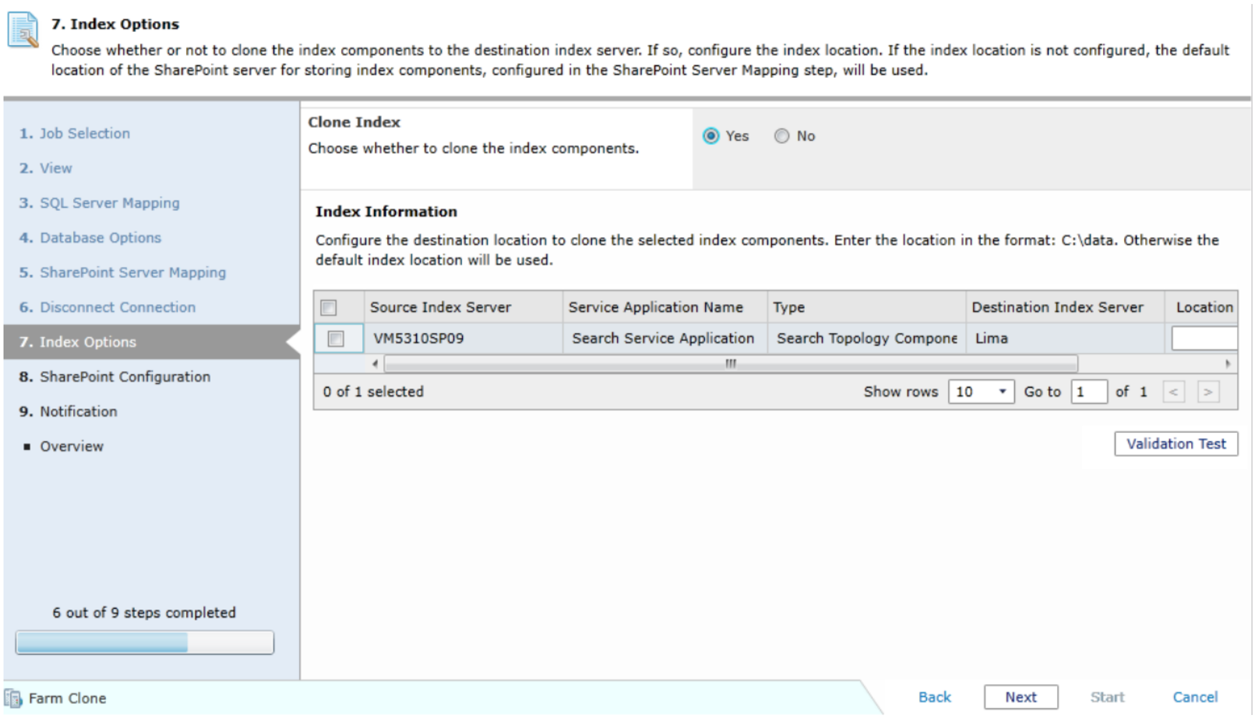

7. Index Options

Choose whether or not to clone the index components to the destination index server. If so, configure the index location. If the index location is not configured, the default location of the SharePoint server for storing index components, configured in the SharePoint Server Mapping step, will be used.

1. Job Selection
2. View
3. SQL Server Mapping
4. Database Options
5. SharePoint Server Mapping
6. Disconnect Connection
7. Index Options
8. SharePoint Configuration
9. Notification
- Overview

6 out of 9 steps completed

Clone Index

Choose whether to clone the index components.

● Yes ○ No

Index Information

Configure the destination location to clone the selected index components. Enter the location in the format: C:\data. Otherwise the default index location will be used.

□	Source Index Server	Service Application Name	Type	Destination Index Server	Location
□	VM5310SP09	Search Service Application	Search Topology Compone	Lima	

0 of 1 selected Show rows 10 Go to 1 of 1

Validation Test

Farm Clone

Back Next Start Cancel

Figure 15) BLOB device mapping in full Farm Clone operation.

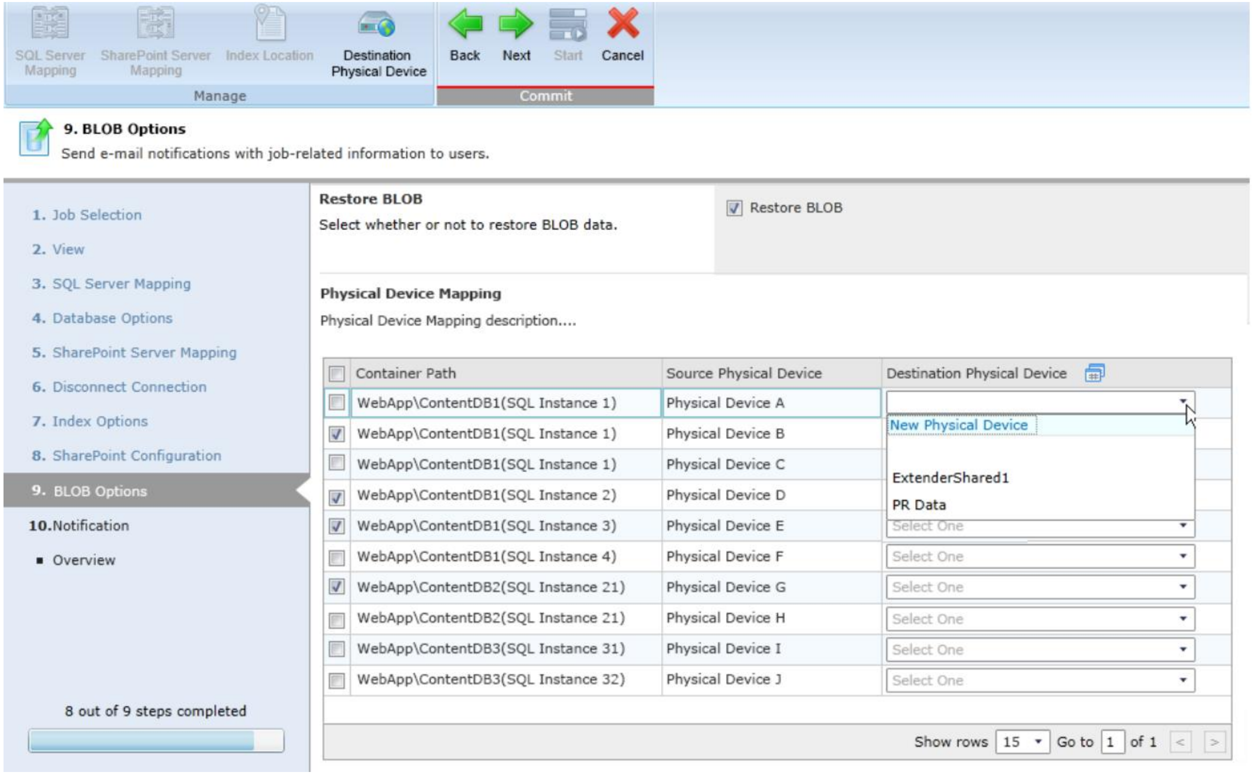

9. BLOB Options

Send e-mail notifications with job-related information to users.

1. Job Selection
2. View
3. SQL Server Mapping
4. Database Options
5. SharePoint Server Mapping
6. Disconnect Connection
7. Index Options
8. SharePoint Configuration
9. BLOB Options
10.Notification
- Overview

8 out of 9 steps completed

Restore BLOB

Select whether or not to restore BLOB data.

☑ Restore BLOB

Physical Device Mapping

Physical Device Mapping description....

□	Container Path	Source Physical Device	Destination Physical Device
□	WebApp\ContentDB1(SQL Instance 1)	Physical Device A	New Physical Device ExtenderShared1 PR Data
☑	WebApp\ContentDB1(SQL Instance 1)	Physical Device B	
□	WebApp\ContentDB1(SQL Instance 1)	Physical Device C	
☑	WebApp\ContentDB1(SQL Instance 2)	Physical Device D	
☑	WebApp\ContentDB1(SQL Instance 3)	Physical Device E	Select One
□	WebApp\ContentDB1(SQL Instance 4)	Physical Device F	Select One
☑	WebApp\ContentDB2(SQL Instance 21)	Physical Device G	Select One
□	WebApp\ContentDB2(SQL Instance 21)	Physical Device H	Select One
□	WebApp\ContentDB3(SQL Instance 31)	Physical Device I	Select One
□	WebApp\ContentDB3(SQL Instance 32)	Physical Device J	Select One

Show rows 15 Go to 1 of 1

Note: All the WFE data (configurations) need to be restored from the production farm backups on completion of the full Farm Clone restore wizard.

Restore of SQL Availability Group Replica

Replica restore options in SMSP 8.2 allow you to restore only one of the backed-up replicas at random or to restore all of the replicas of the AlwaysOn AG, as shown in Figure 16.

Figure 16) Replica restore options for SQL replicas.

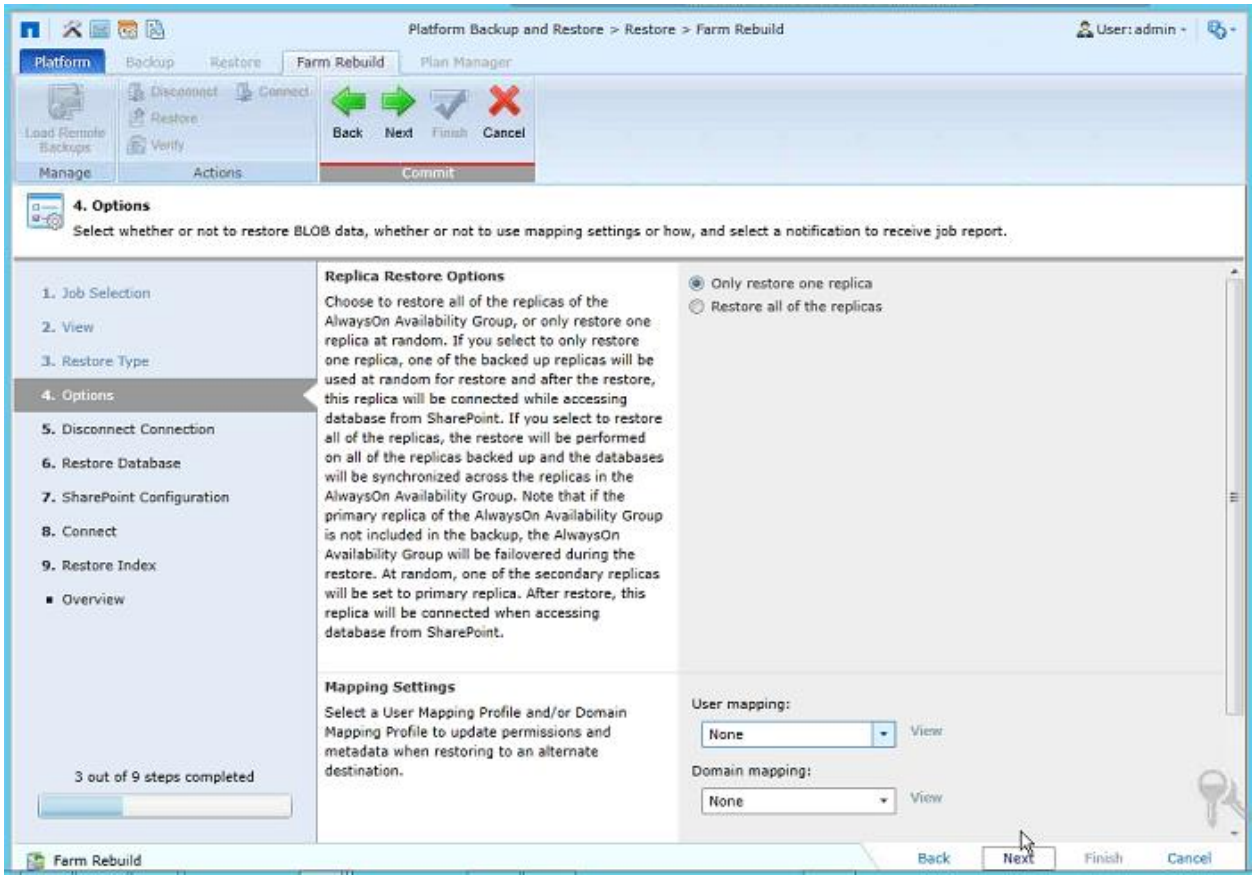

4. Options

Select whether or not to restore BLOB data, whether or not to use mapping settings or how, and select a notification to receive job report.

1. Job Selection
2. View
3. Restore Type
4. Options
5. Disconnect Connection
6. Restore Database
7. SharePoint Configuration
8. Connect
9. Restore Index
- Overview

3 out of 9 steps completed

Replica Restore Options

Choose to restore all of the replicas of the AlwaysOn Availability Group, or only restore one replica at random. If you select to only restore one replica, one of the backed up replicas will be used at random for restore and after the restore, this replica will be connected while accessing database from SharePoint. If you select to restore all of the replicas, the restore will be performed on all of the replicas backed up and the databases will be synchronized across the replicas in the AlwaysOn Availability Group. Note that if the primary replica of the AlwaysOn Availability Group is not included in the backup, the AlwaysOn Availability Group will be failovered during the restore. At random, one of the secondary replicas will be set to primary replica. After restore, this replica will be connected when accessing database from SharePoint.

● Only restore one replica
○ Restore all of the replicas

Mapping Settings

Select a User Mapping Profile and/or Domain Mapping Profile to update permissions and metadata when restoring to an alternate destination.

User mapping: None View

Domain mapping: None View

- **Only restore one replica.** Restores only one of the backed-up replicas at random. This replica will be connected when databases are accessed from SharePoint after the restore job.
- **Restore all of the replicas.** Restores all the backed-up replicas and the databases are synchronized across the replicas in the AlwaysOn AG. The original AG listener will still be connected when databases are accessed from SharePoint. To select this option, you should have backed up all of the replicas of the AlwaysOn AG. If the primary replica of the AlwaysOn AG is not backed up, the AlwaysOn AG will fail over during the restore. One of the secondary replicas will be set to primary at random, and it will be connected when databases are accessed from SharePoint.

Best Practices

- Confirm that the SharePoint file system resources are restored prior to restoring the farm components.
- Verify that the source and destination servers are the same version and patch level for SharePoint. You cannot restore backed-up SharePoint 2010 data to SharePoint 2013 or restore backed-up SharePoint 2013 data to SharePoint 2010. If the site within SharePoint 2013 is a SharePoint 2010 mode site, the content can be restored only to a SharePoint 2013 site that is in the SharePoint 2010 mode.
- Confirm that you periodically test the restore from SMSP backups to validate that the backups can be used in the event of a disaster.
- SMSP supports the verification of backups on SnapMirror destinations and SnapVault secondary locations, thus offloading the read I/O from the production database servicing users.
- To enable a Farm Clone, make sure that the config and admin content databases reside on a NetApp storage system.
- Restore BLOB data first, and then restore content and stub databases. Also, disable the garbage collection (BLOB retention) before the database restore finishes.
- An out-of-place restore or a restore to an alternate location can be done to any SharePoint Server, provided it has the SMSP agent installed.
- To restore customizations successfully, NetApp recommends that you deploy the .wsp file for both the trusted and sandboxed solutions to the destination.
- Make sure that the control service does not run on the same host as the SMSP agent because, during an SMSP farm rebuild operation, the agent is disconnected and IIS on the WFE agent is reset.
- For SharePoint content, SMSP can create a granular out-of-place restore to a different farm at a different granular level (site collection to item/item version level). However, SMSP does not support out-of-place restores of SharePoint services and components because SMSQL-based DB backup is based on local Snapshot copies that might not be available to a SQL Server agent on another farm. For the same reason, SMSP does not support out-of-place database restore to a different SQL Server instance, even though this is possible through SMSQL.
- SQL Server mappings for full Farm Clone only support many-to-one server mapping and one-to-one server mapping and not one-to-many server mappings.

For additional information, refer toSnapManager 8.2 for Microsoft SharePoint Platform Backup and Restore User's Guide.

6.4 Storage Optimization

Microsoft offers remote BLOB storage (RBS) as the official offloading technique for BLOB externalization, implemented by SQL Server. RBS is available in SharePoint 2013 and SharePoint 2010, based on the API supported by SQL Server 2012 and SQL Server 2008 R2. SMSP 8.2 includes storage optimization solutions to keep your SQL Server resources optimized with intelligent archiving and BLOB offloading to a NetApp SMB share on tiered storage. Deduplication and compression enabled on NetApp storage work on externalized BLOBs to provide improved I/O operation, deduplication, and/or compression. However, RBS does not increase the storage limits of content databases. The supported limits still hold true for SharePoint databases.

Note: SMSP supports SQL Server FILESTREAM provider for Microsoft SQL Server, with content externalized to a local SQL Server as direct-attached storage or iSCSI-attached NetApp SAN/NAS storage. Remote RBS FILESTREAM provider is not supported.

Note: SMSP does not support third-party RBS providers.

Note: SMSP Storage Manager and connector rely on RBS BLOB provider to externalize content to CIFS shares; you must use the Enterprise Edition of SQL Server 2008 R2 with SP1, SQL Server 2012, or SQL Server 2014. If you use the SQL Server Standard Edition, you can use the local SQL Server file stream.

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

Storage Manager

SharePoint can contain unstructured data called BLOBs consuming up to 95% of content database space. If ignored, BLOBs can lead to degradation of database performance and the user experience. With SMSP Storage Manager, organizations can mitigate the negative consequences of exponential data growth. They can do this by combining multiple real-time and scheduled business rules to externalize BLOB content based on file size, type, or other document properties to NetApp CIFS shares. Also, with shredded storage in SharePoint 2013, a single document will be externalized as multiple “shreds” to the BLOB store, shown in externalization details in Figure 17.

Figure 17) Document stored as shreds.

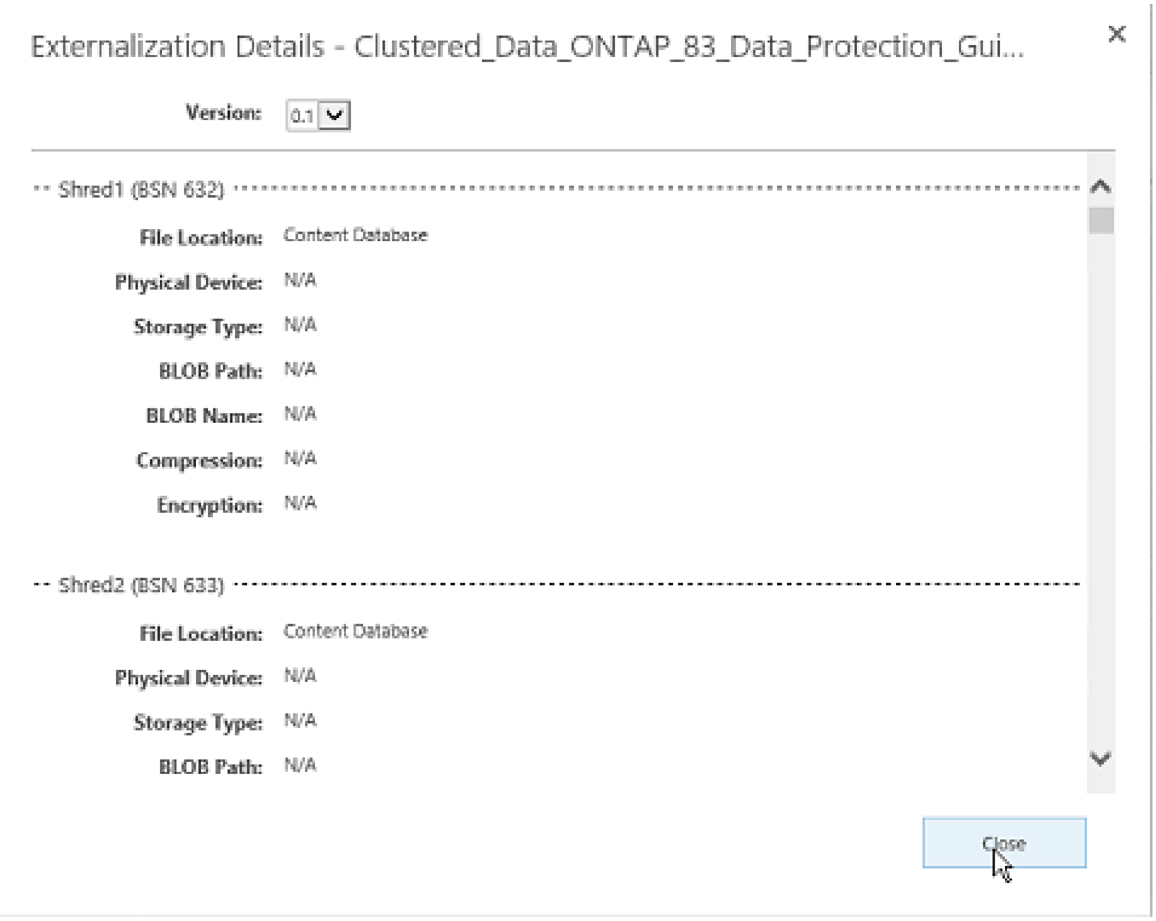

Externalization Details - Clustered_Data_ONTAP_83_Data_Protection_Gui...

Version: 0.1

-- Shred1 (BSN 632) ----------

File Location: Content Database

Physical Device: N/A

Storage Type: N/A

BLOB Path: N/A

BLOB Name: N/A

Compression: N/A

Encryption: N/A

-- Shred2 (BSN 633) ----------

File Location: Content Database

Physical Device: N/A

Storage Type: N/A

BLOB Path: N/A

The intent of shredded storage is to improve SQL I/O by utilizing smaller files, but this often results in poor performance when combined with RBS, as depicted in Figure 18.

Figure 18) Shred size versus performance.

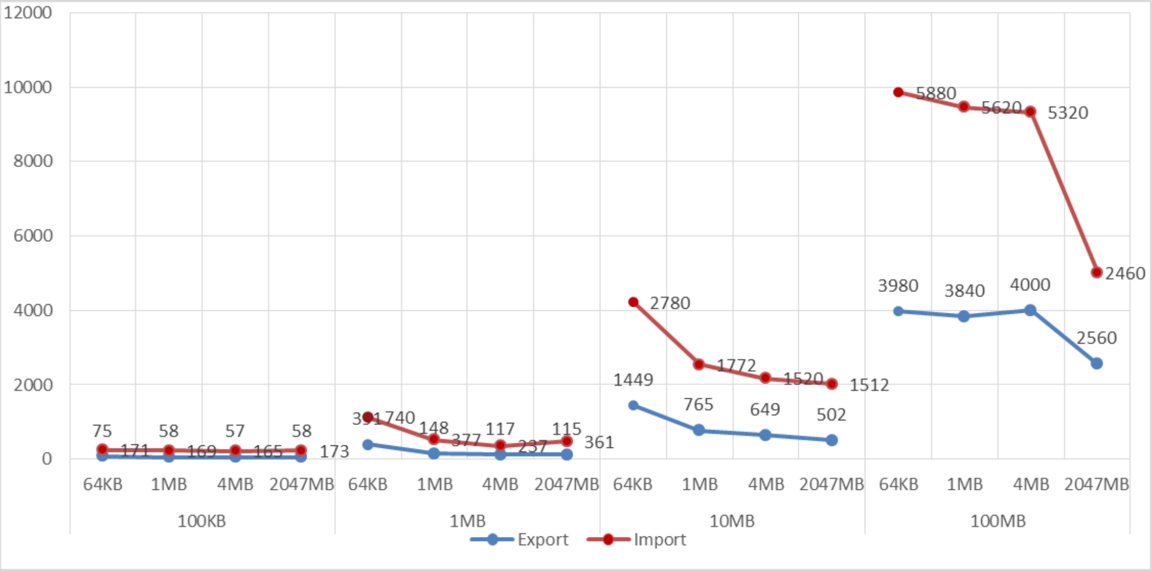

Although increasing the shred size (file write chunk size) in SharePoint from its default 64KB to 1MB and then to 4MB does result in better performance, it is unlikely that customers will be willing to modify this setting for their entire farm. To provide comparable performance between files surfaced from the database with small shreds and files surfaced over RBS, SMSP 8.2 Storage Manager now stores *full files* on NetApp CIFS shares each time BLOB is externalized. In the event of modification, even though only a shred is added to the database, the scheduled Storage Manager action will still retrieve and store the entire version of that file. Although this is less space efficient when multiple versions are used (since the full file is stored instead of shreds), there are several advantages:

- There is no performance penalty for leveraging RBS and keeping the current file write chunk size at 64KB.
- RBS allows smaller databases and easier maintenance.
- File system maintenance (such as archiving to tape) is much simpler with fewer files.
- Single item recovery from legacy BLOB store backups is simpler to perform to identify a single file versus multiple shreds.
- Single-instance storage saves substantially more space than shredded storage.
- After working with several customers in storing files greater that 50TB, NetApp recommends using this option over shredded storage.

The setting “Generate non-shredded BLOBs,” as shown in Figure 19 is available when creating rules in the Scheduled Storage Manager:

Figure 19) Generate non-shredded BLOBs.

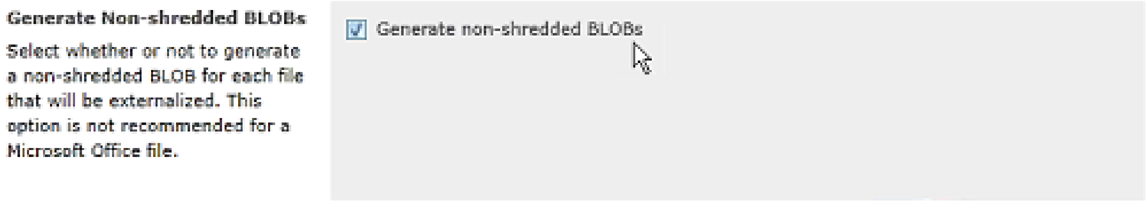

Generate Non-shredded BLOBs

Select whether or not to generate a non-shredded BLOB for each file that will be externalized. This option is not recommended for a Microsoft Office file.

☑ Generate non-shredded BLOBs

With this option enabled, clicking “Externalization Details” for a list item in the document library shows the complete file in the BLOB path, as shown in Figure 20.

Figure 20) Externalized data stored as nonshredded BLOB.

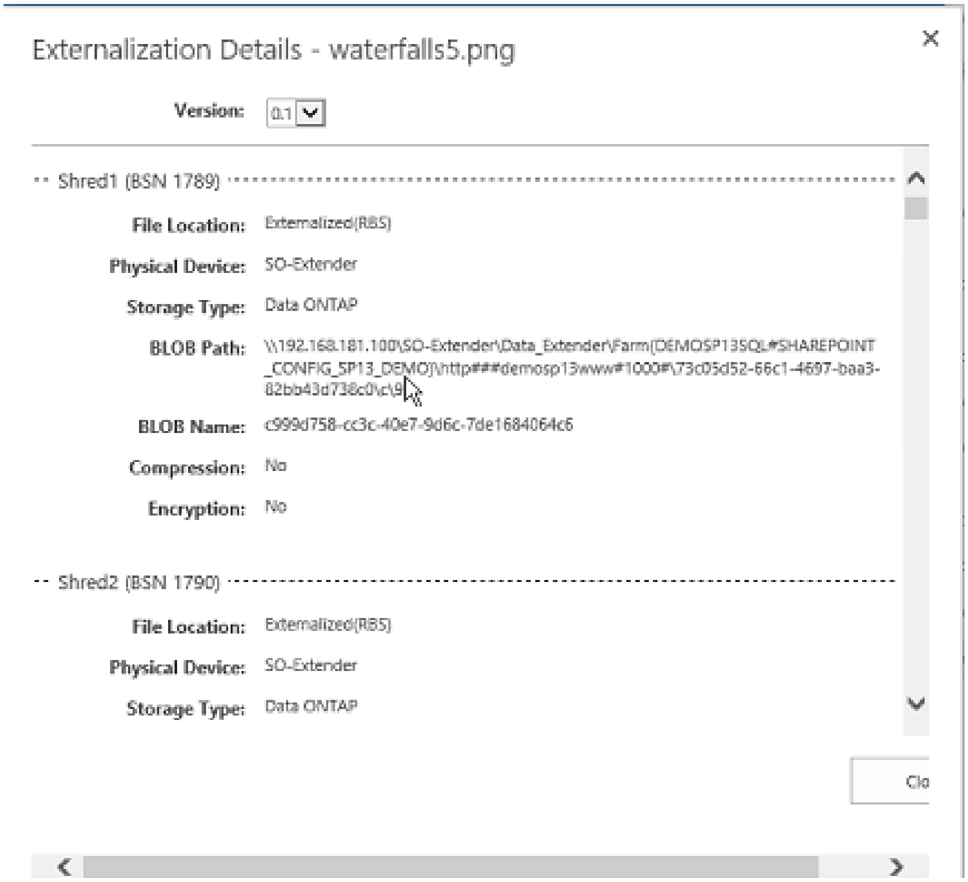

Externalization Details - waterfalls5.png

Version: 0.1

-- Shred1 (BSN 1789) --

File Location: Externalized(RBS)

Physical Device: SO-Extender

Storage Type: Data ONTAP

BLOB Path: \\192.168.181.100\SO-Extender\Data_Extender\Farm(DEMOSP13SQL#SHAREPOINT_CONFIG_SP13_DEMO)\http###demosp13www#1000#\73c05d52-66c1-4697-baa3-82bb43d738c0\c\9

BLOB Name: c999d758-cc3c-40e7-9d6c-7de1684064c6

Compression: No

Encryption: No

-- Shred2 (BSN 1790) --

File Location: Externalized(RBS)

Physical Device: SO-Extender

Storage Type: Data ONTAP

Because certain files such as Microsoft Office documents perform better with shredded storage, the setting to “Generate non-shredded BLOBs” is optional for each rule created in the Scheduled Storage Manager. By default it is disabled.

Note: NetApp does not recommend using Real-Time Storage Manager in combination with shredded storage in SharePoint 2013 and will not benefit from this change.

SMSP 8.2 also displays the content databases in the Storage Manager tree view, as shown in Figure 21, by setting isEnabledDBTreeMode to true in C:\Program Files\NetApp\SMSP8\Manager\Control\Config\StorageManager\ControlStorageManagerSettings.config.

Note: There is no need to restart SMSP services.

.

Figure 21) View content databases in Storage Manager.

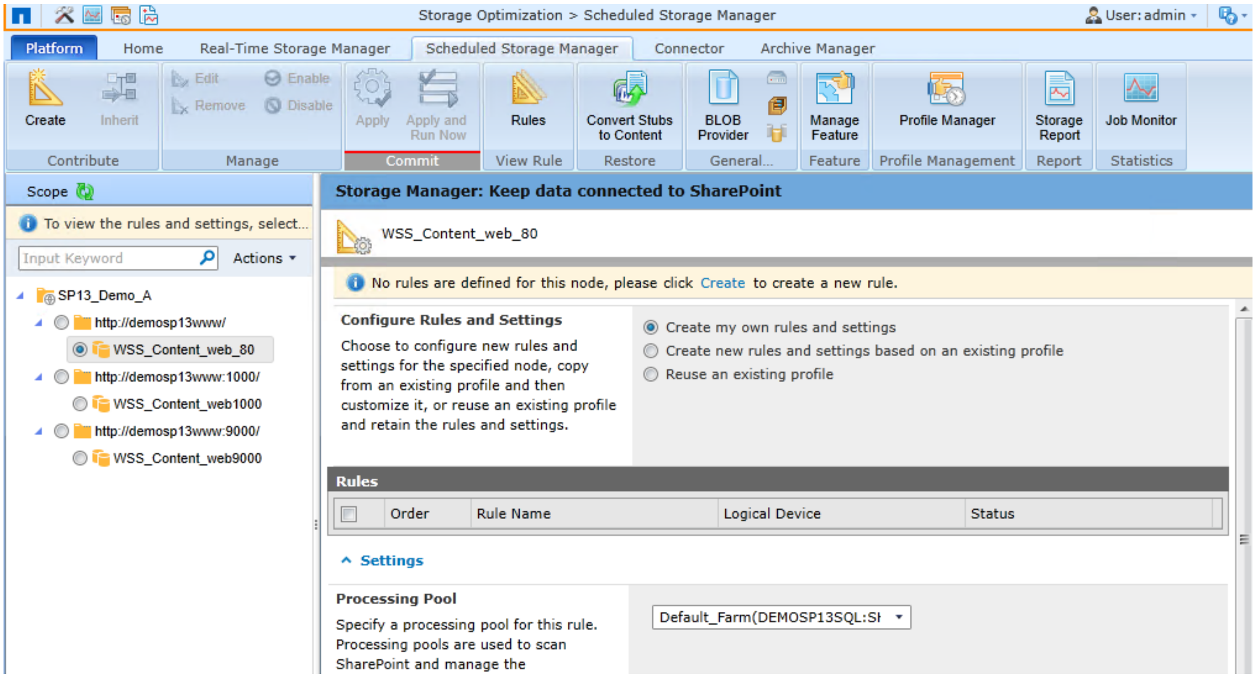

Storage Manager: Keep data connected to SharePoint

WSS_Content_web_80

No rules are defined for this node, please click Create to create a new rule.

Configure Rules and Settings

Choose to configure new rules and settings for the specified node, copy from an existing profile and then customize it, or reuse an existing profile and retain the rules and settings.

● Create my own rules and settings
○ Create new rules and settings based on an existing profile
○ Reuse an existing profile

Rules

□	Order	Rule Name	Logical Device	Status

Processing Pool

Specify a processing pool for this rule. Processing pools are used to scan SharePoint and manage the

Default_Farm(DEMOSP13SQL:SI

Connector

The SMSP connector is used to collaborate on NetApp CIFs shares and cloud storage resources directly to map security and metadata and to connect files to SharePoint without migration. Connected content appears as regular SharePoint content and can be leveraged exactly as if it resided within a SharePoint document library. All of SharePoint's powerful document management functionality–including permissions management, workflows, alerts, and versioning–can be applied to connected content.

The flow for granular restore when BLOB resides on a CIFS share, especially in the case of the connector, is as follows:

1. If BLOB was selected during backup:
   a. If the source volume is online, the restore process tries to find BLOB from the source volume Snapshot copy or backup device.
   b. If the source volume is offline, the restore process tries to find BLOB from the current path saved in the connector library settings.
2. If BLOB was not selected during backup, the restore process tries to use the current setting path to find BLOB data and tries to read device information from the cache file "`C:\Program Files\NetApp\SMSP8\Agent\data\SP2013\Arch\Cache\AvePoint.GCommon.Contract.Storage.Entity.PhysicalDeviceDto.Cache.`"

Connector for NFS Data

With SMSP 8.2, the NFS connector extends the same connector capabilities to connect NFS shares to SharePoint, as shown in Figure 22.

Figure 22) Connector settings to connect to NFS share.

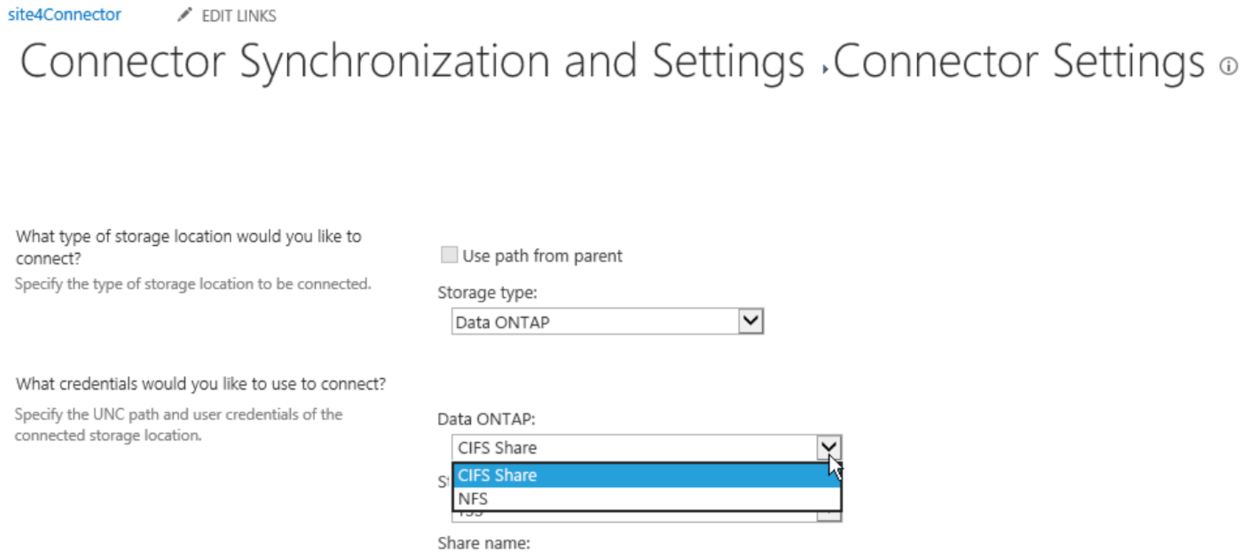

site4Connector   EDIT LINKS

Connector Synchronization and Settings › Connector Settings

What type of storage location would you like to connect?

Specify the type of storage location to be connected.

□ Use path from parent

Storage type: Data ONTAP

What credentials would you like to use to connect?

Specify the UNC path and user credentials of the connected storage location.

Data ONTAP: CIFS Share

Share name:

The variations for the user interface are located only in the connection information provided by the user in both the SharePoint library interface and the SMSP management interface.

Figure 23) Connector settings on SharePoint document library settings page.

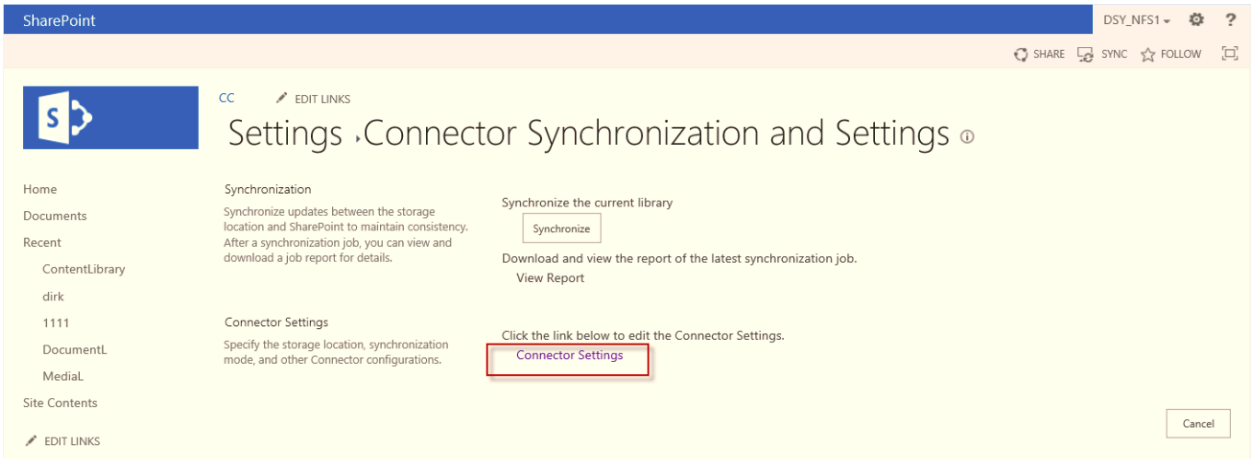

Settings › Connector Synchronization and Settings

Synchronization

Synchronize updates between the storage location and SharePoint to maintain consistency. After a synchronization job, you can view and download a job report for details.

Synchronize the current library

Synchronize

Download and view the report of the latest synchronization job.

View Report

Connector Settings

Specify the storage location, synchronization mode, and other Connector configurations.

Click the link below to edit the Connector Settings.

Connector Settings

Cancel

Figure 24 shows the SharePoint document library connected to an NFS share in connector settings.

Figure 24) Connected to NFS share.

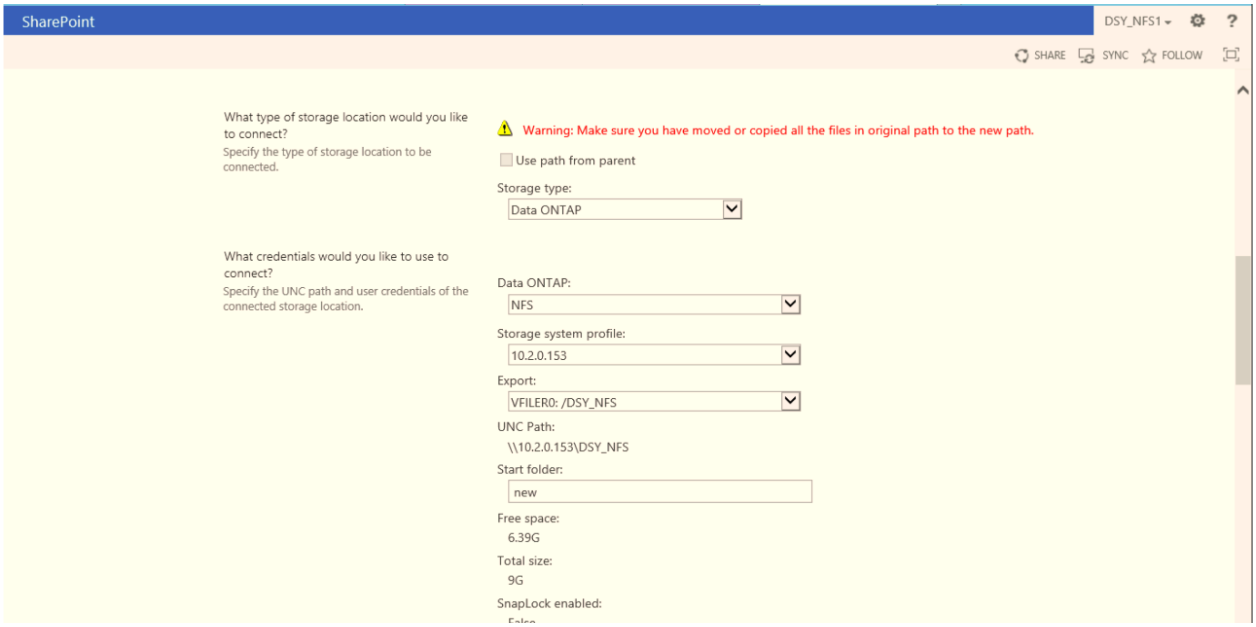

What type of storage location would you like to connect?
Specify the type of storage location to be connected.

⚠ Warning: Make sure you have moved or copied all the files in original path to the new path.

□ Use path from parent

Storage type: Data ONTAP

What credentials would you like to use to connect?
Specify the UNC path and user credentials of the connected storage location.

Data ONTAP: NFS

Storage system profile: 10.2.0.153

Export: VFILER0: /DSY_NFS

UNC Path: \\10.2.0.153\DSY_NFS

Start folder: new

Free space: 6.39G

Total size: 9G

SnapLock enabled: False

Only a select set of metadata will be loaded on synchronization, especially when connected to an NFS share, as shown in Figure 25.

Figure 25) Metadata synchronized from NFS share.

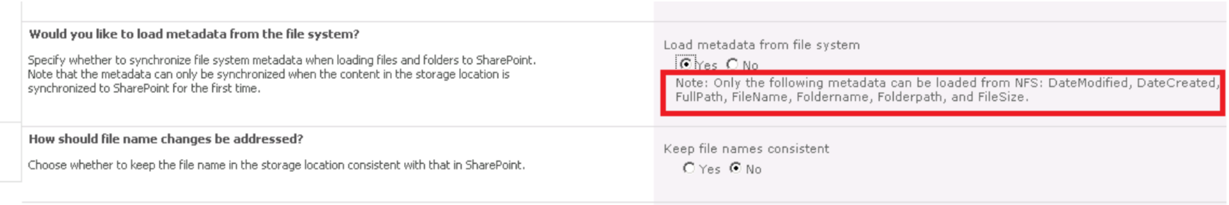

Would you like to load metadata from the file system?

Specify whether to synchronize file system metadata when loading files and folders to SharePoint. Note that the metadata can only be synchronized when the content in the storage location is synchronized to SharePoint for the first time.

Load metadata from file system
● Yes ○ No

Note: Only the following metadata can be loaded from NFS: DateModified, DateCreated, FullPath, FileName, Foldername, Folderpath, and FileSize.

How should file name changes be addressed?

Choose whether to keep the file name in the storage location consistent with that in SharePoint.

Keep file names consistent
○ Yes ● No

Note: The connector for NFS is tested and verified only for Data ONTAP 7-Mode storage systems until clustered Data ONTAP supports Windows native NFS clients.

Migration of Connected Share

Migration of connector BLOB data residing on a NetApp CIFS or NFS share using SnapMirror requires remapping the connector library to a new CIFS share location in order to perform database site collection and folder or item-level restores. This path can be changed only from SharePoint site > Configure Connector Library Settings page > Connector Settings.

Note: Restart the SMSP agent service each time you modify the physical device used by storage optimization to clear the cache.

Best Practices

- NetApp recommends creating a CIFS share for BLOB on a volume that does not contain an operating system, paging files, database data, log files, or the tempdb file.
- Keep the CIFS share for BLOB separate for each farm or SQL Server instance for ease of backup data retention management.
- Confirm that the volume used by the SMSP storage policy device that contains backup data (including restore index) is separate from that of the volume used by the CIFS share for BLOB, especially when you use SnapVault with SMSP backups.
- Set the stub database to simple recovery model to avoid big transaction log size. If the user needs to change to full recovery model, make sure that the LUN has enough space for transaction logs and shrink the logs if necessary.
- Verify that the SMSP stub database is included in the backup and added as a custom database.
- Use a separate stub database per web application when:
  - You need to divide the web applications into multiple backup plans.
  - Also, when BLOB backup is part of the backup plan, the restore does not overwrite the stub data from a different backup plan.
- The SMSP agent uses the stub DB for BLOB access; therefore, make sure that the stub database in the SQL instance:
  - Is close to the SharePoint WFE and SMSP manager.
  - Is in the same SQL Server instance as the SharePoint content database.
- Do not access or change the externalized BLOB content manually outside of normal SharePoint operations.
- Confirm that the volumes used by the CIFS shares are big enough to hold a large amount of BLOB data.
- Confirm that the RBS provider that comes with the SMSP agent is installed on every server in the SharePoint farm, because these DLLs implement methods for the RBS application programming interface (API) and perform the actual externalization of BLOB to the NetApp CIFS share.
- If you currently use SQL Server FILESTREAM and want to move to RBS, use the SMSP data import wizard to convert the FILESTREAM RBS BLOB to SMSP RBS BLOB. After BLOB conversion, the FILESTREAM BLOB becomes orphaned; you have to run the RBS garbage collection task outside of SMSP.
- If you choose to use SQL Server authentication when creating the SMSP databases, make sure the user has db_creator and security admin database roles assigned.
- When using SMSP connector libraries, make sure to comply with software boundary limits as specified by Microsoft for SharePoint.
- To make sure connector synchronization operation does not become time consuming, spread the documents onto different libraries and site collections.
- In the case of the 7-Mode Transition Tool (7MTT), make sure you migrate to the CIFS LIF to be able to access the connector libraries successfully.
- Review and update the storage system profile appropriately unless all the suitable LIFs were migrated.

For additional information, refer to SnapManager 8.2 for Microsoft SharePoint Storage Optimization User's Guide.

6.5 High Availability

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

Control Service High Availability

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

- The first control service installed is the master, which can be changed.
- Because the control database for the control service is now in SQL Server, clustering and log shipping apply for HA.
- Make sure you register the agents and media service to the control service. In case this control service becomes unavailable, reregister the agents and media service to another control service in order to continue to access SnapManager for SharePoint.
- Also make sure to configure a report location in the job monitor before you use the log manager and job monitor with SMSP control service high availability. Otherwise, each server in which the control service is installed will retain its own log only for the jobs carried out by the control service installed on the server.

The following requirements must be met:

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

Media Service High Availability

If you use SMSP to manage your SharePoint farm, media service plays a very important role. Media service is used only to store the generated index and run postprocessing after the index has been generated and transferred by the SMSP agent on SQL Server. You can configure high availability of the media service by using a Microsoft Windows cluster failover configuration for load-balanced access to the data storage locations. This requires that all LUN/SMB physical devices have the same drive letter and mount point on all nodes. In cluster administrator, set all SMSP manager services as cluster generic services. Set control service or media service as a dependent on the shared drives. Use the media service server cluster name and IP address for any interaction with it.

When many SQL Server agents run index generation in parallel, configure the backup plans to use different storage policies that leverage multiple media servers to make sure that the media servers used within a storage policy are polled sequentially.

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

Refer to Hardware and Software Requirements for SharePoint 2013.

7 SMSP 8.2 and NetApp SnapVault

SMSP provides a rapid solution to archive backup Snapshot copies to the secondary storage using NetApp SnapVault technology, which is disk-to-disk backup software that safeguards data at the block level with a fast and streamlined solution. This enables retention of weeks of SMSP backup Snapshot copies of databases, BLOB, search index, and storage policy device. These SnapVault relationships are created using NetApp OnCommand® System Manager. The SnapVault backup data (or the remote backup data) is used only when the local backup has been deleted based on retention settings.

The SMSP backup retention of SharePoint databases is mapped to SMSQL, which actually creates the database Snapshot copies. The retention of search index, BLOB data, and storage policy device relies on the database backup retention. The Snapshot copies of search index and BLOB CIFS volumes at the local site will be deleted when SMSQL returns the related database backup that was deleted by retention. However, the local Snapshot copy of the storage policy device will not be deleted and needs to be manually removed. The remote SnapVault Snapshot copies need to be manually deleted, or deleted using vault policy.

In SMSP 8.1, the SMSP storage policy retention rule, as shown in Figure 26, only configures the SMSQL local backup retention. SMSP queries SMSQL for the database backup status to determine if the local data (backup metadata, granular restore index, and so on) should be deleted. SMSQL itself does not manage the archived SnapVault backup retention. It either relies on the NetApp Unified Manager (for 7-Mode) or the vault policy (on clustered Data ONTAP).

Figure 26) SMSP 8.1 storage policy retention rule.

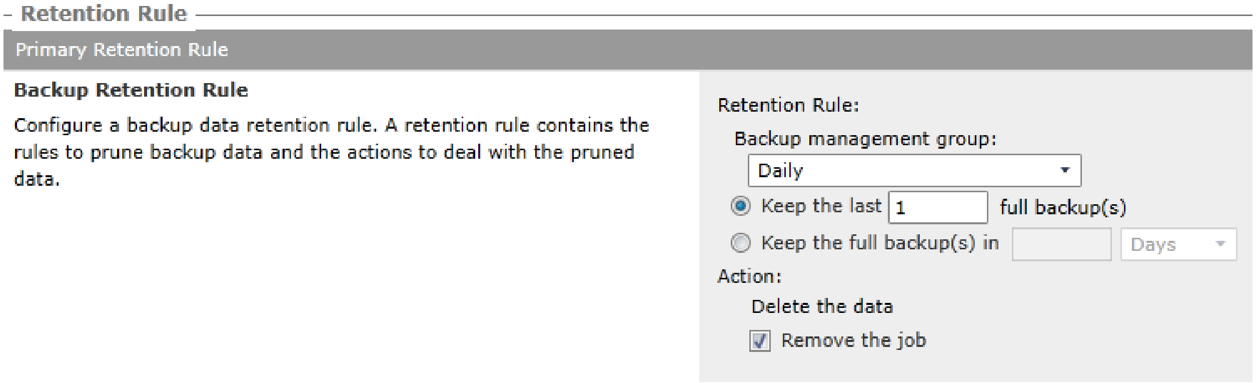

Retention Rule

Primary Retention Rule

Backup Retention Rule

Configure a backup data retention rule. A retention rule contains the rules to prune backup data and the actions to deal with the pruned data.

Retention Rule:

Backup management group: Daily

● Keep the last 1 full backup(s)

○ Keep the full backup(s) in [ ] Days

Action:

Delete the data

☑ Remove the job

Some potential improvement areas of the SnapVault mechanism in SMSP 8.1 are as follows:

1. When SMSQL (7.1 or earlier versions) loses connection to Unified Manager, it will delete the SMSQL backup job metadata, thereby causing SMSP to delete the backup data as well, because SMSQL database backups no longer exist.
2. In clustered Data ONTAP, the vault policy needs to be created and assigned for each volume, with appropriate rules matching the SnapMirror label, as shown in Figure 27. This will become a cumbersome task in a SharePoint farm using data LUNs on different volumes and different storage virtual machines with storage administrator engagement required to manage SnapVault backup retention. It becomes even more laborious in large customer environments that have multiple SharePoint farms.
3. Search index, BLOB, and storage policy device Snapshot copies do not use the selected SMSQL remote backup group as a SnapMirror label for archived backups. You need to separately assign a vault policy for these volumes with the same retention setting as the SMSQL database vault policy so that the SnapVault Snapshot copy can be automatically deleted.

4. The vault policy is capable of maintaining only the “Destination Retention Count,” as shown in Figure 27, for the number of backups to retain in the SnapVault destination and not the number of days that the backups can be retained.

In addition, this Destination Retention Count can cause accidental deletion of backup data. For example, if the user chooses to retain 30 daily backups as defined in the SnapVault policy and if the tests run from SMSP add several Snapshot copies to SnapVault, this setting can cause deletion of necessary Snapshot copies, as defined in the vault policy.

Figure 27) SnapVault policy with destination retention count.

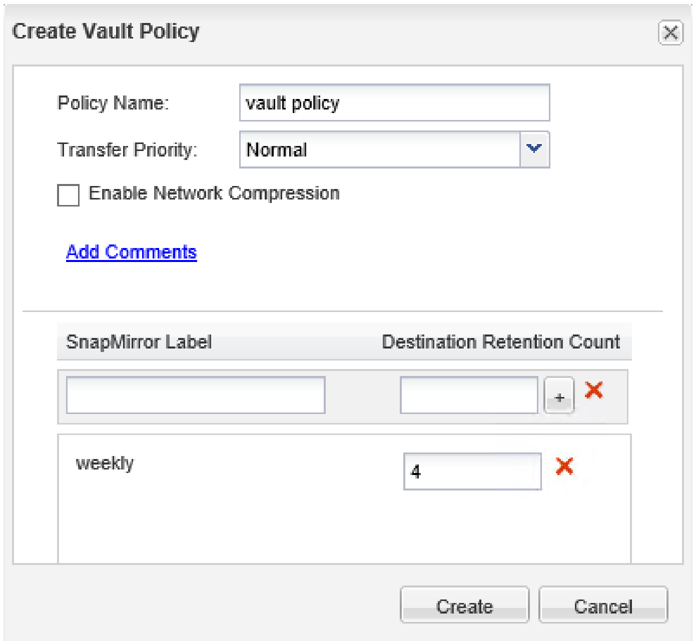

Policy Name: vault policy

Transfer Priority: Normal

□ Enable Network Compression

SnapMirror Label	Destination Retention Count
weekly	4

Hence, SMSP 8.2 allows you to set SnapVault retention in the storage policy to delete the remote SnapVault Snapshot copies for SharePoint search index, CIFS volumes containing BLOB, and storage policy based on SMSQL database backup status. The setting “Enable SnapVault destination retention rule,” as shown in Figure 28, allows all SharePoint volumes to be managed with the same SnapVault retention policy. This greatly reduces the overhead of manually configuring the retention for each individual volume and deletion of SnapVault Snapshot copies, as was the case in prior releases.

Figure 28) SMSP 8.2 SnapVault retention setting in storage policy.

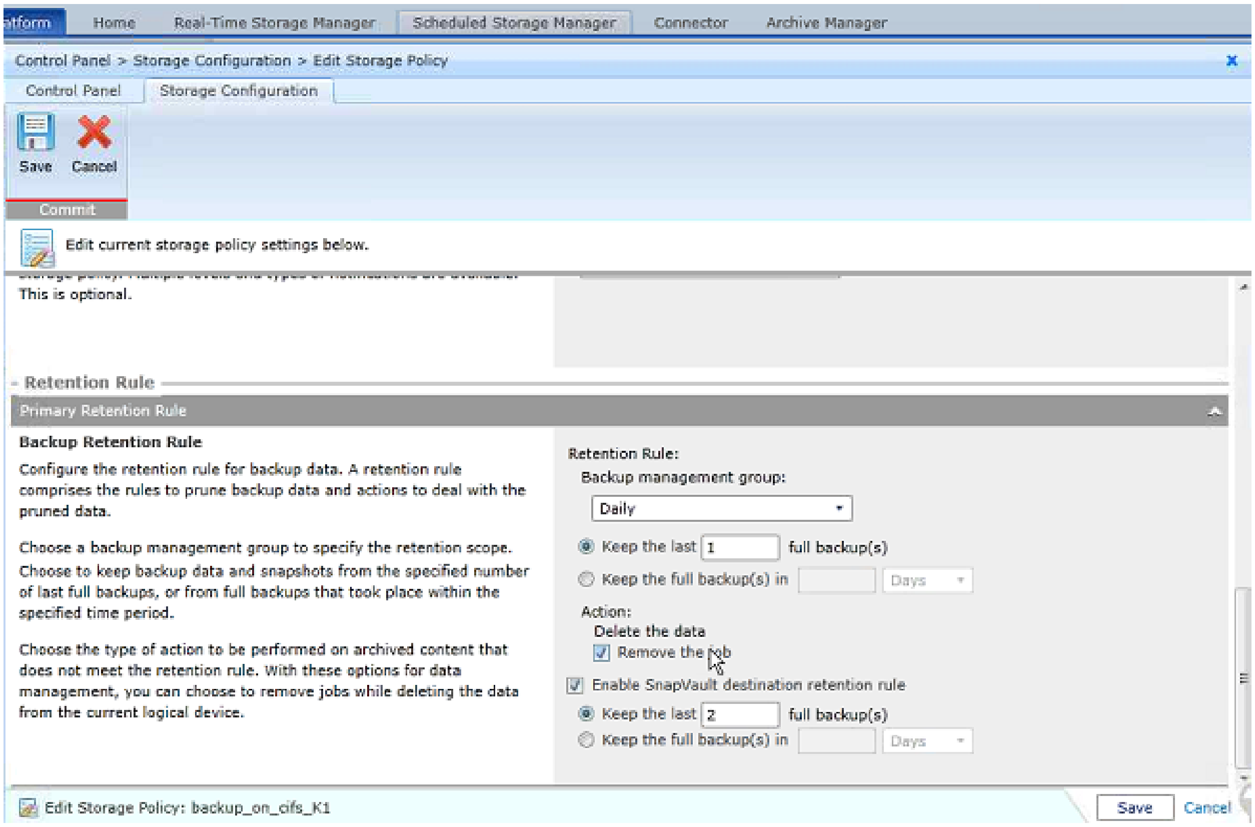
atform
Home
Real-Time Storage Manager
Scheduled Storage Manager
Connector
Archive Manager
Control Panel > Storage Configuration > Edit Storage Policy
Control Panel
Storage Configuration
Save
Cancel
Commit

Edit current storage policy settings below.

This is optional.

Retention Rule

Primary Retention Rule

Backup Retention Rule

Configure the retention rule for backup data. A retention rule comprises the rules to prune backup data and actions to deal with the pruned data.

Choose a backup management group to specify the retention scope. Choose to keep backup data and snapshots from the specified number of last full backups, or from full backups that took place within the specified time period.

Choose the type of action to be performed on archived content that does not meet the retention rule. With these options for data management, you can choose to remove jobs while deleting the data from the current logical device.

Retention Rule:

Backup management group: Daily

● Keep the last 1 full backup(s)

○ Keep the full backup(s) in Days

Action:

Delete the data

☑ Remove the job

☑ Enable SnapVault destination retention rule

● Keep the last 2 full backup(s)

○ Keep the full backup(s) in Days

Edit Storage Policy: backup_on_cifs_K1

SnapVault with SMSP 8.2 works as follows:

1. The option to “Enable SnapVault destination retention rule” is disabled by default. This implies that the local database snapshot and index/BLOB Snapshot copy will follow the retention rule managed by SMSQL with Unified Manager for 7-Mode and the vault policy for clustered Data ONTAP.
2. When “Enable SnapVault destination retention rule” is enabled, SMSP deletes the SnapVault Snapshot copy based on the retention rule defined and the Snapshot copies by name on the SnapVault destination. It ignores the SMSQL status of the remote database backup. You can now choose SnapVault retention based on the number of backups or days.

   Note: Make sure that the storage administrator does not configure the vault policy rule for Snapshot copies on the SnapVault destination.

Figure 29) SMSP 8.2 option to delete SnapVault destination data.

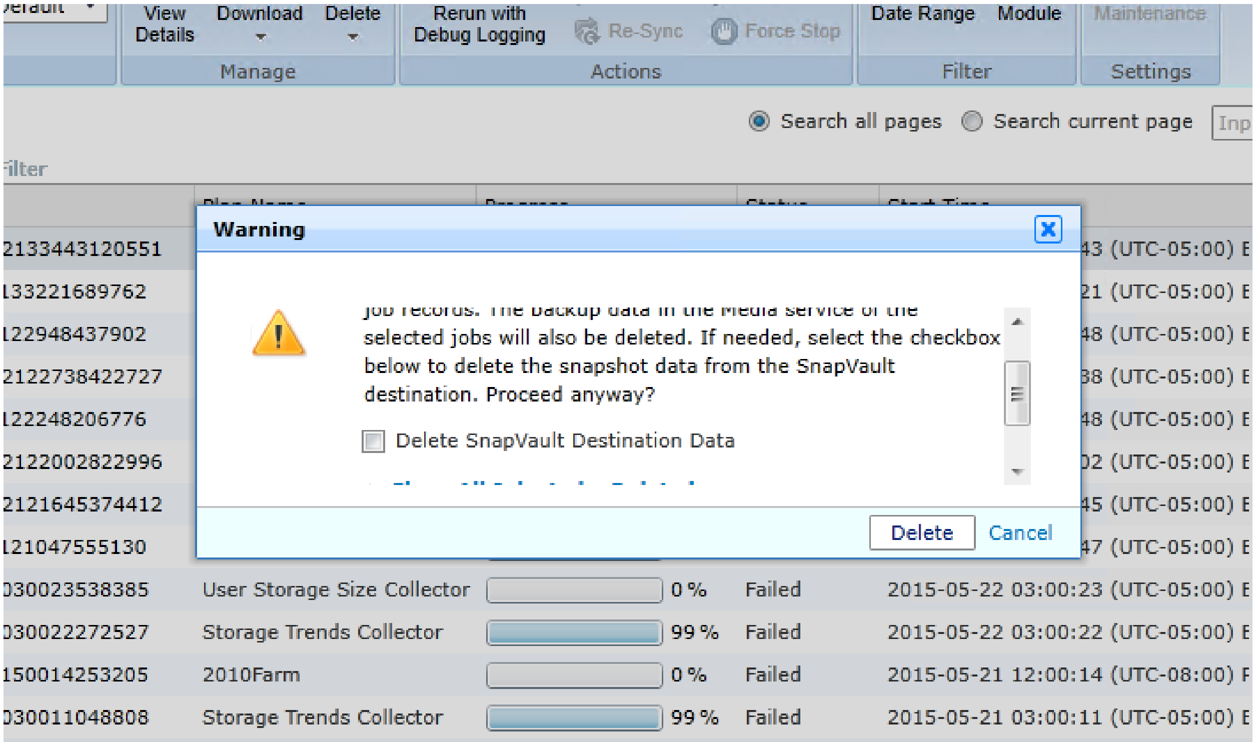

	Plan Name	Progress		Status	Start Time
2133443120551					43 (UTC-05:00) E
133221689762					21 (UTC-05:00) E
122948437902					48 (UTC-05:00) E
2122738422727					38 (UTC-05:00) E
122248206776					48 (UTC-05:00) E
2122002822996					02 (UTC-05:00) E
2121645374412					45 (UTC-05:00) E
121047555130					47 (UTC-05:00) E
030023538385	User Storage Size Collector		0 %	Failed	2015-05-22 03:00:23 (UTC-05:00) E
030022272527	Storage Trends Collector		99 %	Failed	2015-05-22 03:00:22 (UTC-05:00) E
150014253205	2010Farm		0 %	Failed	2015-05-21 12:00:14 (UTC-08:00) F
030011048808	Storage Trends Collector		99 %	Failed	2015-05-21 03:00:11 (UTC-05:00) E

Warning

job records. The backup data in the Media service of the selected jobs will also be deleted. If needed, select the checkbox below to delete the snapshot data from the SnapVault destination. Proceed anyway?

□ Delete SnapVault Destination Data

Delete   Cancel

Best Practices

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

For additional information, refer to SnapVault Best Practices Guide for Clustered Data ONTAP.

8 SharePoint Disaster Recovery with SMSP

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

8.1 NetApp SnapMirror

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

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

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

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

For more information, refer to TR-4372 SnapManager for Microsoft SharePoint Server Disaster Recovery Guide.

9 Virtualization

Businesses of all sizes perform server consolidation across their application infrastructure to lower cost, improve scalability, and improve service-level agreements through virtualization. SharePoint as an application supports virtualization, and so can SnapManager for SharePoint.

Best Practice
The SMSP manager server hosting the control and media services can be virtualized. Make sure that you have sufficient memory allocated for each VM, as defined for system requirements in the SnapManager 8.2 for Microsoft SharePoint Installation Guide.

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

9.1 Microsoft Hyper-V

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

Best Practices

- To reduce disk contention, store system files on aggregates dedicated to storing VM data. Keep the SharePoint content on a separate aggregate. This makes sure that SharePoint I/O is separate from that of VMs.
- SMSP VHDs should be created only as thin fixed-type VHDs.
- NetApp recommends limiting the use of pass-through disks in Hyper-V except wherever considered necessary. This is because a limitation of pass-through disks is that Hyper-V Snapshot copies are not supported.

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

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

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

9.2 VMware ESX

SMSP uses NetApp Virtual Storage Console (VSC) in addition to SDW for LUN provisioning and application-consistent backups and recovery, leveraging NetApp Snapshot copies for VMs hosted in a VMware vSphere environment. The NetApp VSC, which is a server-side plug-in, needs to be installed on the vCenter system. Make sure the ports used by SMSP manager (control and media services) are open on the guest OS VM. If the user plans to have the SMSP manager service VM OS disk on NetApp storage, NetApp recommends following TR-3749: NetApp Storage Best Practices for VMware vSphere.

Best Practices

- Always use SMSP 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.
- Have only FC/iSCSI-attached datastores in the same ESX or ESXi host or in different hosts in the same cluster. Do not mix them.
- Use the Data ONTAP operating system PowerShell Toolkit (PSTK) to automate the test bubble (SRM replicated farm) and SDCLI.
- Use VMFS and NFS datastores for OS file and SharePoint binaries for VMware HA and separate RDMs for SharePoint databases, BLOB, search index, and media storage.
- When using VMware HA, make sure that the net share path used for the SMSP job report location is accessible from the failover machine as well.

AutoSupport Logging Changes in Storage Configuration

NetApp AutoSupport™ (ASUP®) logging in a storage system profile, as shown in Figure 30, helps to understand customer deployment scenarios and encourages adoption.

Figure 30) ASUP logging in storage system profile.

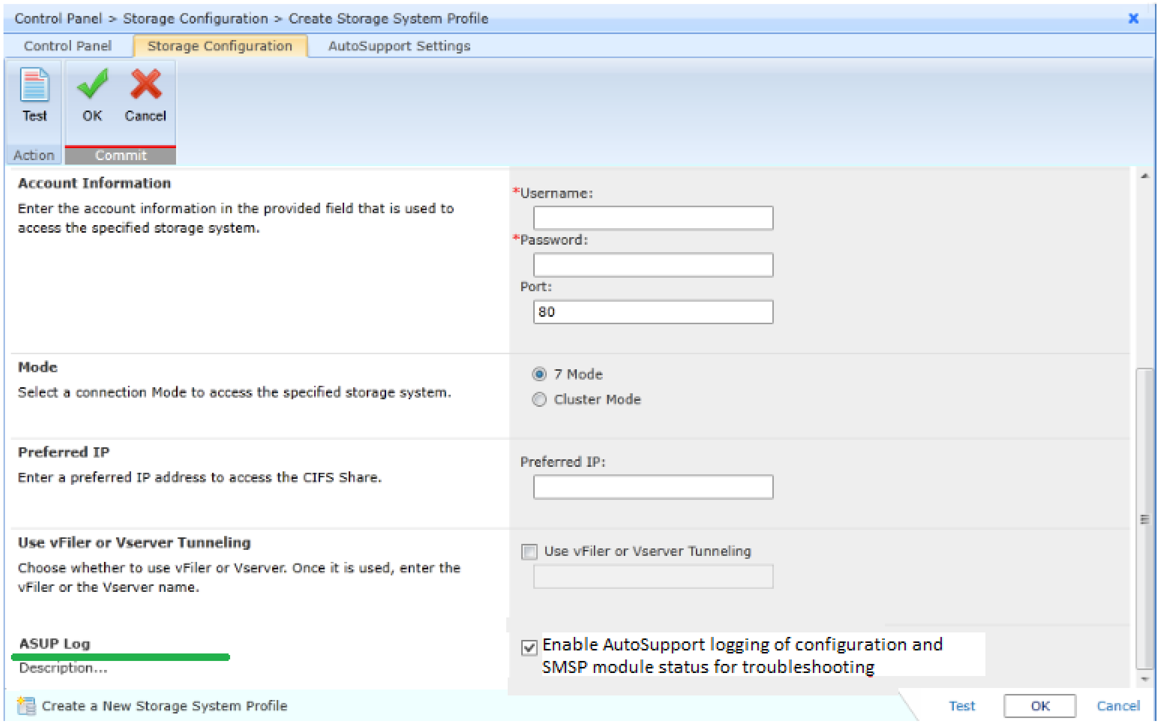

Account Information

Enter the account information in the provided field that is used to access the specified storage system.

*Username:

*Password:

Port: 80

Mode

Select a connection Mode to access the specified storage system.

● 7 Mode

○ Cluster Mode

Preferred IP

Enter a preferred IP address to access the CIFS Share.

Preferred IP:

Use vFiler or Vserver Tunneling

Choose whether to use vFiler or Vserver. Once it is used, enter the vFiler or the Vserver name.

□ Use vFiler or Vserver Tunneling

ASUP Log

Description...

☑ Enable AutoSupport logging of configuration and SMSP module status for troubleshooting

Best Practice

Restart the SMSP Control Service to make sure that the message is sent to the registered NetApp storage system immediately. If you do not restart SMSP Control Service, the message is sent weekly to the registered NetApp storage system.

References

This section lists useful resources and references that can assist you in planning and managing your SharePoint storage environment.

NetApp Storage Systems and Clustered Data ONTAP

- NetApp Data Storage Systems:
  http://www.netapp.com/us/products/storage-systems/
- Data ONTAP 8 Documentation:
  http://support.netapp.com/documentation/productlibrary/index.html?productID=30092
- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization:
- http://www.netapp.com/us/system/pdf-reader.aspx?m=tr-3702.pdf&cc=us
- Deploying VMware vCenter Site Recovery Manager 5 on Data ONTAP Operating in 7-Mode:
  http://www.netapp.com/us/media/tr-4064.pdf
- SharePoint Community:
  https://communities.netapp.com/community/products_and_solutions/microsoft/sharepoint

NetApp SnapMirror

- SnapMirror How-to Guide:
  http://mysupport.netapp.com/NOW/knowledge/docs/olio/guides/Snapmirror.shtml

SnapDrive for Windows and SnapManager for SQL Server

- NetApp SnapDrive for Windows:
  http://support.netapp.com/documentation/productlibrary/index.html?productID=30049
- SnapManager for Microsoft SQL Server:
  http://support.netapp.com/documentation/productlibrary/index.html?productID=30041

Microsoft SharePoint Server

- Capabilities and Features in SharePoint 2013:
  https://www.microsoft.com/en-us/download/details.aspx?id=34023
- Hardware and Software Requirements for SharePoint 2013
  http://technet.microsoft.com/en-us/library/cc262485.aspx
- Plan for SharePoint 2013:
  https://technet.microsoft.com/en-us/library/cc261834.aspx
- Overview of Shredded Storage in SharePoint 2013:
  https://www.microsoft.com/en-us/download/details.aspx?id=39719
- Software Boundaries and Limits for SharePoint 2013:
  https://technet.microsoft.com/en-us/library/cc262787.aspx
- Capacity Planning for SharePoint Server 2013:
  https://technet.microsoft.com/en-us/library/ff758645.aspx
- Plan Service Deployment in SharePoint 2013:
  https://technet.microsoft.com/en-us/library/jj219591.aspx
- What's New in SharePoint 2013 Upgrade:
  https://technet.microsoft.com/en-us/library/ee617150.aspx
- Overview of the Upgrade Process to SharePoint 2013:
  https://technet.microsoft.com/en-us/library/cc262483.aspx
- Verify Database Upgrades in SharePoint 2013:
  https://technet.microsoft.com/en-us/library/cc424972.aspx
- Introduction to Shredded Storage in SharePoint 2013:
  http://blogs.technet.com/b/wbaer/archive/2012/11/12/introduction-to-shredded-storage-in-sharepoint-2013.aspx
- Restore Web Applications in SharePoint 2013:
  https://technet.microsoft.com/en-us/library/ee748647.aspx
- Plan for High Availability and Disaster Recovery for SharePoint 2013:
  https://technet.microsoft.com/en-us/library/cc263031.aspx
- SnapManager for Microsoft SharePoint:
  http://mysupport.netapp.com/documentation/productlibrary/index.html?productID=30036
- Restore Farms in SharePoint 2013:
  https://technet.microsoft.com/en-us/library/ee428314.aspx
- Use Best Practice Configurations for the SharePoint 2013 Virtual Machines and Hyper-V Environment:
  https://technet.microsoft.com/en-us/library/ff621103.aspx

Version History

Version	Date	Document Version History
Version 1.0	June 2015	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.

Copyright Information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC. SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at http://www.netapp.com/us/legal/netapptmlist.aspx. TR-4431-0615

NetApp®

www.netapp.com