



Technical Report

SnapManager 6.1 for SharePoint with Data ONTAP 8.1 Operating in 7-Mode

Soumen De, Microsoft Solutions Engineering, NetApp
February 2012 | TR-3887

EXECUTIVE SUMMARY

This document discusses the planning considerations and best practices when deploying Microsoft® SharePoint® Server 2010 on NetApp® storage systems. It also covers the best practices for the NetApp enterprise content management solution for SharePoint, which is called SnapManager® 6.1 for SharePoint. These best practices also apply to SnapManager 6.1 for SharePoint.

TABLE OF CONTENTS

1	INTRODUCTION	4
1.1	PURPOSE AND SCOPE	4
1.2	INTENDED AUDIENCE	4
2	SHAREPOINT SERVER 2010 ARCHITECTURE	4
3	SHAREPOINT SERVER 2010 PLANNING CONSIDERATIONS	8
3.1	SYSTEM REQUIREMENTS	8
3.2	PLANNING CONSIDERATIONS	8
3.3	SHAREPOINT UPGRADE CONSIDERATIONS	12
4	STORAGE PLANNING FOR MICROSOFT SHAREPOINT SERVER 2010	13
4.1	NETAPP STORAGE SOFTWARE AND TOOLS	13
4.2	PRINCIPLES OF DESIGNING STORAGE FOR MICROSOFT SHAREPOINT 2010	15
5	NETAPP SOLUTION FOR SHAREPOINT SERVER	26
5.1	SNAPMANAGER 6.1 FOR SHAREPOINT SERVER OVERVIEW	26
5.2	SNAPMANAGER 6.1 FOR SHAREPOINT SERVER ARCHITECTURE	26
5.3	PREPARING SHAREPOINT DATABASES FOR MIGRATION TO NETAPP STORAGE	31
5.4	UPGRADING SMSP 6.0 TO SMSP 6.1	32
5.5	INSTALLING SNAPMANAGER 6.1 FOR SHAREPOINT SERVER IN A NEW SHAREPOINT ENVIRONMENT	33
5.6	SECURITY PERMISSION REQUIRED IN THE CONFIGURATION OF THE USER ACCOUNT IN SQL	34
5.7	STORAGE OPTIMIZATION	35
5.8	BACKUP GUIDELINES	37
5.9	RETENTIONS	38
5.10	RESTORE GUIDELINES	39
5.11	MONITORING	40
5.12	TROUBLESHOOTING	41
6	PERFORMANCE	43
7	HIGH AVAILABILITY	43
7.1	APPLICATION RESILIENCY	43
7.2	STORAGE RESILIENCY	45
10	VIRTUALIZATION	47
11	DISASTER RECOVERY	50
11.1	DISASTER RECOVERY TESTING	53
11.2	NETAPP METROCLUSTER	55
12	REFERENCES	56

LIST OF TABLES

Table 1) Database details.....	10
Table 2) Aggregate layouts.....	17
Table 3) Ports used by SMSP components.....	29
Table 4) SharePoint Servers 2007 and 2010.....	31
Table 5) SMSP components mapped to SharePoint farm hosts.....	33
Table 6) Server roles in SharePoint Server 2010 environment.....	44
Table 7) SharePoint server role and their virtualization.....	49
Table 8) Replication models.....	50
Table 9) LUN status on SnapMirror destination systems.....	53
Table 10) Replication modes.....	55

LIST OF FIGURES

Figure 1) Single server farm (graphic supplied by Microsoft).....	5
Figure 2) Two server farm (graphic supplied by Microsoft).....	5
Figure 3) Two-tier small farm (graphic supplied by Microsoft).....	6
Figure 4) Three-tier small farm (graphic supplied by Microsoft).....	7
Figure 5) Shared services.....	8
Figure 6) Different components required for SnapManager.....	15
Figure 7) NetApp technologies for storage efficiency.....	16
Figure 8) SMSP data protection and storage optimization architecture.....	27
Figure 9) Account Manager.....	30
Figure 10) Login properties-DETDC\administrator.....	34
Figure 11) Windows permission.....	34
Figure 12) SnapManager for SharePoint 6.1 events.....	40
Figure 13) SnapManager for SharePoint Log Manager.....	41
Figure 14) Log level setting.....	42
Figure 15) Services-log level setting.....	42
Figure 16) NetApp technologies for storage resiliency.....	45
Figure 17) MetroCluster and SnapMirror deployment.....	56

1 INTRODUCTION

1.1 PURPOSE AND SCOPE

This document describes best practices and offers insight into design considerations when deploying Microsoft SharePoint Server 2010 on NetApp storage systems. The goal of this document is to provide guidance for effective and efficient storage deployment planning, end-to-end data protection, and retention planning.

The scope of this guide is limited to technical design considerations based on the design principles and preferred standards recommended by NetApp for storage infrastructure when deploying Microsoft SharePoint Server 2010. This document does not cover end-to-end SharePoint implementation details.

1.2 INTENDED AUDIENCE

This document offers guidance in planning and deploying Microsoft SharePoint Server 2010 on NetApp storage systems. The best practices and recommendations described enable Microsoft SharePoint Server architects and NetApp storage administrators to plan a highly available and easy-to-manage SharePoint environment to meet stringent service-level agreements. It is assumed that the reader has working knowledge of the following:

- NetApp Data ONTAP[®] operating system and SnapMirror[®], SnapVault[®], and SnapRestore[®] technologies
- NetApp SnapDrive[®] data management software
- NetApp SnapManager for SQL Server[®] (SMSQL)
- NetApp SnapManager 6.0 for SharePoint or NetApp SnapManager 5.0 for Microsoft Office SharePoint Server
- Microsoft SharePoint Server 2010 architecture and administration
- Microsoft Office SharePoint Server (MOSS) 2007
- Microsoft SQL Server 2005, 2008 or 2008 R2
- Windows[®] Hyper-V[™] and VMware[®] ESXi

2 SHAREPOINT SERVER 2010 ARCHITECTURE

The traditional three-tier roles of a Microsoft[®] SharePoint[®] Server 2010 farm can be deployed on a single server or many servers.

The three-tier roles include:

- Web server role
- Application server role
- Database server role

In a small farm, server roles can be combined on one or two servers. For example, the Web server and application server roles can be combined on a single server or on two or more servers to achieve redundancy.

Here are some of the different farm topologies used in SharePoint 2010 deployments.

Single Server Farm

This is ideal for a user base of less than 100. This farm consists of a single server catering the requirement of the Web Front End (WFE), Database Server and the Application Server.

Figure 1) Single server farm (graphic supplied by Microsoft).



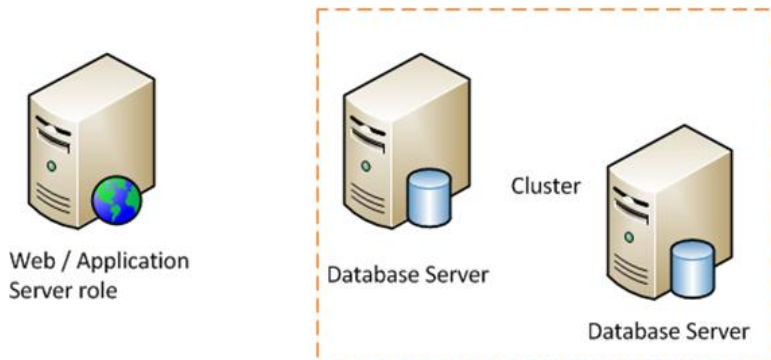
Web Front End,
Database Server,
Application Server all
runs in a single Server

In the small farm environment it still recommended to deploy the SnapManager for SharePoint media server on a separate machine so that the box hosting sharepoint has room to server user requests.

Two Server Farm

This farm is equipped with one Database Server and one Webserver that performs all the application services. For high availability a clustered or mirrored database server is recommended. This farm can cater to the requirement of 100 to 10,000 user base.

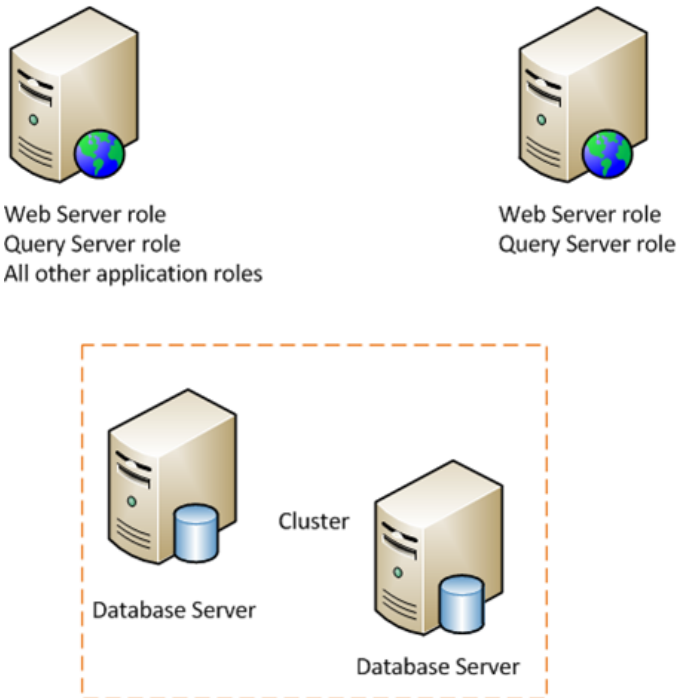
Figure 2) Two server farm (graphic supplied by Microsoft).



Two-Tier Small Farm

This environment is adequate for a user base of 10,000 to 20,000 with low service usage. Two WebServers, one Web Server performing the tasks of the Application Server and the Database Server (clustered or mirrored).

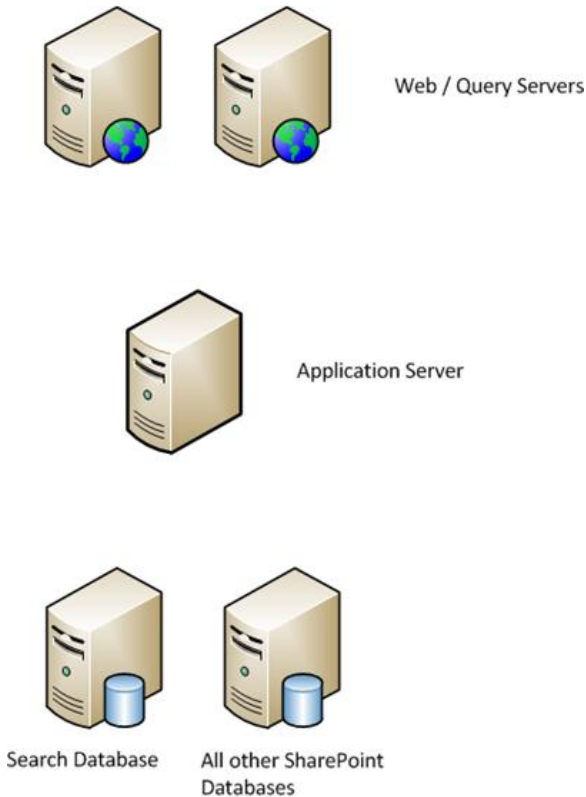
Figure 3) Two-tier small farm (graphic supplied by Microsoft).



Three-Tier Small Farm

Three-Tier small farm is much more similar to a Two-Tier Small Farm. A dedicated Application Server is added to cater the requirements of the Services. To improve the performance of the Search Service the Search Database can be moved to a dedicated Database server. This Farm Topology is ideal for a solution with a Search Database containing nearly 10 million items which can be considered as a large Search Database. Though it is optional, it's a good practice to have a Database Cluster or a mirror.

Figure 4) Three-tier small farm (graphic supplied by Microsoft).



There has been a huge architectural change in SharePoint Server 2010 from its predecessor, SharePoint Server 2007, primarily in its performance and scalability. One of the major architectural changes in SharePoint 2010 is the shared service architecture. To know more about the service architecture visit: <http://technet.microsoft.com/en-us/library/cc560988.aspx>.

Figure 5 shows the services that can be shared across Web applications. These services are called service applications. Shared services replace Shared Services Provider (SSP), offering greater ease and flexibility in administering SharePoint. Unlike SSP, shared services are based on the Windows Web Services framework. SharePoint Server 2010 flattens an SSP, so that each service is an independent entity. For example, Search and Profile are considered shared services. The number of shared services available in a farm depends on which SKU is used. Here is the link to provide you more details on the features available in different SKU. <http://sharepoint.microsoft.com/en-us/buy/Pages/Editions-Comparison.aspx>.

Figure 5) Shared services.



Some of the service applications can be shared across farms. Sharing service applications across Web applications and farms greatly reduces the resources required to provide these services across multiple sites.

For details about the service application on TechNet, visit: <http://technet.microsoft.com/en-us/sharepoint/ff686757.aspx>.

3 SHAREPOINT SERVER 2010 PLANNING CONSIDERATIONS

Planning is required for SharePoint Server 2010 implementation to make sure that all features are completely used.

3.1 SYSTEM REQUIREMENTS

SharePoint Server 2010 is the new-generation application from Microsoft that offers a number of installation scenarios. These installations can be single-server (Test environment) or multiple-server Installations. To know more about the Hardware and software requirements of the server, refer to the TechNet Article <http://technet.microsoft.com/en-us/library/cc262485.aspx>.

3.2 PLANNING CONSIDERATIONS

The physical architecture of the SharePoint deployment is determined by the size and the topology of the farm. (Size means the number of users; farms are categorized as small, medium, or large).

When planning to deploy a SharePoint 2010 environment on NetApp storage, it is important to analyze the existing deployment, if any, for a working dataset model, and to estimate the expected workload demands.

Take the following factors into consideration:

- **Workload** is the user base, usage characteristics, and demand that the system needs to sustain. This involves factors like average daily requests per second (RPS), average RPS at peak time, average daily concurrent users, Office clients, OneNote[®] clients, and so on.
- **Concurrent users** are an important measure for concurrency of operations that users execute on the server farm in a given time period. This also includes daily average and the number of concurrent users at peak load.
- **Requests per second** is the indicator used to describe the demand on the server farm expressed in the number of requests processed by the farm per second.
- **RPO** is an acronym for **Recovery Point Objective**. This is the maximum amount of data loss that is tolerated in the event of a disaster. RPO target vary widely depending on volatility and criticality of the data stored within the SharePoint farm.
- **RTO** is an acronym for **Recovery Time Objective**. RTO drives how SharePoint data protection should be approached prior to a disaster. RTO denotes the timeline within which post disaster farm and data recovery must be completed. If the data that is stored within SharePoint is highly critical to business operations, then RTOs are generally going to trend towards hours, minutes, or maybe even real-time (that is, an RTO that mandates transferring to a hot standby farm or “mirrored” data center for zero recovery time and no interruption in service). For SharePoint data and farms that are less business critical (maybe a publishing site that contains “nice to have” information), RTOs could be days or even weeks.

It is important to understand the performance impacts of the estimated workload. The following best practices help to avoid performance issues.

Best Practices

- Limit the content database size to 200GB to maintain system performance.
- 300 content databases are supported per Web application. NetApp recommends using Windows PowerShell to manage Web applications with a large number of content databases.
- Limit the number of site collections in a content database to 2,000. However, the supported limit is 5,000.
- For optimum performance, NetApp recommends 250,000 site or subsites per site collection.
- NetApp recommends having 100GB per site collection.
- A single document library should not have more than 30,000,000 items. This size varies based on the structure of the library and the type and size of documents stored.
- It is good to have multiple instances of SQL on the backend; this allows more flexible backup and recovery schedule. Having a single monolithic SQL Server can have implications for RTO/RPO times for data recovery as for the backup window.
- SharePoint 2010 IO requirement of 0.25 IOPS per GB content stored. 2 IOPS per GB is recommended for optimal performance.

After Sharepoint 2010 Sp1 the content database can hold data up to 4TB. However, these are limitations and constraints to the content databases that exceed 4TB. Databases that exceed 4TB are referred to as Document Archiving and/or Records Management use:

- The content DB should be distributed across multiple data files.
- Content DB exceeding 4TB is only supported with site collections using the Document center and Record center templates.
- There should less than 5% content read and less than 1% actively written on the site collection hosted within these content DB.
- Alerts, workflow, link Fix up and Item level security are not supported.

Table 1 details the databases that are created as a part of the SharePoint deployment, based on the product version and edition.

Table 1) Database details.

Product Version and Edition	Databases
SharePoint Foundation 2010	<ul style="list-style-type: none"> • Configuration • Central Administration content • Content (one or more) • Usage and Health Data Collection • Business Data Connectivity • Application Registry service (if upgrading from Microsoft Office SharePoint Server 2007 Business Data Catalog) • Subscription Settings service (if it is enabled through Windows PowerShell)

Product Version and Edition	Databases
Additional databases for SharePoint Server 2010 Standard Edition	Search service application: <ul style="list-style-type: none"> • Search administration • Crawl (one or more) • Properties (one or more) User Profile service application: <ul style="list-style-type: none"> • Profile • Synchronization • Social tagging Web Analytics service application: <ul style="list-style-type: none"> • Staging • Reporting • Secure Store • State • Managed Metadata • Word Automation services
Additional database for SharePoint Server 2010 Enterprise Edition	Performance Point
Additional databases for Project Server 2010	<ul style="list-style-type: none"> • Draft • Published • Archive • Reporting
Additional database for FAST Search Server	Search Administration

For more information, visit: <http://technet.microsoft.com/en-us/library/cc678868.aspx>.

Warning: Do not use the default naming for content databases from the SharePoint as it contains the Global Unique Identifier (GUID) as part of the name. The reason for this is the maxlength for the database name is 128 characters. The best way to name the content database is to name it on the content that will be housed in those content db.

CALCULATING CONTENT DATABASE SIZE

Make sure that the server running SQL has the fastest response from the I/O subsystem. Faster disks or arrays provide sufficient I/O operations per second while maintaining low latency and queuing on all disks. SharePoint 2010 requires disk subsystem performance of 0.25 IOPs per GB, 2 IOPS per GB is recommended for optimal performance.

Capacity planning for SharePoint is not a precise task, but here is a formula to find the approximate capacity of the content database:

$$\text{Database Size (ContentDBSum)} = ((D * V) * S) + (10 * (L + (V * D)))$$

Where,

D = Expected number of documents

S = Average size of the document in the storage

L = Estimated number of list items in the environment

V = Average number of versions

10kb is the constant that roughly estimates the amount of metadata required by SharePoint 2010. If the system requires significant use of metadata, that number may be increased.

The following factors affect the size of the content database:

- **Recycle bins.** Until the document is deleted from the first stage and second stage recycle bins, it continues to occupy the content database. By default, items are deleted after 30 days.
- **Audit.** The number of days to keep the audit logs is based on compliance requirements and internal regulatory needs.
- **Office Web Apps.** The Office Web Apps cache can be taxing on the size of the content database. Cache size could be configured up to 100GB. Best practice is to isolate the Office Web Apps cache from other content. To do so create a separate content database, set it to contain only one site collection then configure the office web apps cache to use that database.

CALCULATING INDEX AND CRAWL DATABASE SIZE

During capacity planning for SharePoint, it's important to consider the size of the index and the crawl database.

Use the following formulas to estimate the index size:

$$\text{TotalIndexsize} = \text{ContentDBSum} * 0.35$$

This also reserves room for merges and partitioning of the index. Best practice is to have index partition between 5M and 10M items.

$$\text{QueryIndexsize} = \text{TotalIndexsize} / (\text{number of index partition})$$

Next, calculate the crawl database size:

$$\text{TotalCrawlDBSize} = \text{ContentDBSum} * 0.046$$

For more information about sizing the SharePoint deployment, download SharePoint Server Caches Performance from Microsoft Download Centre:

<http://www.microsoft.com/downloads/en/details.aspx?FamilyID=fd1eac86-ad47-4865-9378-80040d08ac55&displaylang=en>

Remote Blob Store (RBS) is a library API set that is incorporated as an add-on feature pack for Microsoft SQL Server. SnapManager 6.0 for SharePoint has an RBS provider that helps in moving the content from structured data source to NetApp storage. Using RBS has the following benefits:

- Binary large object (BLOB) data can be stored on a NetApp storage system that is configured as a CIFS share.
- The administration of the BLOB storage is controlled by SnapManager for SharePoint (SMSP).
- Database server resources are freed for database operations.

Note: RBS is supported in SQL Server 2008 Enterprise version.

3.3 SHAREPOINT UPGRADE CONSIDERATIONS

If the planning includes an upgrade from SharePoint 2007 to SharePoint 2010, consider the pointers in this section.

- Update your servers to Service Pack 2 (SP2) of Microsoft Office SharePoint Server 2007 or later. To run the upgrade process, as either an in-place or database attach upgrade, the environment must be updated to Service Pack 2 for Microsoft Office SharePoint Server 2007.

Upgrade your operating system to a 64-bit version of Windows Server 2008 R2 or Windows Server 2008 with SP2. If you are using SQL Server, upgrade or migrate to a 64-bit version of SQL Server 2008 R2,

SQL Server 2008 with SP1 and Cumulative Update 2, or SQL Server 2005 with SP3 and Cumulative Update 3.

- Run the pre-upgrade checker to look for potential issues.
The pre-upgrade checker reports missing customizations, issues with orphaned sites, and more, so that you can address these issues before performing the upgrade.
- Perform a trial upgrade on a test farm first.
Back up the live farm, restore it to test servers, and then perform the upgrade. Testing is really important as it's hard to predict what the size of the SharePoint 2010 database is after upgrade. SMSP offers the capability of out of place restore to perform this option on a test bed and trial before upgrade.
- Plan for capacity.

Make sure that you have sufficient disk, processor, and memory capacity to handle upgrade requirements. Refer to the TechNet link for Hardware and software requirements.

<http://technet.microsoft.com/en-us/library/cc262485.aspx>.

- Back up your environment.
Perform a full backup of your environment before upgrading. Use SnapManager for SharePoint to take a complete backup of your server farm so that you can go back to the earlier settings and perform a full farm restore if necessary.
- (Optional) If you are using the database attach upgrade method, set the original databases to read-only.
If you expect a long outage while you perform a database attach upgrade, you can set the databases in the original environment to be read-only so that users can continue to access their data without changing it. It is recommended to keep in mind that before we perform the database attach upgrade it could increase the size of the database. There are third-party tools available to perform the upgrade and can help to avoid the capacity issues and reduce the complexity on a migration For more information, see [Attach databases and upgrade to SharePoint Server](#).

Note: After you begin the upgrade process, do not add any servers to the server farm.

Note: After upgrade, review the Upgrade Status page and upgrade logs.

From the logs, you can determine whether there are any issues that must be addressed. Then review the upgraded sites. The Upgrade Status page reports on the upgrade progress, and the upgrade logs list any errors or warnings that occurred during the upgrade process.

4 STORAGE PLANNING FOR MICROSOFT SHAREPOINT SERVER 2010

4.1 NETAPP STORAGE SOFTWARE AND TOOLS

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

- **SnapDrive for Windows.** This application helps with storage provisioning and managing disks, in both physical and virtual environments. SnapDrive manages the LUNs on the storage system, making them available as local disks on Windows hosts.

Here are the key features of SnapDrive for Windows:

- Enables online storage configuration, LUN expansion, and streamlined management.
- Integrates Data ONTAP Snapshot technology, which creates point-in-time images of data stored on LUNs.
- Works in conjunction with SnapMirror software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes.

- Enables management of SnapDrive on multiple hosts.
- Enables support on Microsoft cluster configurations.
- Enables iSCSI session management.
- Provides support for role-based access control with DataFabric Manager server to separate server administrator and storage administrator functions, and to limit SnapDrive actions and operations depending on the role or job function of the user.
- Supports VMware ESX based iSCSI initiator, LUN migration with Vmotion and RDM LUNs on VMs hosted of ESX servers.
- **SnapManager for SharePoint.** This application helps to reduce storage costs and to manage the SharePoint environment efficiently. Its browser-based interface enables you to automate the data backup process and other administrative functions. SnapManager also makes it easy to migrate and store SharePoint files on Server Message Block (SMB or CIFS) shares outside of Microsoft SQL Server to improve the scalability of large SharePoint 2007 or 2010 deployments.

Here are some important benefits of SnapManager for SharePoint:

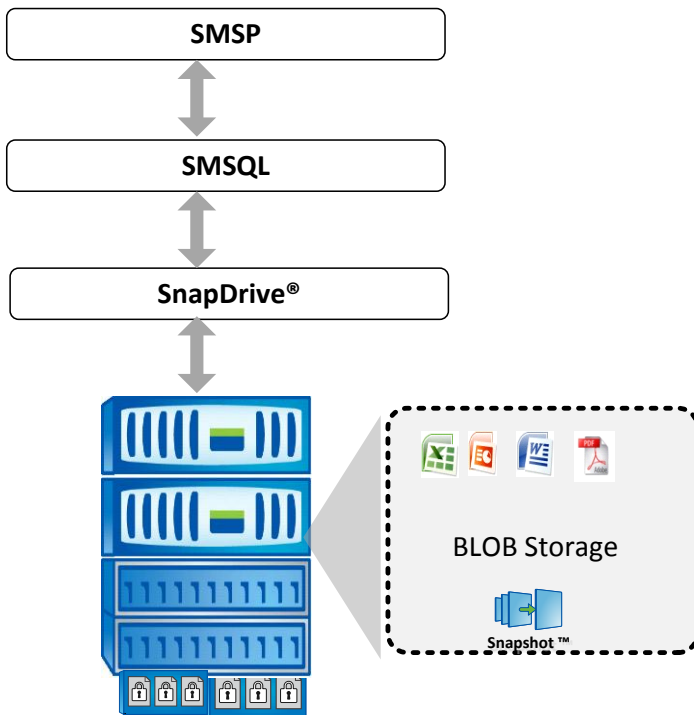
- Improves the scalability of SharePoint deployments by storing and managing large data files on networked-attached storage
- Easily migrates data from SMB or CIFS shares and Microsoft Exchange public folders to SharePoint for improved data management
- Reduces storage costs with space-efficient backup and data deduplication
- Simplifies SharePoint data management and automates routine tasks, including backup and replication
- **SnapManager for SQL Server.** This application is tightly integrated with SQL Server to help streamline database storage management while simplifying storage layout planning, backup, and restore operations for SQL Server databases.

Here are some of the important benefits of SnapManager for SQL Server:

- Reduces storage costs significantly with NetApp's space-efficient backup capabilities
- Streamlines data management and automate routine tasks to increase DBA productivity
- Increases backup frequency to protect more data automatically without affecting performance
- Restores a failed database of any size to full production in minutes
- Creates complete database clones in seconds without the need for additional storage space

Figure 6 depicts the different components that are required for SnapManager for SharePoint solution.

Figure 6) Different components required for SnapManager.

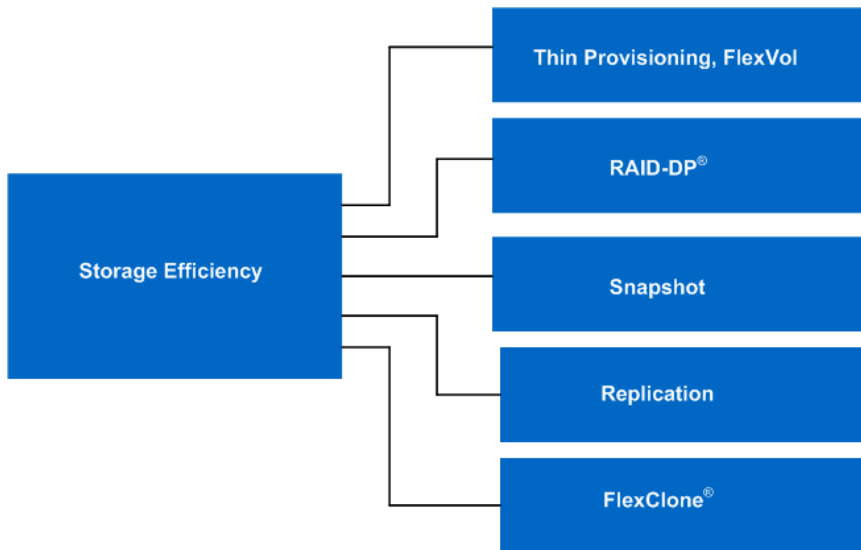


4.2 PRINCIPLES OF DESIGNING STORAGE FOR MICROSOFT SHAREPOINT 2010

Storage efficiency is the ability to store and manage data in a way that consumes the least amount of space with little or no impact on the overall performance of the system. Storage efficiency goes beyond just data deduplication; it is a combination of RAID, provisioning (overall layout and utilization), mirroring, and other data protection technologies. The deployment of SharePoint 2010 on the storage needs to take into consideration these technologies.

Figure 7 illustrates the technologies that NetApp offers to implement storage efficiency and reap its cost-savings benefits by optimizing existing storage in the infrastructure as well as deferring or avoiding future storage expenditures. The more these technologies are used in conjunction, the larger the savings.

Figure 7) NetApp technologies for storage efficiency.



RAID-DP

The performance acceleration provided by NetApp Write Anywhere File Layout[®] (WAFL) and the double-disk protection provided by NetApp RAID-DP[®] make economical and large-capacity (Serial Advanced Technology Attachment) SATA drives practical for production application use. In addition, to negate the read latencies associated with large drives, SATA drives can be used with the NetApp Performance Acceleration Module (PAM) card also known as Flash Cache, which significantly increases the performance of random read-oriented workloads.

Without high-performance double-disk RAID protection such as RAID-DP, large SATA disk drives are prone to failure, especially during their long RAID reconstruction time for a single drive failure.

RAID-DP is NetApp's implementation of double-parity RAID 6, an extension of NetApp's original Data ONTAP WAFL RAID 4 design.

Unlike other RAID and enhanced RAID technologies such as RAID 10, RAID 5, RAID 6, RAID 50, and so on, what makes RAID-DP special is its ability to provide a higher level of data protection without any performance impact, and at the same time consume a minimal amount of storage.

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

SNAPSHOT TECHNOLOGY

NetApp Snapshot technology provides zero-cost, near-instantaneous point-in-time copies of the file system (volume) or LUN by preserving Data ONTAP WAFL consistency points.

There is no performance penalty for creating Snapshot copies, because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not 100% for each backup, as with mirror copies. This means savings in storage costs for backup and restore purposes and opens up a number of possibilities for efficient data management.

STORAGE THIN PROVISIONING

In a shared storage environment, *thin provisioning* is a method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data, rather than the traditional method of allocating all the blocks up front. Thin provisioning eliminates almost all whitespace, which helps avoid

poor utilization rates. The enabling technology behind NetApp thin provisioning is Write Anywhere File Layout (WAFL), which can be thought of as the virtualization layer of Data ONTAP. When a LUN or volume is created, WAFL does not dedicate specific blocks out of the NetApp volume for the LUN or volume, or for Snapshot copies of the LUN or volume. Instead, it allocates the blocks from the NetApp aggregate when the data is actually written. This allows the administrator to provision more storage space, as seen from the connected servers, than is actually physically present in the storage system.

Pooling all available disks into a single, large aggregate may maximize performance; however it may not meet the data availability requirements set forth in the SLA.

Creating separate aggregates for SharePoint 2010 databases and transaction log or SnapInfo volumes can meet the performance requirements of SharePoint Server 2010 while still providing the data availability required by most SLAs. However, hosting the database and the transaction logs or SnapInfo volumes in a single volume can have benefits from a storage efficiency perspective. This can further be improved by the use of FlashCache or PAM module on the controller

To attain the optimal layout for SQL Server databases with NetApp storage systems, it is a best practice to have a separate aggregate for each SQL Server instance

SharePoint configuration databases are not very read or write intensive. SharePoint configuration databases and SharePoint admin content databases can be placed on a single volume and resources can be shared.

Best Practice

NetApp recommends having at least 10% free space available in an aggregate that is hosting SharePoint data. This allows optimal performance of the storage system.

Error! Not a valid bookmark self-reference. details the benefits and tradeoffs of two possible aggregate layouts.

Table 2) Aggregate layouts.

Number of Aggregates	Benefit	Tradeoff	Recommended Configuration
A single aggregate for all SharePoint data	Minimize the number of parity disks required	In the unlikely event of losing an aggregate, SharePoint data is completely lost from the affected aggregate	Smaller storage controller deployments
Multiple aggregates for SharePoint data (database and transaction logs on separate aggregates)	Dedicated aggregate handles the I/O for a database or for transaction logs	More parity disks required because there are more aggregates	Large enterprise deployments with stringent SLAs

VOLUME PLANNING

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

- This functionality results in the following additional benefits for SharePoint environments:
- A large number of volumes can be created, all with independent Snapshot[®] copy schedules, mirroring policies, and so on.

- All volumes can be managed independently, while receiving the maximum I/O benefit of a much larger pool of disks.

Volume layout is critical in creating and sustaining a highly available SharePoint environment. Careful consideration of backup groups, disaster recovery scenarios, and archiving solutions helps to determine the placement of volumes onto aggregates, and the corresponding LUNs onto those volumes. The content database data and log files should reside in separate volumes and the layout should follow the configuration rules for SnapManager for SQL Server.

Best Practice

NetApp recommends separating database and transaction logs from different servers into separate volumes. The best practice to separate volume is that we have dedicated storage for each SQL server.

VOLUME SIZING

Volume sizing has two parts: database volume sizing and transaction log volume sizing. When sizing SharePoint and SQL Server volumes, NetApp recommends calculating the appropriate amount of space needed for each LUN.

SPACE GUARANTEE

Space guarantee enables thin provisioning. This option can be set at the aggregate, volume, or LUN level, depending on the requirements of the application. If the space guarantee at the volume level is set to Volume, which typically makes sure that the amount of space required by the FlexVol volume is always available from its aggregate. This is the default setting for FlexVol volumes. With the space guarantee set to Volume, the space is subtracted, or reserved, from the aggregate's available space at volume creation time. If the space guarantee for the volume is set to None, the space is reserved from the aggregate regardless of whether or not it is actually used for data storage. The volume reserves no space from the aggregate during volume creation. Space is first taken from the aggregate when data is actually written to the volume. Write operations to space-reserved LUNs in a volume with guarantee=none will fail if the containing aggregate does not have enough available space. LUN reservation makes sure that the LUN has space in the volume, but setting guarantee=none **does not** ensure that the volume has space in the aggregate. When the space guarantee for the volume is set to File, the aggregate makes sure that space is always available for overwrites to space-reserved LUNs.

FRACTIONAL_RESERVE

Fractional_reserve is a volume option that specifies how much space Data ONTAP reserves for Snapshot overwrite data for LUNs to be used after all other space in the volume is used. The default value for fractional_reserve is 100%. However, by using the autodelete setting, the fractional reserve can be set to 0.

AUTOSIZE

This volume setting (available in Data ONTAP 7.1 and later) defines whether a volume should automatically grow to avoid filling up to capacity. This option is available only for flexible volumes. It is possible to define how fast the volume should grow by using the `-i` option. The default growth increment is 5% of the volume size at creation. It is also possible to define how large the volume is allowed to grow to by using the `-m` option. If Volume Autosize is enabled, the default maximum size to grow to is 120% of the original volume size.

Best Practice

There must be enough space available in the aggregate for the Autosize option to succeed. NetApp recommends planning for additional buffer space when using thin provisioning for SharePoint 2010 environments

Always keep the user and system databases separate. When creating new MDF and LDF files, pre-size them to minimize auto-growth events. SQL Server supports the ability to grow both data and log files automatically as they fill. Recovery models are designed to control transaction log maintenance and a SharePoint database can be configured with either FULL recovery or SIMPLE recovery mode. To know more about the database and their recommended recovery model, follow this link.

For Database types and descriptions (SharePoint Foundation 2010), visit:
<http://technet.microsoft.com/en-us/library/cc974471.aspx>.

NetApp recommends setting the file sizes appropriately and growing them manually at times of minimal system use on a planned basis. Autogrow should be used only as a safety net to prevent the files from becoming full and making the database read-only at times when unpredicted substantial growth occurs.

When database files are expanded, there is an impact on performance. This impact is minimized but not eliminated through fast file initialization.

Observe the following points about the location of files:

- MDF files should be located on their own LUN.
- LDF files should be located on their own LUN.
- BAK and TRN backup files should be located on their own disks.

Install SQL Server binaries in the `C:\Program Files\Microsoft SQL Server` folder across all the boxes. It is advisable to deploy data and log file folders in separate disks instead of in the folder path where SQL Server binaries are installed.

For maximum performance, pool all of the available disks into a single large aggregate. Creating separate volumes for the SharePoint content and SnapInfo meets the performance requirements. Use thin provision and Flexible volume in the volume planning. The thin provisioning that we are recommending is to guarantee volumes and LUNs, fractional reserve to zero with volume autosize and snapshot autodelete policies enabled.

AUTODELETE

This volume setting (available in Data ONTAP 7.1 and later) allows Data ONTAP to delete Snapshot copies if a threshold is met. This threshold, called a *trigger*, can be set so that Snapshot copies are automatically deleted under one of the following conditions:

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

The order in which Snapshot copies are deleted is determined by the following three options:

- **Delete_order.** Specifies whether the oldest or newest Snapshot copies should be deleted first.
- **Defer_deleted.** Allows the user to define a group of Snapshot copies that should first be deleted when no other Snapshot copies are available. It is possible to defer the deletion of user-created

Snapshot copies, scheduled Snapshot copies, or Snapshot copies beginning with a configurable prefix.

- **Commitment.** This option determines how Snapshot copies used for SnapMirror and dump operations should be handled. If this option is set to Try, these Snapshot copies are deleted only if they are not locked. If set to Disrupt, these Snapshot copies are deleted even if they are locked.

Best Practice

When using SnapMirror to replicate a SharePoint database, NetApp recommends not using the Disrupt option for commitment, because SnapMirror baseline Snapshot copies can be destroyed by Autodelete. The last SnapMirror Snapshot copies will always be deleted. In many configurations, deleting the last SnapMirror Snapshot copy is not desirable because a new full baseline copy is required to resume mirroring operations. For example, if the source and destination are at different sites, recreating this baseline can be a time-consuming and costly process.

CONFIGURATIONS WHEN USING THIN PROVISIONING

There are many ways to configure NetApp storage appliances for LUN thin provisioning; each has its advantages and disadvantages. It is possible to have thinly provisioned volumes and non-thinly provisioned volumes on the same storage system, or even the same aggregate. LUNs for critical production applications like SharePoint 2010 might be configured without thin provisioning, while LUNs for other types of applications might be thinly provisioned. When considering whether to use thin provisioning for a SharePoint 2010 environment, consider the following best practice configurations.

Option One: Volume Guarantee Set to None

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

This configuration has the advantage that the free space in the aggregate is used as a shared pool of free space. The disadvantage is that the high level of dependency between volumes and the level of thin provisioning cannot easily be tuned on an individual volume basis. When using this configuration, the total size of the volumes would be greater than the actual storage available in the host aggregate. With this configuration, the storage administrator generally sizes volume so that they need to manage and monitor only the used space in the aggregate.

Option Two: Using Autogrow and Autodelete

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

The advantage of this configuration is that it is possible to finely tune the level of thin provisioning for a SharePoint 2010 environment. With this configuration, the volume size defines or guarantees an amount of space that is available only to LUNs within that volume. The aggregate provides a shared storage pool of available space for all the volumes contained within it. If the LUNs or Snapshot copies require more space than is available in the volume, the volumes grow automatically, taking more space from the

containing aggregate. An additional advantage of having the LUN space reservation disabled in that case is that Snapshot copies can use the space that is not needed by the LUNs. The LUNs themselves are not in danger of running out of space because the Autodelete feature removes the Snapshot copies that are consuming space.

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

NETAPP FLEXCLONE

A NetApp FlexClone volume is a writable point-in-time Snapshot copy of a FlexVol volume or another FlexClone volume. FlexClone uses space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and the clone. FlexClone volumes are useful in any situation where testing or development occurs, any situation where progress is made by locking in incremental improvements, and any situation where it is necessary to distribute data in changeable form without endangering the integrity of the original.

FlexClone technology can be leveraged both at the primary storage system and at the SnapMirror destinations for effective use of resources. FlexClone can also be used for disaster recovery testing without affecting the operational continuity of the SharePoint 2010 environment.

FlexClone can be created from the storage controller console by using the following command:

```
vol clone create cl_vol_name [-s {volume | file | none}] -b f_p_vol_name [parent_snap]
```

Example:

```
vol clone create SharePoint_clone -s file -b SharePoint SMSP_snap
```

This example creates a new read-write FlexClone volume named `SharePoint_clone`, based on an existing volume named `SharePoint`, from the point-in-time Snapshot copy named `SMSP_snap`.

For more information about how FlexClone works and the command line reference, refer to the FlexClone documentation in the [Data ONTAP Administration Guide](#).

Best Practice

NetApp recommends creating FlexClone volumes by leveraging SnapDrive for Windows in SharePoint environments. This automates the creation of the FlexClone volumes and connecting the LUNs within the clone to the test and dev host.

NETAPP DEDUPLICATION

Data deduplication is a data compression technique for eliminating coarse-grained redundant data, typically to improve storage utilization. NetApp deduplication is a fundamental component of Data ONTAP, NetApp's core operating architecture. NetApp deduplication combines the benefits of granularity, performance, and resiliency with a significant advantage in the race to improve storage utilization demands.

The deduplication process stores only unique blocks of data in the volume and creates additional metadata in the process.

Each 4k block in the storage system has a digital fingerprint, which is compared to other fingerprints in the volume. If two fingerprints are found to be the same, a byte-for-byte comparison is done of all bytes in the block. If they are an exact match, the duplicate block is discarded and the space is reclaimed.

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

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

Note:

- Deduplication is transparent to SharePoint, which does not recognize the block changes, so the SharePoint database remains unchanged in size from the host, even though there are capacity savings at the volume level.
- Tests show that space savings on a CIFS share can be up to 35% to 40%, and space savings in SharePoint content databases are approximately 30% on average.
- Adequate space must be available on the flexible volume for the `sis on` command to complete successfully. If the `sis on` command is attempted on a flexible volume that already has data and is completely full, it fails because there is no space to create the required metadata.

Deduplication Metadata Overhead

Although deduplication offers substantial storage savings in many environments, there is a small amount of storage overhead associated with it that should be considered when sizing the flexible volume. The total storage used by the deduplication metadata files is approximately 1% to 6% of the total data in the volume.

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

Best Practice

Replication of deduplicated volumes is supported by using SnapMirror. Do not use synchronous SnapMirror to replicate volume that use deduplication.

SIZING THE MEDIA SERVER

The Media Server is a component of the SnapManager for SharePoint architecture. In SMSP 6.1, the Media Server is responsible for storing a lot more data than in the previous versions of SMMOSS, 2.X and 6.x.

The Media Server manages the following items:

- Backup job metadata, backup index
- Backup IIS metadata
- Backup SharePoint 12 Hive or 14 Hive
- Archived content from Archiver and Extender

The size of the backup-set indexes created by the Media Server depends on the level of granularity chosen when creating the backup set. There are five levels of granularity in SMSP: site, Web, folder, item, and item version.

- **Site.** This level of granularity makes sure that individual root-level sites within the SharePoint Web application can be recovered.
- **Web.** Makes sure that subsites under the root-level sites can be recovered.
- **Folder.** Makes sure that document libraries and lists can be recovered.
- **Item.** Makes sure that each item in the SharePoint Web application can be individually recovered, including lists, announcements, documents, calendars, events, and so on.
- **Item version.** Makes sure that all versions of each item (especially documents) in the SharePoint Web application can be individually recovered, including lists, announcements, documents, calendars, events, and so on.

In a normal SharePoint Web application, as the level of granularity becomes finer, the number of objects that must be indexed increases, and therefore the size of the index increases.

Normally, it is difficult to get a count of the number of objects at each level of granularity, which makes sizing the Media Server very difficult. However, the size of each entry in the index is roughly 300 bytes. This information can be used to size the Media Server data store. Note that the change in the number of objects increases as the level of granularity becomes finer. This means that it is possible to closely determine the size of the Media Server indexes up to the folder granularity level by taking a few runs of the backup plan set at the required granularity levels and then checking the size of the index folder that is created.

The folder where SMSP stores the index created as a part of running a backup plan is located at `<Media Server backup data path>\<agent name>\<Plan ID>\<Job ID>`. To find the PlanID of a backup plan, click the View hyperlink of the job status row (in the BackupJob monitor) for the job. The same row also contains the JobID column.

Sizing for SMSP Backup Index

Let the number of backup plans that use the Media Server = N

Let the number of online backups per backup plan = X (averaged)

Let the number of archived backups per backup plan = Y (averaged)

The LUN size for the Media Server datastore can be approximated as:

$$\text{LUN_SIZE} = N \times (X+Y) \times 0.1 \times (\text{Avg. content DB size for the databases involved in the backup plans})$$

It has been observed that, on average, each backup index consumes up to 10% of the content DB size.

Therefore the volume size for the Media Server datastore can be approximated as:

$$\text{VOLUME_SIZE} = (\text{LUN_SIZE} \times 1.1) + [(\text{Number of Snapshot copies to keep online}) \times (0.1 \times (\text{Avg. content DB size for the databases involved in the backup plans}))]$$

Assumption: 0% fractional reserve.

Sizing for BLOB Storage

- Size of CIFS volume = [active BLOBs + orphan BLOB within delayed deletion window] + Snapshot copies x Snapshot average size
- Size of LUN = 1.1 x [active BLOBs + orphan BLOB within delayed deletion window] + Snapshot copies x Snapshot average size

Note: BLOB within delayed deletion window is the Garbage collection for SMSP. The important thing to notice is that the garbage collection window needs to be greater than the SMSP backup retention.

Sizing for BLOB Store Index

- General index is about 1% of BLOB size
- Full-text index is about 15% of BLOB size

DESIGNING STORAGE LAYOUT FOR SHAREPOINT DATA

Storage must be carefully planned and laid out for the SharePoint content, considering the objectives of faster access and recoverability. SharePoint content can be categorized in two categories: content databases and SharePoint BLOB data on external Server Message Block (SMB) shares.

Content Databases

SharePoint content databases that need to be migrated to NetApp storage system must first be laid out according to the best practices for keeping the SQL databases on the NetApp LUNs. In a large SharePoint deployment, it is common for the Web applications to have multiple content databases, as

recommended by Microsoft. For SharePoint 2010, the suggested content DB limit for optimum performance is 200GB. This could go up to 1 TB for environments with noncollaborative I/O and usage patterns, such as record centers.

Best Practices

To improve the backup and restore performance for content databases that belong to the same Web application, follow these guidelines:

- Multiple content databases should be stored on a single LUN to use fewer mount points for large SharePoint deployments. A maximum of 35 databases can be stored in a single LUN or a single volume.
- Storing multiple databases in a single LUN or volume results in faster backups due to fewer backup groups, because fewer volume are used for all the databases. However, the databases on the LUN should be dictated by the Recovery Time Objective of the databases.
- Similarly, each Web application should be allocated a dedicated volume for all its content databases.
- LUNs should be created on the faster disks, such as Fibre Channel disks or SAS disks with 15k rpm.
- Restores tend to be faster when the entire LUN are restored. This is because SMSP can take advantage of LUN Clone Split Restore for much faster restores. Restoring a single database from a LUN is snapshot can take a long time since it is basically a big file copy
- Consolidation of multiple LUNs on a single volume can ease monitoring and administrative activities at the storage level.

SharePoint BLOB Data on External CIFS Shares

The BLOB data, archived from the SharePoint content databases, should be placed on a NetApp NAS device (that is, the CIFS shares). The key factors for designing the layout of the CIFS/SMB shares is faster recoverability.

Best Practices

- Create a CIFS share in its own dedicated volume, for leveraging NetApp SnapRestore[®] at the volume level for faster restores.
- Compression and de-duplication can be used together and are transparent to Microsoft SharePoint.
- Create multiple CIFS shares, one per volume, for each Web application in each farm.
- Create separate CIFS shares for archived BLOB data and for extended BLOB data due to their different I/O activity patterns. This restriction makes it easier to organize and manage the archived data based on the archiver/extender index, instead of having data from the same source scattered among different devices
- Do not share the BLOB storage NAS volumes with any other data or LUNs.
- Keep multiple CIFS shares available for each Media Server to store the BLOB data. Make sure that the security should be set to the SMSP user account used to access these shares.
- Each volume designated for CIFS share should be minimum 1GB. There is no defined maximum limit; the maximum depends on the number of volumes carved out from the storage.
- CIFS shares can be created on the volumes/aggregates residing on the SATA disk drives.

HANDLING CONCURRENT REQUESTS

In large SharePoint environments, it is possible that many users may access or upload archived content concurrently. To reduce the impact on the system, SnapManager for SharePoint has a throttling control mechanism running over these concurrent requests. By default, the maximum number of concurrent

requests is 20 per Web front-end server and 50 per SnapManager Media Service, based on a server configuration of PIII-700MHz CPU and minimum 2GB of RAM. If there are more concurrent requests than this limit, they are put into a queue and processed after current requests are served. This setting should be sufficient for most situations and can be adjusted if needed. (Consult technical support about the best setting for your environment).

Default Maximum Concurrent Number of Requests

- **On Web Front-End Agents:** 20 requests
- **On Media Service:** 50 requests

Generally these limits are related to the server hardware specification and server load at the time of operation, so there is no fixed limit for SMSP. The concurrent limit can be changed by taking the following points into consideration, because they are mainly related to available sockets, threads, and memory on the system.

- Each user download (Archiver/Extender stub data) requires:
 - Two sockets on Archiver Agent, one socket on Media Service
 - Two threads on Archiver Agent, two threads on Media Service
 - 320KB memory on Archiver Agent, 128KB memory on Media Service
- Each user upload (Extender) requires:
 - Two sockets on Archiver Agent, one socket on Media Service
 - Three threads on Archiver Agent, two threads on Media Service
 - 320KB memory on Archiver Agent, 128KB memory on Media Service
- Each end-user archiving request requires:
 - Two sockets on Archiver Agent, one socket on Media Service
 - Three threads on Archiver Agent, two threads on Media Service
 - 320KB memory on Archiver Agent, 128KB memory on Media Service

The memory size is mainly the buffer size used for asynchronous transfer; it can be increased if the network condition is not optimal. The concurrent limit can also be changed according to environment needs. For deployment with a large user base, NetApp recommends using SMSP Media Service on 64-bit servers, where Java® JVM has higher resource limits. Media Server with 16GB memory and dual core CPU can support about 300 to 500 concurrent user requests.

One of the most important decision points is the number of Media Servers and managers to install. In making that decision, follow these guidelines.

Best Practices

- SnapManager for SharePoint is cluster aware and all of the components like the manager and the media server can be made Highly Available.
- Alternatively, Media Servers can be deployed on virtual machines (VMs). NetApp recommends provisioning at least two Media Servers in a VM-based setup. The benefit of using a virtualized platform for deploying Media Servers is the ease of scaling out. This means that if the resources on a Media Server are maxed out, additional Media Servers can easily be provisioned.
- The processing activities performed by the Media Server are resource intensive. The most important resources are CPU (the process is CPU intensive) and memory (the process is multi-threaded and deals with fair amount of data). Media Server uses at least 35% of CPU time during its processing cycles; therefore NetApp recommends a dual core CPU and a minimum of 16GB memory.
- A single SMSP Manager can manage multiple farms. At a minimum, one SMSP Manager is required for the entire SharePoint infrastructure in a single farm. Ideally, the manager should be installed on a physical server. However, it can be installed on a virtual machine as well.

5 NETAPP SOLUTION FOR SHAREPOINT SERVER

Because SharePoint Server is an enterprise-level application, there should be an enterprise-level content management solution for managing the SharePoint content. The NetApp solution for SharePoint offers data protection, including backup and recovery, and storage optimization and migration as well.

5.1 SNAPMANAGER 6.1 FOR SHAREPOINT SERVER OVERVIEW

SnapManager 6.1 for SharePoint Server is an enterprise-level backup and recovery solution for Microsoft SharePoint versions including Windows SharePoint Services 3.0, Microsoft Office SharePoint Server 2007, SharePoint Foundation 2010, and SharePoint Server 2010. This management software can make life easier for SharePoint administrators and lower the data center administration cost as well.

SnapManager for SharePoint can perform full SharePoint farm-level backups and restore different levels of SharePoint components, from a single document up to the contents of the entire farm.

SnapManager 6.1 for SharePoint Server offers the following capabilities for SharePoint farms:

- Integrates with and leverage the EBS/RBS framework from Microsoft in SharePoint 2007 and 2010 and SQL Server 2008
- Enhanced command line support for migration and data protection along with the integration with windows Powershell v3
- Provides content archiving and content migration capabilities
- Centralized management of backup and recovery for multiple SharePoint farms
- Ability to back up all objects and entities of a SharePoint farm
- Integrates with NetApp SnapMirror and SnapVault
- Integrates with NetApp SnapLock[®] technology
- Integrates with Microsoft operation manager MOM and System Center operation manager SCOM event log, NetApp ASUP[™] (AutoSupport[™]) system for streamlined support

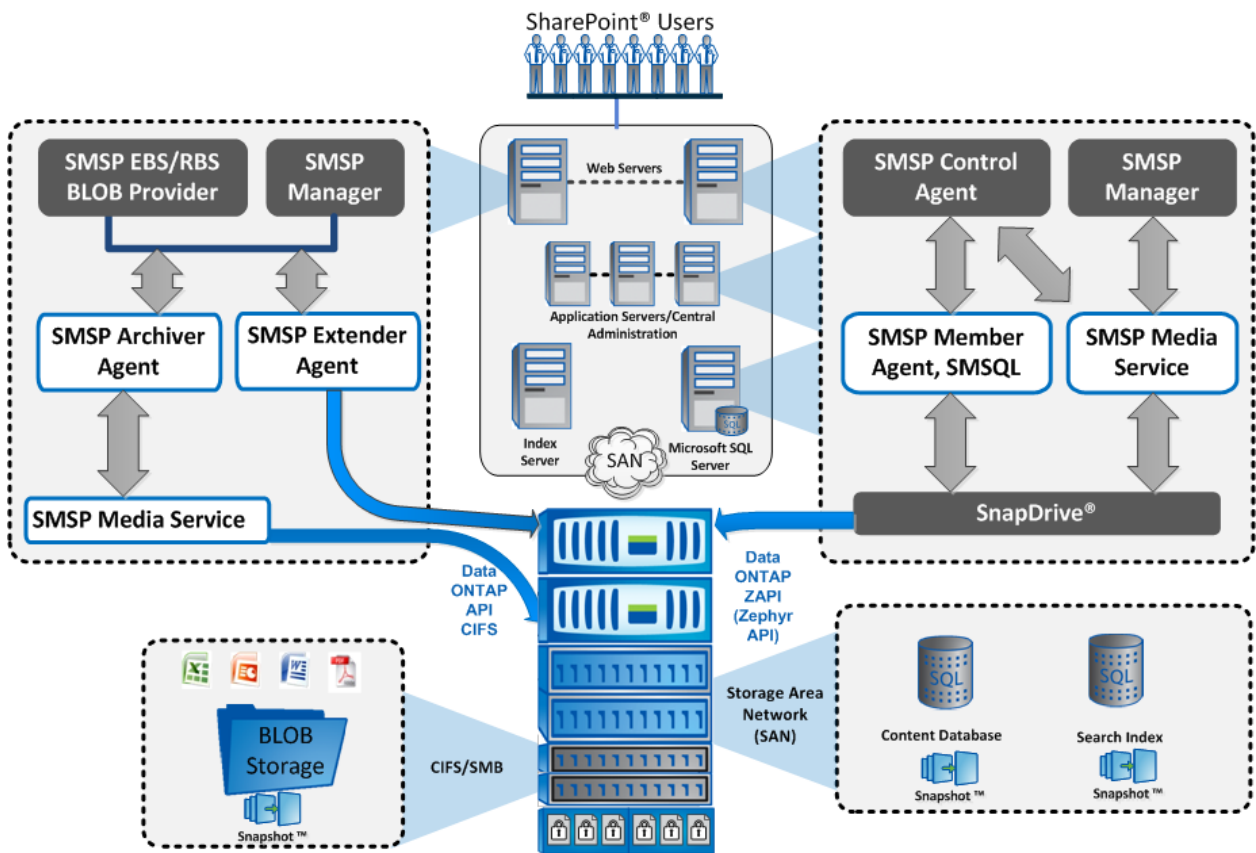
5.2 SNAPMANAGER 6.1 FOR SHAREPOINT SERVER ARCHITECTURE

SnapManager for SharePoint Server is a 64-bit application that is capable of protecting SharePoint 2007 and SharePoint 2010 databases. It was designed with the following objectives:

- To provide content archiving and content migration capabilities
- Centralized management of backup and recovery for multiple SharePoint farms
- Ability to back up all objects and entities of a SharePoint farm
- Integration with NetApp SnapMirror and SnapVault

Figure 8 illustrates the SnapManager for SharePoint data protection and storage optimization architecture showing the relation with the SMSM Manager, Media Server, Control Agent, Archiver/Extender Agents, Member Agents, and so on.

Figure 8) SMSP data protection and storage optimization architecture.



To meet the objectives listed under section 5.2, SMSP uses an agent-based architecture. These agents help provide centralized management, and they automate most backup and recovery tasks.

This section describes the components of SnapManager for SharePoint:

- **SMSP Manager.** This component provides central backup and restore management by using the services of the control and member agents (discussed later in this section). It also provides the central graphical user interface that is used to initiate backup and restore tasks for SharePoint Web applications.
- **SMSP Media Server.** This is the most important component of SMSP, responsible for managing data in SMSP. It performs the following key functions:
 - Generates and stores backup set indexes and metadata
 - Manages all data archiving and data retention activities, including restore of CIFS locations
 - Archived Data is transferred to the media server directly from related agents, but not for the extended content.
- **SMSP Control Agent.** SMSP Control Agent is part of the SMSP Data Protection module. It runs as a service on each SharePoint Web front-end (WFE) server and is responsible for discovering the SharePoint Web applications that run on that WFE. It is also responsible for coordinating backup and restore tasks for the Web applications on its WFE server. It receives task control from SMSP Manager, analyzes the SharePoint farm structure, and sends the job control to related SMSP Member Agents in the farm to complete the operations.
- **SMSP Member Agent.** SMSP Member Agent is part of the SMSP Data Protection module. Its main purpose is to interact with SharePoint related components to perform backup and restore operations. There are three types of Member Agents:

- Member Agent. Installed on SQL Server, uses the SMSQL cmdlet to backup and restore databases. It also generates a backup index from a database Snapshot copy and sends it to SMSP Media Service.
- Member Agent. Installed on SharePoint Index Server, uses the SnapDrive CLI to create Snapshot backups of index LUNs.
- Member Agent. Installed on SharePoint front-end server, can back up IIS settings, SharePoint 12 Hive or 14 Hive, and resources in the file system. Backup data is transferred to SMSP Media Service.
- **SMSP Archiver and Extender.** These components enable policy-based content archiving, content loading, and content migration to SharePoint farms. These components can be incrementally enabled in the SMSP Agent configuration tool on appropriate hosts.

Archiver ages out content based on business rules set by the customer. The scope of the Archiver rules are granular and allows to create rich business rules around what content is Archived and eventually aged out.

It is recommended that the customer have good record retention policies in place prior to using Archiver.

The general value proposition of Archiver is that you are using less storage capacity because you are no longer storing stale unused content.

Extender is a scalability and storage efficiency feature. The extender rules are based on content size. Both new and existing content can be externalized. The scope of the Extender rules (with RBS) are at the database level and are not granular.

The general value proposition of Extender is that you are using smaller SQL databases and can store BLOBs on SATA. This allows for better scalability and when used with NetApp storage, you can leverage storage size compression and de-duplication

- **SMSP EBS/RBS Provider.** The SMSP EBS Provider makes it possible to interact with externalized SharePoint content by leveraging the EBS framework provided by SharePoint Server 2007 SP1 or the RBS framework provided by SQL Server 2008 for SharePoint Server 2010.

Note: SMSP does not support any third –party EBS/RBS providers.

NETWORK PORT USAGE AND SECURITY IN SNAPMANAGER 6.1 FOR SHAREPOINT

All SMSP components use TCP/IP ports to communicate with each other. This method makes sure that communication between the agents can flow easily even with complex deployment scenarios such as demilitarized zone (DMZ) deployments where the SharePoint WFE server is typically present in an external public domain and the application and database servers are in the internal private domain.

SMB protocol is the basis for Microsoft file and print sharing and many other networking operations. To prevent attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. To use it on the filer SMB signing has to be enabled on the NetApp filer as well as on the WFE servers.

Note: Performance impact on the controller and the WFE server has to be accounted before implementing SMB signing. There is a performance overhead of 10%-15% after SMB signing is enabled. Make sure that the server and storage controllers are sized to accommodate SMB signing.

For a detailed description of the ports used by SMSP, see pages 16 and 20 of the [SnapManager 6.0 for SharePoint Installation and Administration Guide](#). Table 3 summarizes the ports used by the SMSP components.

Table 3) Ports used by SMSP components.

Component	Service Name	Default Port	Description
Manager	Web Service	8080 (HTTP port for external communication)	Web server for admin console
		8443 (HTTPS port for external communication)	
		12002 (internal communication) 12011 (access the CLI)	
	Media Service	12001	Media service controls messages and data port
	Control Services	12000	Communication port with agent
Agent	Communication Service	10103	Communication port for agents
		10107	Communication port for Archiver, Migrator, Extender

Best Practice

It is recommended that the SMSP default port settings should be used as much as possible. This helps to reduce the time needed to configure agents' port settings and helps to avoid errors in that process.

Generally, SMSP Manager can use a range of the port numbers 12000-12012, but if there is a conflict with any other application, then you can use a port number beyond this range as well.

Note: Customer with Co-existence of SnapManager for SharePoint with Avepoint DocAve, note that the ports will need to be changed if DocAve is used along with SnapManager for SharePoint in the SharePoint Farm.

SECURED NETWORK ACCESS

In the case of a secured network environment using firewalls, the SMSP Manager and Agent ports should be open for the communication.

For remote installation, the secured network relies on Windows Management Instrumentation on the agent machine. The following ports should be open through the firewalls:

- Remote Procedure Call (RPC): 135
- Windows Management Instrumentation: 445

Connection to the NetApp storage controller requires the following ports to be open through the firewalls.

- HTTP:80
- HTTPS:443
- RPC: 135

For SnapDrive for Windows and SMSQL:

- 808

SECURITY MODEL

SMSP has a very comprehensive security model that supports two modes of authentication:

- **SMSP-based authentication.** In this model, the logins are specific to SMSP and are internal to it. This is referred to as the *local system* mode of authentication.
- **Domain-based authentication.** In this model, the logins used are domain-level logins and therefore are authenticated by the active directory. This is referred to as the *active system* mode of authentication.

Best Practice

NetApp recommends using domain-based authentication to integrate SMSP authentication with the active directory system. This eases the security administration for SMSP.

The SMSP security principles are categorized as follows:

- **User.** These are logins that are used to gain access to the SMSP interface. As discussed earlier in this section, logins can either be SMSP-based (local) or AD-based (domain).
- **Groups.** Groups are specialized roles in SMSP that have certain specific access rights. There are three types of groups in SMSP: SMSP administrators, SMSP operators, and SMSP managers. Each user must be part of one of these three groups. SnapManager administrators are the group that has admin privileges, and the initial default local login administrator is also a member of this group.

Security in SMSP is managed through the Account Manager portion of the SMSP GUI, which is accessed through the Control Panel. Figure 9 shows the Account Manager section.

The SMSP-based username admin (password is also admin) is the default login whenever SMSP is installed. This login is a member of the SMSP administrators group, and is the only one that can be used to log in to SMSP Manager for the first time. After that, other users can be created and placed in one of the three user groups just described.

Figure 9) Account Manager.

The screenshot displays the 'Users and Groups' management interface. On the left is a navigation sidebar with categories like Data Protection, Storage Optimization, Job Monitor, Control Panel, and SMSP Services. The main area is split into two panes: 'System Defined' and 'User Defined'. The 'System Defined' pane contains a table with columns for 'Account Type' and 'Account Name', listing 'Admin', 'Administrators', 'Managers', and 'Operators'. The 'User Defined' pane is currently empty. To the right of the table is a 'General' tab with a form for creating a new user. The form includes fields for 'Username' (pre-filled with 'admin'), 'Old Password', 'New Password', 'Confirm Password', and 'Email'. 'Save' and 'Reload' buttons are located at the bottom right of the form.

Account Type	Account Name
	Admin
	Administrators
	Managers
	Operators

General Member Of

Username:

Old Password:

New Password:

Confirm Password:

Email:

Save Reload

Best Practice

Change the default password for the SMSP-based admin account immediately after installation. This will help avoid any security issues after the installation.

5.3 PREPARING SHAREPOINT DATABASES FOR MIGRATION TO NETAPP STORAGE

SharePoint databases must be located on a NetApp system (LUNs) to be backed up by SnapManager for SharePoint. If the databases are not located on the NetApp LUNs, they must be migrated to the NetApp LUNs. There are two ways to migrate the SharePoint databases:

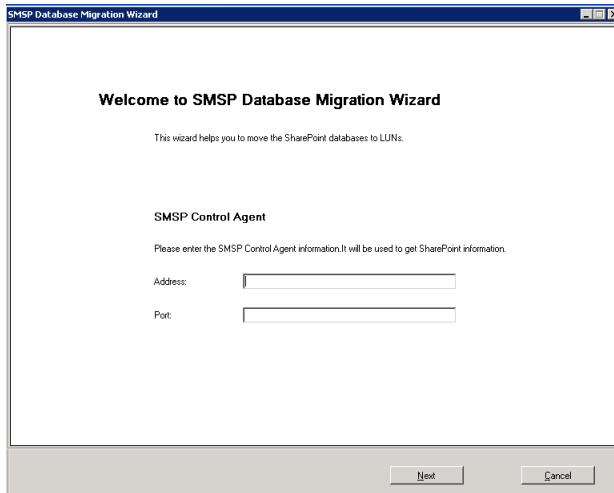
- SharePoint Database/Index Migration Wizard in SnapManager for SharePoint
- Configuration Wizard in SnapManager for SQL Server

SHAREPOINT DATABASE/INDEX MIGRATION WIZARD IN SNAPMANAGER FOR SHAREPOINT

This is a single interface to seamlessly migrate both SharePoint content databases and indexes to NetApp LUNs. This interface can be accessed from the SMSP Agent Tools under the SMSP group in Program Manager. Depending on the SQL Server and/or SharePoint Server being installed on the host, select the SharePoint Database/Index Migrator for SharePoint 2007 or 2010 and then, in the wizard, select the databases or index migration options.

The Migration Wizard will migrate SharePoint databases and search indexes. Alternatively, these content databases can be migrated to NetApp LUNs by using SMSQL in the following order:

1. Log in to the SharePoint Database/Index Migration Tool.



2. Stop all the following services in the SharePoint farm:

Table 4) SharePoint Servers 2007 and 2010.

SharePoint Server 2007	SharePoint Server 2010
Windows SharePoint Services Administration	SharePoint 2010 Administration
Windows SharePoint Services Timer	SharePoint 2010 Timer
Windows SharePoint Services VSS Writer	SharePoint 2010 VSS writer
Windows SharePoint Services Tracing	SharePoint Tracing

SharePoint Server 2007	SharePoint Server 2010
IIS Service	IIS Admin Service
Office SharePoint Server Search	SharePoint Server Search 14
Windows SharePoint Services Search	SharePoint Foundation Search V4

- Use the SMSQL Configuration Wizard to move the SharePoint databases, such as the configuration database, content database, central administration console database, search database, SSP database, and so on.



Note: If it is intended to use the SnapVault integration feature of SMSP 6.1 to archive backups of the SharePoint databases, then make sure that the LUN is a qtree-based LUN and is the only LUN residing on the qtree.

- Start the services listed in step 2.

Best Practice

SharePoint VSS service is always manual; do not change it to automatic or start it manually.

5.4 UPGRADING SMSP 6.0 TO SMSP 6.1

SMSP 6.1 supports seamless upgrade from the earlier version of SnapManager for Microsoft Office SharePoint Server (SMMOSS) 2.0 and later. To perform a smooth upgrade of SMSP 6.1 Manager and Agents, upgrade the manager first before upgrading the agent. During the upgrade process, SMSP automatically uninstalls the previous version, and previously created plans and configurations are automatically upgraded as well. Uninstalling the legacy version does not remove configurations or backup data.

When upgrading from SMSP 6.0 and SharePoint 2007 with EBS-enabled sites to SMSP 6.0 and SharePoint 2010, the EBS remains enabled on those sites. If SMSP RBS Provider is enabled after the

upgrade on other content DBs in other sites, it becomes active, is used for BLOB creation and retrieval from those sites, and can co-exist with EBS in SharePoint 2010 environments.

Note: Configuration data for SMSP is stored in the directory `SnapManager for SharePoint Server\VaultServer\ SMSPData`.

Best Practice

Take a system backup of SMSP every time you make any changes in the SMSP Manager configuration for rollback purposes by using the System Recovery feature in the SMSP Manager Control Panel.

Some examples of different stages:

1. Once SMSP agents and managers are installed and configured.
2. When RBS provider is installed and configured on specific databases.
3. When any additional databases are added and RBS is enabled on them.
4. When any changes are made to the backup schedules, backup plan configuration, archiver schedule, retention, and so on.
5. When any media servers or agents are added or removed.

5.5 INSTALLING SNAPMANAGER 6.1 FOR SHAREPOINT SERVER IN A NEW SHAREPOINT ENVIRONMENT

SMSP has a number of components that serve various functions. Some of these components—the Archiver and Extender agents—are optional and can be enabled after installation.

To make sure that SMSP is correctly installed in the SharePoint environment, use the information in Table 5 makes you understand the role of the servers in a SharePoint farm and the SMSP components that should be installed on them.

Table 5) SMSP components mapped to SharePoint farm hosts.

Server Role	SMSP Component	Remarks
Central Administration Server	SMSP Manager, Control Agent	Mandatory
Media Server	Media Service	Mandatory. Should be installed on every media server by using SMSP Manager software.
Web Front-End (WFE) Servers	Archiver Agent	Optional. Enable only if you want to archive the contents of one or more of the Web apps hosted on the WFE.
	Extender Agent	If you are not using EBS/RBS you do not need to enable it.
Index Server	Member Agent	Install only if you want to back up the SharePoint search indexes and search databases.
SQL Server host	Member Agent	Mandatory

5.6 SECURITY PERMISSION REQUIRED IN THE CONFIGURATION OF THE USER ACCOUNT IN SQL

There are a set of permissions that are required on Microsoft SQL Server for the SnapManager for SharePoint service account.

The following permissions for the SnapManager for SharePoint Accounts are:

- The SnapManager for SharePoint account must be a member of the Farm Administrator Group in Central Administration
- This account must be in a Farm Administrator group as a member, but not a member of a group which is then applied into this group.

SQL permissions required for the “same” SMSP account as above are:

- SYSADMIN

Figure 10) Login properties-DETDC\administrator.

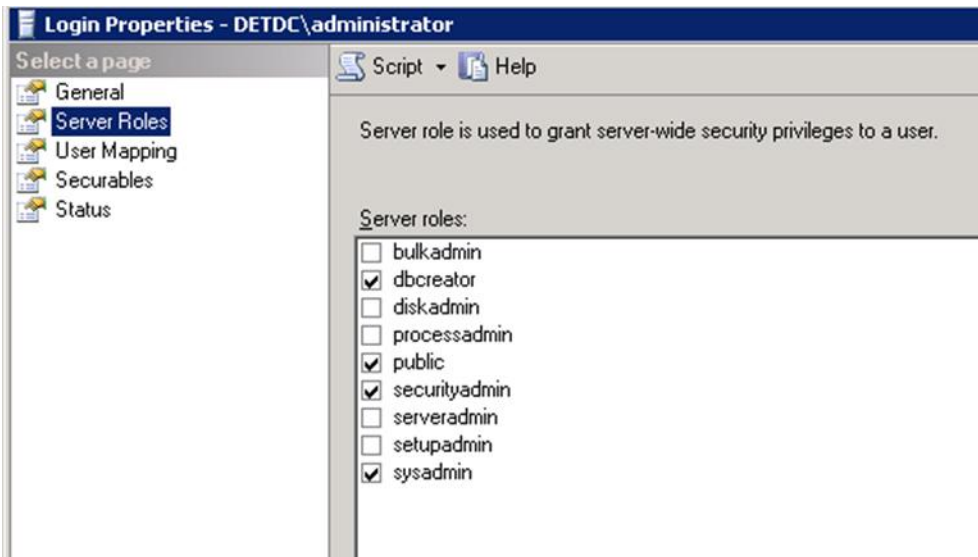


Figure 11) Windows permission.

Permission	Grantor	Grant	With Grant	Deny
Create DDL event n...	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Create endpoint	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Create trace event ...	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
External access ass...	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Shutdown	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Unsafe assembly	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
View any database	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
View any definition	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
View server state	DETDC\administrator	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Windows Permissions for the "same" SMSP account as above are. Local Administrator for any servers that have SMSP agents installed, including:

- SQL Servers
- Application Servers
- SMSP Manager Server
- Web Front End Servers

Where Group Policy is applied to Windows machines with restrictive settings, the Local Administrator does not suffice. Edit the local security policy (and/or default domain policy) for the above servers and add the following:

- Log on Local
- Log on as Batch Job
- Log on with System Interactively
- Log on as a Service

5.7 STORAGE OPTIMIZATION

SMSP 6.1 includes storage optimization solutions that provide the tools you need to keep your SQL Server resources optimized with intelligent archiving and real-time BLOB offloading to a CIFS share on lower cost storage. BLOB offloading is done in two ways, either using EBS (External BLOB Storage) Provider or RBS (Remote BLOB storage) Provider technology in SharePoint. In SMSP 6.1, EBS and RBS Providers are installed as part of the Archiver/Extender Agent on the SharePoint Web front-end servers.

EBS VERSUS RBS

EBS is available in both SharePoint 2007 SP1 and SharePoint 2010. It is enabled at farm level, providing Extender granularity at the site collection level, which means that different rules can be applied to individual sites.

RBS is available in SharePoint 2010, based on the API supported by SQL Server 2008 and SQL Server 2008R2. To run RBS on a remote server, you must be running SQL Server 2008R2 Enterprise edition. The advantages of RBS over EBS include much tighter integration with SQL Server 2008R2 ENT. Microsoft recommends using RBS for any new deployments, because EBS is supported until SharePoint 2010 only. RBS also has the following limitations:

- Extender granularity is limited to the content database level, so it is not possible to apply different rules at the site level.
- RBS with SQL Server native file stream supports only the local file system. The RBS Provider in SMSP 6.0 does not have file stream and can store the BLOB on a remote SMB share on NetApp storage system.
- RBS is not supported with SQL Server 2005. SQL Server 2008 Enterprise is required for RBS.
- To perform migration from SharePoint 2007 EBS stubs to SharePoint 2010 RBS stubs, you must upgrade the SharePoint 2007 EBS stubs to SharePoint 2010 EBS stubs first and then upgrade the SharePoint 2010 EBS stubs to SharePoint 2010 RBS stubs.

SnapManager for SharePoint 6.1 will allow the migration and upgrade capability from EBS to RBS:

- EBS to RBS SharePoint 2010
- SQL RBS File stream to SnapManager for SharePoint RBS
- SharePoint 2007 to SharePoint 2010 EBS

If EBS and RBS coexist in a SharePoint farm, then the workflow for uploading or downloading the content is slightly different.

- **For downloads.** SharePoint calls the right provider from where the stub was created.
- **For uploads.** If RBS is enabled, BLOB is stored through the RBS Provider. If RBS is not enabled but EBS is enabled, BLOB is stored through the EBS Provider.

Best Practices

BLOB Storage Guidelines

- Using EBS or RBS is beneficial in the following cases:
 - The existing content databases are 500GB or larger. Using EBS or RBS would be helpful in bringing the content databases down to an optimal size, as supported by Microsoft.
 - The BLOB data files are larger than 256 kilobytes (KB).
 - The BLOB data files are at least 80KB and the database server is a performance bottleneck. In this case, EBS/RBS reduces the both I/O and processing load on the SQL Server database server.
- If the current SharePoint version is 2007 or earlier, you want to move to SharePoint 2010, and you have no externalized content and then choose RBS to store the BLOBs.
- If your current SharePoint version is 2007 SP1 with EBS enabled, and you want to move to SharePoint 2010, then the best option is to use SMSP 6.1 and perform EBS-to-RBS direct stub migration options are available.
- If your current SharePoint version is 2007 SP1 with EBS enabled, and you really need to move to RBS in SharePoint 2010 by using SMSP 6.1, then you would follow the procedure described in section 8.1, "Migration from Different Source Environments to SharePoint."

DEVICE MANAGER IN SMSP

SMSP Media Service manages data by using devices. There are four device types:

- **Backup device.** For backing up job data; can be on local disk, LUN, or UNC path, NetApp recommends using a NetApp LUN.
- **Archived index device.** For Archiver/Extender index and full text index; can be on a LUN or SMB share on NetApp storage. NetApp recommends using a NetApp LUN.
- **Archived data device.** For Archiver/Extender data; can be on a LUN or SMB share on NetApp storage. NetApp recommends using a NetApp SMB share.
- **SnapLock archived data device.** For Archiver data with compliance requirements; must be NetApp SMB share.

In each device type, multiple physical devices can be defined. Each physical device is assigned an order ID. Before running a job or writing a file, SMSP Media Service checks that there is enough free disk space. If the free disk space of the device is less than 1GB (configurable), data is written to the next physical device on the list. If the number of files reaches the maximum limit for the folder or volume, new files are created in the next physical device on list. The SMSP Manager checks the space of the physical device every 30 minutes. When there is no free space in the device, if there's enough space on the storage side, it is possible to expand the same device space.

On each physical device, the following folder structure is used for the archived data device and the SnapLock archived data device:

- **Archiver.** `data_archiver\DataVolume\farm\webapp\site collection`
- **Extender.** `data_extender\DataVolume\farm\webapp\site collection`

The archived index device folder structure is:

- **Index.** data_archiver\IndexVolume\farm\webapp\site collection
- **Full text index.** full_text_index\farm\webapp\site collection

SnapManager for SharePoint 6.1 allows administrator to configure logical device for backup jobs and archive jobs. There are five kinds of logical device

- Local
- Network
- LUN
- CIFS Share
- Snaplock

For each logical device type, multiple physical devices can be defined. Once a physical device is full, data will automatically be written to the next physical device with enough free space.

Best Practices

- Make sure that the volumes for the SMB shares are big enough to hold a large amount of BLOB data.
- Create multiple physical devices with the SMB shares with dedicated volumes (one SMB share per volume) for fast recovery with SnapRestore.
- Use different media server and BLOB stores if you want to keep BLOB from different content databases separate.
- If you expanded the device space and want to reuse the device immediately, select the device in Device Manager and save it again. Otherwise the device will be reused after the checking.
- If EBS and RBS coexist in a SharePoint farm, create separate physical devices for them.

5.8 BACKUP GUIDELINES

SQL database backup and BLOB storage backup serve different purposes in SMSP. Database backup is for both disaster recovery (DR) and non-DR (item-level recovery). BLOB storage backup is only for DR. Following are some backup schedule best practices.

Best Practices

- Increase the frequency of backups of the SQL databases that hold the metadata. NetApp recommends taking these backups every couple of hours.
- Create separate backup plan for backup of WFE servers from the database backup. This will help the SMSP backup index from getting bloated, especially if you have large amount of GAC data.
- Backup maintenance plans are very important. As database verification and granular restore index generation during backup job create overhead in the server and increases the time for backup job. Always perform this operation at a different time from the backup schedule.
- Create a single backup plan for SharePoint database and BLOB data in SMSP 6.1 as the BLOB data is backed up with the SharePoint content

SHAREPOINT COMPONENTS BACKUP

The following SharePoint 2007 and SharePoint 2010 components are covered in the SMSP backup plans:

- All SharePoint databases
- Project Server 2007 databases

- FAST Databases
- Nintex Databases
- SharePoint search index
- SharePoint components and settings
- SharePoint solutions
- SharePoint front-end resources, including IIS settings, inetpub folder, SharePoint Hive, and Global Assembly Cache (GAC)

SnapManager for Microsoft SQL Server is used to perform database Snapshot copy backups, and SnapDrive is used to perform Snapshot copy backups of the search index. Backup data of other SharePoint components is sent to the SnapManager for SharePoint Media Service, where it is stored together with backup job metadata and index.

By default, the backup data is stored locally at the following location:

C:\Program Files\NetApp\SnapManager for SharePoint\Vault Server\Media\data-backup

SMSP does not create Snapshot copies of the backup data on the Media Server volume as part of the backup process; it is streaming-based backup.

For disaster recovery purposes, this backup data should be placed on a NetApp LUN or CIFS share, which can be automatically replicated to a SnapMirror destination volume. For longer retention of the backup data, database backups can be archived to a SnapVault destination when Protection Manager is configured for SMSQL.

Best Practices

- Use SnapMirror to replicate the Media Server backup data to a secondary storage system at a DR location for protecting the Media Server from failures at the primary site.
- When keeping the backup data on a LUN, schedule SnapDrive Snapshot copies (using SDCLI) of the Media Server backup data LUN, to be run after any database backup.
- When keeping the backup data on a SMB share, schedule a Data ONTAP Windows PowerShell script to be run after any database backup.
- For disaster recovery purposes, all of the backup data should be automatically replicated to a SnapMirror destination volume. For longer retention, database backups can be archived to a SnapVault destination when Protection Manager is configured for SnapManager for SQL Server.

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

5.9 RETENTIONS

The Data Retention component of SMSP enables you to define data retention and expiration policies. The following list describes the types of data retentions:

- Backup data retention. The backup data retention policies are a method of pruning old backup data to make room for new backups. The Backup Data Retention rule specifies the number of backup sets to keep and the recycle frequency with the number of backup cycles.
- Archived data retention. The archived data retention policies are a method of pruning the old archived data to make room for new archived data. The Archived Data Retention rule specifies the retention time to keep the archived data in the given interval for one Archiver plan.

Archiver should be used in scenarios where data retention policies are well defined as you cannot recover BLOB that are deleted after the policies have expired.

The retention policy also works together with the retention setting in a SnapLock device.

- Orphan Retention. The orphan retention policies, or stub restore policies, or garbage collection, are a method of pruning Archiver data whose stub does not exist in SharePoint after a specified time. Once this option is enabled in Deleted Stub Policy under Storage Optimization, the SharePoint environment is periodically scanned to find deleted content. This process is called *Sync Deletion*, and the identified deleted content is called *orphan data*. The schedule for running the Sync Deletion process is called *Delay Deletion*.

For example, if the delay time is set to 6 months, after the first time a stub is found to be deleted in SharePoint, the stub is marked and archived from SharePoint. If the file is not accessed or the stub is not restored in 6 months, the data is deleted.

Best Practices

- Configure minimum delay time to make sure that contents are not inadvertently deleted by SharePoint users.
- Set the Backup Retention length shorter than the Archiver Orphan Retention length, and set the Archiver Orphan Retention shorter than the Archiver Retention.

5.10 RESTORE GUIDELINES

There are two restore levels available in SnapManager for SharePoint: farm component level and granular level.

FARM COMPONENT LEVEL

The SharePoint backup farm structure is displayed under the Farm Browser tab. Multiple farm components, such as content databases, Web applications, SSP, and even the entire farm and its settings can be selected for restoration.

GRANULAR LEVEL

If granular indexing options were selected during backup, individual site collections, sites, lists, folders, items, and item versions can be restored from the content databases.

RESTORE FROM SNAPVAULT

The SnapVault backup data (or the remote backup data) is used only when the local backup has already been deleted. However, restoring from SnapVault data is much slower than using SnapRestore from the local backup data.

Note: Only database backup to SnapVault is supported. Restoring the SharePoint search index from SnapVault is not supported. Restore from SnapVault is only for the SharePoint database and not for RBS/EBS

OUT OF PLACE RESTORE

SMSP restores backup data to any alternate SharePoint location, either in the same farm or across different farms.

To restore to a location that does not exist in SharePoint, use the blank box at the bottom of each level in the SharePoint tree. To create a new site collection, the full URL is needed, but for other levels (from site down to folder), only the name needs to be specified.

You should always select a container for the content that is either on the same level or one level higher than what is being restored. For example, a site should always be restored either to a site or to a site collection, and a list or library should always be restored either to a site or to another list or library.

To accomplish an Out of Place Restore for content from SnapVault and/or SnapMirror destinations, the network must allow the source to connect to the destination over FC/IP and/or iSCSI.

Best Practices

- NetApp recommends verifying the backups before performing restore operations.
- For verification on the SnapVault and SnapMirror destinations, the network is required to allow the source to connect to the destination over FC/IP and/or iSCSI.
- Use NetApp FlexClone to restore from SnapMirror or SnapVault content databases. Databases from the SnapMirror or SnapVault destination Snapshot copies can be cloned by using FlexClone and mounted on the active SQL Server node to go down to any item level and recover any items or the entire database itself. This is performed on active SQL Servers (supported only in Windows 2008, SMSQL 5.2, and SnapDrive 6.4 for Windows). For clustered SQL Servers in Windows 2008 and later, downtime of the SQL Server instance is not necessary. However, with Windows 2003, downtime would be required.
- Front-end resources cannot be restored together with other farm components. They need to be restored separately after any other farm components are restored.
- Before the restore, make sure that the following Windows services are running: Windows SharePoint Services Timer and Windows SharePoint Services Administration (in SharePoint 2007); or SharePoint 2010 Timer and Windows SharePoint Services Administration (in SharePoint 2010).

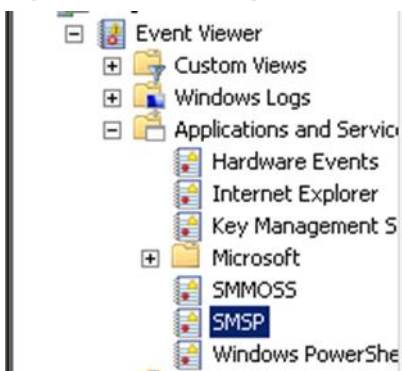
For detailed steps and procedures for the different restore options, see the [SnapManager 6.0 for SharePoint Installation and Administration Guide](#).

5.11 MONITORING

SMSP logs events to the Windows Application Event log for error identification, error description, and confirmation of successful completion of operations. SMSP integrates with NetApp AutoSupport services for sending the event generated reports. SMSP also integrates with Microsoft System Center Operations Manager (SCOM) and Operations Manager (MOM) to send operation details as event logs for the centralized monitoring. Configure these settings under the “Reporting” module in the SMSP Management GUI, through AutoSupport settings, SCOM and MOM settings.

Monitoring of SnapDrive for Windows, SnapManager for SQL Server and SnapManager for SharePoint is possible with the management packs that will be available for SCOM. These will be made available in the NetApp communities with the instructions so that the user can monitor all the events created by Snap Manager products. SnapManager for SharePoint 6.1 events are stored under the separate tree as shown in Figure 12.

Figure 12) SnapManager for SharePoint 6.1 events.



Best Practices

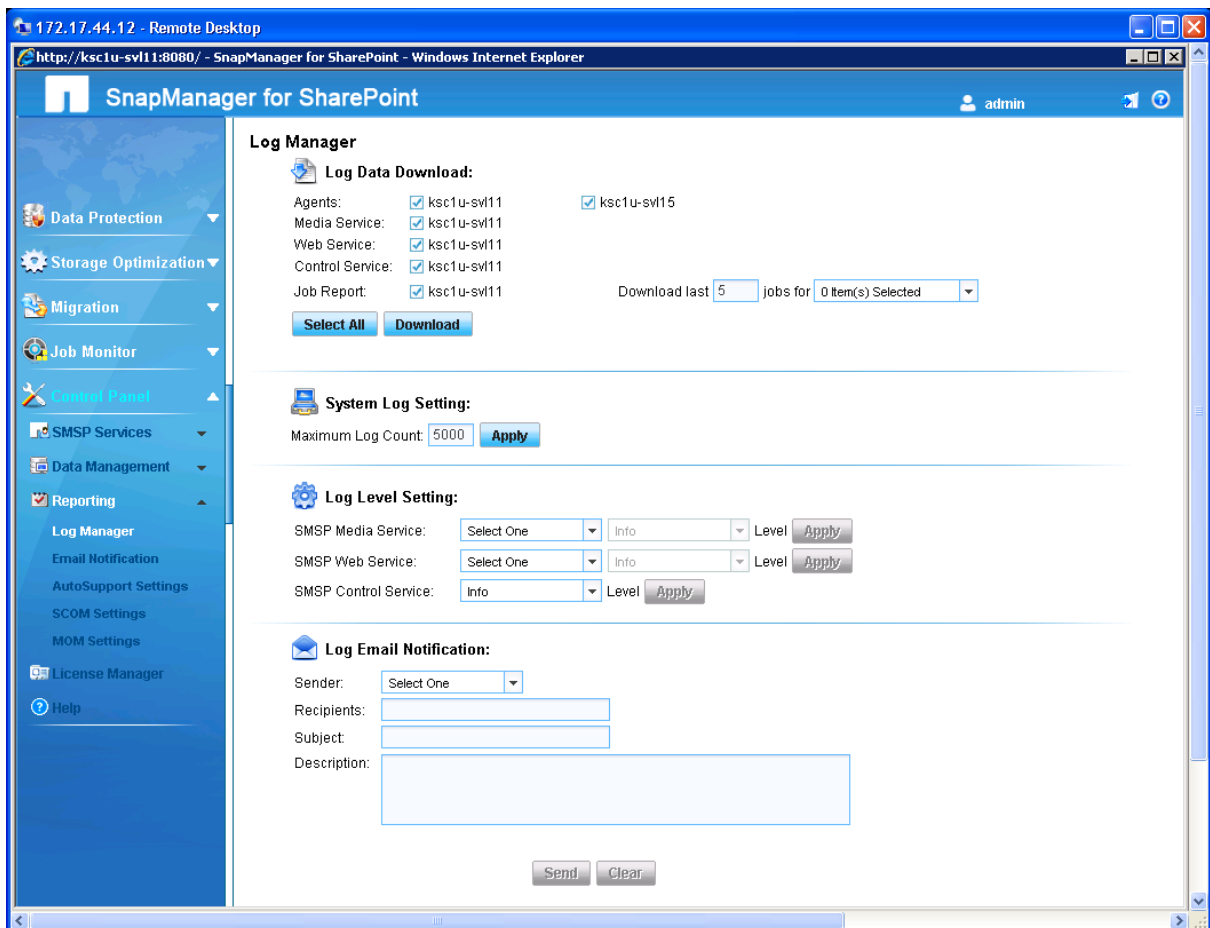
- The AD account specified in Manager Configuration tool will be used for SCOM integration. This account must be the DB Owner for SCOM DB.
- The status of OpsMgr SDK Service on SCOM Server should be “**Started.**”
- The user account entered in the manager configuration tool must be the IIS (MOM Site) pool user or local admin.

5.12 TROUBLESHOOTING

SnapManager for SharePoint Log Manager provides classifies logs associated with SMSP components for troubleshooting and monitoring purposes as follows:

- **System logs.** These logs include all actions performed under SnapManager for SharePoint Manager services. These are stored in the internal Manager databases and can be viewed by clicking Job Monitor > Log Viewer.
- **Manager logs.** Detailed log files of each Manager Service are stored on the Manager machine.
- **Agent logs.** Detailed logs of each Agent are stored on the Agent machine's event viewer called SMSP.

Figure 13) SnapManager for SharePoint Log Manager.



Log Data Download is used to centrally download logs from multiple SnapManager for SharePoint components, including the control, media, Web services, and agents. Select all agents or services for which you require logs and click **Download** and save the consolidated .zip file to your chosen location.

By selecting appropriate values, you can limit the job numbers and job type to be downloaded in the report. By default, all items are selected.

During the course of troubleshooting or calling up Technical support it is recommended to create debug logs.

Debug logging could be enabled by setting the logging under log manager.

Figure 14) Log level setting.

Log Level Setting:

SMSP Media Service:	Select One	Info	Level	Apply
SMSP Web Service:	Select One	Info	Level	Apply
SMSP Control Service:	Debug	Level	Apply	

Here as you can see from this screen there are various levels of logging. To make these changes, choose one of the services from the drop-down list and the appropriate log level setting (either **Debug**, **Info**, **Warning**, **Error**, or **Fatal**). Click **Apply** to save your changes. Because there is only one Control Service, you do not need to select a service to apply the new settings.

Figure 15) Services-log level setting.

Debug Level **Apply**

- Fatal
- Error
- Warning
- Info
- Debug

After the selection is complete you debug level logging for the different service will be set.

6 PERFORMANCE

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

Best Practices

- For SharePoint deployments where the individual content size is small enough (for example, less than 500KB), it's best to keep all of the content in SQL database itself.
- For any larger objects, it's better to deploy EBS or RBS to keep the BLOBs in the external SMB shares on cheaper SATA disk storage.
- Use Flash Cache or PAM module on the controller to optimize the performance.

To maximize performance in SharePoint 2010 deployment with EBS and RBS, follow these best practices.

Best Practices

- Bring up the Media Servers on physical servers instead of virtual machines. If the Media Servers are going to be virtual machines, make sure to have good memory size and CPU power—that is, a heavy host configuration.
- Both the Archived Data Device and the Archived Index Device has to be on a CIFS share.

Note: To learn more about the EBS performance statistics on NetApp storage systems, refer to TR-3874, [Performance Report of SharePoint 2007 EBS Provider for NetApp Storage System](#).

Best Practices

As a best practice always create logical devices with the storage name in lowercase.

<https://kb.netapp.com/support/index?page=content&id=2015280>

To maintain sustained SharePoint deployment, there are performance counters that help to remove bottlenecks and monitor performance. To learn more about the performance counters for SharePoint, see <http://technet.microsoft.com/en-us/library/ff758658.aspx>.

7 HIGH AVAILABILITY

7.1 APPLICATION RESILIENCY

High availability (HA) is an important requirement for business-critical application like SharePoint, where downtime can adversely affect business. SharePoint farm setup may involve load-balanced Web front-end servers, application servers for central administration, search, Excel[®] services, and so on.

Choosing the best solution requires careful planning in terms of deciding whether to create fault-tolerant server hardware or to increase the redundancy of server roles within a farm.

REDUNDANCY WITHIN A FARM

SharePoint Server 2010 supports running server roles on redundant computers (that is, scaling out) within a farm to increase capacity and to provide basic availability.

The capacity required determines both the number of servers and the size of the servers in a farm. After the base capacity requirements are met, you may want to add more servers to increase overall availability.

SQL SERVER FAILOVER CLUSTERING

Failover clustering provides availability support for an instance of SQL Server. A failover cluster is a combination of one or more nodes or servers, and two or more shared disks. A failover cluster instance appears as a single computer, but it provides failover from one node to another if the current node becomes unavailable. SharePoint Server can run on a clustered SQL Server.

SharePoint Server references the cluster as a whole; therefore failover is automatic and seamless from the perspective of SharePoint Server.

SQL SERVER HA MIRRORING

Database mirroring is primarily software solution for increasing database availability. In database mirroring, transactions are sent directly from a principal database and server to a mirror database and server when the transaction log buffer of the principal database is written to disk. This technique can keep the mirror database almost up to date with the principal database. SQL Server Enterprise Edition provides additional functionality that improves database mirroring performance.

High –Available mirroring should be used for mirroring which is also known as high-safety mode with automatic failover. High-availability database mirroring involves three server instances: a principal, a mirror, and a witness. The witness server enables SQL Server to automatically fail over from the principal server to the mirror server. Failover from the principal database to the mirror database typically takes several seconds. SharePoint server is now mirroring aware.

Use of SQL Clustering with NetApp HA Pairs should be used.

Note: Microsoft does not support DB mirroring with RBS. SMSP RBS provider does not support SQL Log shipping. SQL Clustering should be used with NetApp storage controllers to provide HA.

If you are using SnapManager 6.1 for SharePoint to manage your SharePoint farm, then Media Server plays a very important role. It is critical to provide high availability for the Media Server.

Media Servers can be configured into a Microsoft cluster failover configuration by using the LUNs in a shared drive mode. In Cluster Administrator, set all SMSP Manager services as cluster generic services. Set Control Service or Media Server as a dependent on the shared drives.

The current release of SnapManager for SharePoint is cluster aware. The complete configuration of the SnapManager for SharePoint on a cluster is available in the SnapManager for SharePoint 6.1 IAG.

Having discussed all these, let us see the different components that can be made redundant in the SharePoint. Table 6 describes the server roles in a SharePoint Server 2010 environment and their redundancy.

Table 6) Server roles in SharePoint Server 2010 environment.

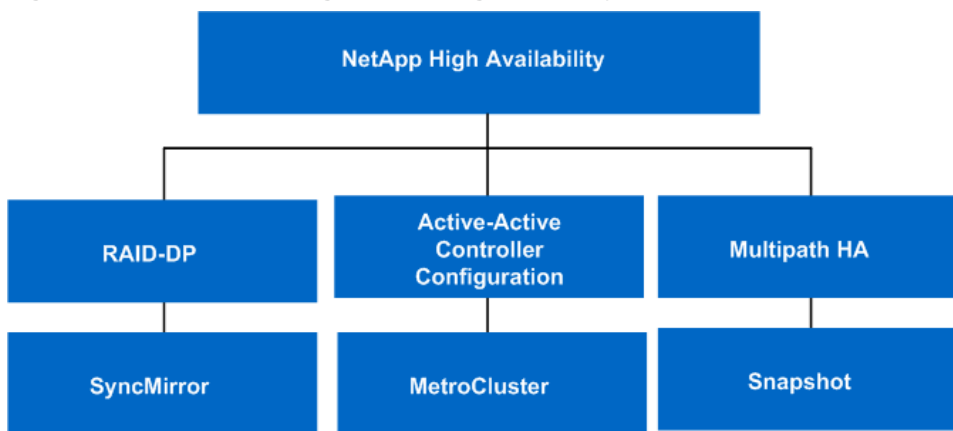
Server Roles and Redundancy Strategy in a Farm Server Role	Redundancy
WFE or front-end server	Multiple WFEs can be deployed and use HW/SW network load balancing
Application server	Multiple application servers can be deployed in a farm
Database server	Clustering or high-availability database mirroring can be deployed for the database server
Media Server	Clustering of Media Servers

7.2 STORAGE RESILIENCY

Storage high availability is a function of data availability and data integrity. Generally, data of higher value requires higher levels, or tiers, of storage resiliency than data of less value. NetApp storage solutions feature a rich set of built-in features and add-on options that offer customers very fine granularity when selecting the storage resiliency tiers that best meet their business needs. Microsoft SharePoint Server 2010 databases typically have varying degrees of business value that in turn dictate the level of storage resiliency required to meet SLAs.

NetApp offers a rich set of features and options that allow customers to meet a wide range of storage resiliency business requirements. Figure 16 shows the NetApp technologies for implementing storage resiliency and high availability that can be used in Microsoft SharePoint 2010 environments for content databases and BLOB storage. The more these technologies are used in conjunction, the better resiliency and high availability can be achieved.

Figure 16) NetApp technologies for storage resiliency.



NETAPP RAID-DP

RAID-DP, the NetApp high-performance implementation of RAID 6, is double-parity RAID that adds a second parity stripe to dramatically increase data availability. With RAID-DP, aggregates and volumes can withstand up to two failed disks in a RAID group, or the more common event of one failed disk followed by an uncorrectable bit read error from the disk drive.

Advantages of Using RAID-DP

- Protects against double disk failures or uncorrectable bit read errors from the disk while in reconstruction mode.
- No impact to required capacity as RAID-DP groups can be double the size of RAID 4 groups.
- RAID-DP enables nondisruptive upgrade of disk firmware that results in zero downtime.

Best Practices

NetApp recommends using RAID-DP always because of its higher reliability with negligible performance costs. NetApp also recommends using the default RAID group sizes with RAID-DP. For even higher resiliency and data reliability, consider using smaller RAID-DP group sizes.

NETAPP SYNCMIRROR

NetApp SyncMirror maintains two copies of the data online so that the data is available and up to date at all times, even in the event of hardware outages, including a very unlikely triple disk failure. When used

with the NetApp active-active controller configuration high-availability solution, NetApp SyncMirror enables you to achieve the highest level of data availability, previously attainable only with mainframe architectures.

Benefits of Using SyncMirror

- Synchronous replication makes sure that a copy of the data is always up to date and available in the local data center.
- Transparent operation enables all applications to access the synchronously replicated data with no programming or system changes.
- Integration with active-active controller configuration offers complete hardware redundancy with automatic failover for the highest level of data availability.

ACTIVE-ACTIVE CONTROLLER CONFIGURATION

Active-active controller configuration (previously known as clustered failover), eliminate the storage controller as a single point of failure. In this configuration, each storage controller has its own dedicated pool of disk drives and handles all I/O operations during normal operation. Each clustered storage controller pair is connected to its partner's disk drives as well and maintains a heartbeat status check of its partner. If a heartbeat check reveals that a paired partner is down, the remaining controller initiates a takeover operation of the failed storage controller's disks and handles all I/O requests until the down controller can be brought back online.

Benefits of Using Active-Active Controller Configuration

- Active-active configurations prevent a storage controller from becoming a single point of failure.
- Active-active configurations increase storage resiliency levels.

Best Practice

Use active-active controller configurations for the highest data availability in SharePoint database environments.

MULTIPATH STORAGE HIGH AVAILABILITY

Multipath storage for an active-active controller configuration provides redundancy for the path from every controller to every disk shelf in the configuration. An active-active configuration without multipath storage has only one path from every controller to every disk, but an active-active configuration with multipath storage has two paths from every controller to every disk, regardless of which node owns the disk.

Advantages of Multipath Storage for High Availability

By providing two paths from every controller to every disk array in the configuration, multipath storage offers the following advantages:

- The loss of a disk shelf module, connection, or host bus adapter (HBA) does not require a failover. The same storage system can continue to access the data by using the redundant path.
- The loss of a single disk shelf module, connection, or HBA does not prevent a successful failover. The takeover node can access its partner's disks by using the redundant path.
- You can replace modules without having to initiate a failover.

Best Practice

If a disk failure occurs, loop traffic generated by the reconstruction operation competes with existing production SharePoint 2010 workload and can result in slower performance of the SharePoint databases. One way to minimize the impact of disk reconstruction on loop performance is to employ more loops than would normally be required to support the amount of storage capacity. With RAID group disks spread across additional loops, any reconstruction activity is also spread across the extra loops, resulting in faster reconstruction times and shorter performance impact on SharePoint 2010 workloads (assuming that the storage processor CPU/memory does not initially bottleneck the reconstructions).

A second way to limit the performance impact of reconstructions is to change the Data ONTAP option setting `raid.reconstruc.perf_impact` to `low`. This approach reduces performance impact on SharePoint 2010 databases, but the reconstruction process takes longer.

NETAPP SNAPSHOT

- NetApp Snapshot copies are backups of how a volume looks at a particular point in time Snapshot copies work almost instantaneously, with very little impact on storage capacity. This approach allows backups to be created in a matter of seconds without affecting the production environment. Backups of the environment can be taken more frequently, and more of these highly space-efficient backup copies can be stored.

10 VIRTUALIZATION

Customers of all sizes are virtualizing and performing server consolidation across their application infrastructure to lower cost, improve scalability and improve service level agreements. SharePoint 2010 as an application supports virtualization so we can similarly virtualize the SnapManager for SharePoint 6.0 components.

During the planning of virtualization, we would have to evaluate and decide between the virtualization technology and the differentiating factors between multiple vendor, namely, the Microsoft Windows Server 2008 R2 or Hyper-V Server 2008 R2 or VMWARE ESX Virtualization stack.

Here are some of the best practices with virtual deployment of SharePoint:

- **Use hardware – assisted virtualization** - Hardware Assisted Virtualization helps improve the performance of virtualization software and the application response times. It has been observed that HAV offers 5% or more throughputs. Hardware-assisted virtualization is provided by the Intel Virtualization Technology (Intel VT) processor family and the AMD Virtualization (AMD-V) processor family
- **Enable hyper-threading on processors that supports this technology** – Hyper-threading is a technology that is available with processors that support symmetric multithreading (SMT). This technology provides two threads for each processor core. The net effect is a doubling of logical processors for each core.
- **Configure Non-uniform memory Access correctly** – Non-uniform memory access (NUMA) is a computer memory design that is used in multiprocessors to improve performance by reducing memory access latency and increasing memory bandwidth. The NUMA boundary is calculated by dividing the memory on the host server by the number of processor sockets. Performance is reduced by approximately 8% when the virtual machine memory allocation is larger than the NUMA boundary.
- **Configure Hyper-V host for optimal performance** – These are some of the points that will help to get the optimum performance from the host server in virtualized environment.
 - a. Dedicate the host computer to Hyper V. Do not run additional applications on the host.
 - b. Install and configure only the required roles and services.

- c. Do not store VM data on system drives. The root partition should only contain the OS and Hyper-V.
- d. Use more than one network adapter on the physical server and dedicate one adapter to Hyper-V Server administration. Do not bind any virtual networks to this adapter. If virtual machines are sharing a network adapter, monitor latency, and throughput to make sure that the adapter is not oversubscribed.
- **Maintain version compatibility for integration services** –Make sure that the Host OS and the guest OS system run the same version of the integration services. Here is the link to the version compatibility for Integration Services. [http://technet.microsoft.com/en-us/library/ee207413\(W.S.10\).aspx](http://technet.microsoft.com/en-us/library/ee207413(W.S.10).aspx).
- **Determine storage options that are based on virtual machine roles and workloads** – Hyper V has three options for virtual hard disks (fixed size, dynamically expanding, and differencing), we recommend fixed-size disks. Fixed-size disks provide the best overall performance and reduce the risk of exceeding physical storage space.
- **Configure general virtual machine settings for performance and stability** – There are few points that help us get the best performance and stability from the Virtualized environment. Prevent boot swarm by using time delay if VM are configured for automatic start. Shut down the VM exactly the same way as the physical computer does not use save state before shutdown. SharePoint Server 2010 features use timer jobs extensively and any latency in time synchronization affects the SharePoint reliability. Do not use time synchronization with the host.
- **Do not use virtual machine Snapshots in a production environment** - Virtual machine snapshots are file based snapshots of the disk data, state and configuration of the VM at a specific point of time. It's recommended not to use VM Snapshots in the production environment. Microsoft SharePoint Server uses timer job extensively, taking the VM snapshots affects time sensitive operations and can result in data corruption or loss. NetApp offers application like SnapManager for SharePoint and SnapManager for Hyper V to remove the limitations by taking block level Snapshots®.
- **Design the virtual topology for optimal performance** –To gain maximum throughput in the SharePoint farm. Design the topology across hosts like mixing farm server role on each host and monitor the performance of the WFE on a single host. For example, mixing WFE with application server as usually they do not write to disk at the same time, this will reduce the disk contention.
- **Do not overload the Hyper-V host** –Make sure that the processor and memory are not oversubscribed. This will have negative effect on performance. Plan the virtual machine configuration and deployment based on the overhead, HA and scale up or scale out
- **Do not run resource intensive jobs on the Hyper-V host and virtual machines at the same time**- Continuous and careful monitoring of the Hyper-V host and virtual machines ensures that resource intensive jobs are not run at the same time. For example, if execution of a backup program on the physical computer and the virtual machines at the same time, there will be resource contention.

Best Practice

For functional areas with more than one connection, such as the multiple network adapters used for VM communications, the connections should be spread across different network adapters, especially if multiple port network adapters are installed. This allows those functional areas to maintain connectivity to the network when configured properly so that, in the event of a port or adapter failure within the Hyper-V server, connectivity is not lost.

Since NIC teaming is not supported for iSCSI communications, NetApp recommends configuring multiple paths for storage connectivity to provide redundancy and additional bandwidth in some configurations using multipathing/MPIO.

For more information on the preceding section follow the links:

- **Best Practices for Virtualization**
<http://technet.microsoft.com/en-us/library/hh295699.aspx>

Perform the Hardening of the Hyper-V role and best practice to reduce the attack surface of Hyper-V and the recommendations to configure the secure virtual networks. Manage and delegate the virtual machine management securely within the organization is essential too. For more information please follow the link

- **Hyper-V Security Guide.**

<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=16650>

To learn more about the Hyper-V implementation on NetApp Storage and the Best practice for Hyper-V and SMHV. Read through these links.

- **NetApp and Microsoft Virtualization: Solution and Implementation Guide**

<http://media.netapp.com/documents/tr-3701.pdf>

- **NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V**

<http://media.netapp.com/documents/tr-3702.pdf>

- **NetApp and VMware vSphere Storage Best Practices**

<http://media.netapp.com/documents/tr-3749.pdf>

Table 7 describes the SharePoint Server role and their virtualization decision that is made in the deployment stage.

Table 7) SharePoint server role and their virtualization.

Role	Virtualization Decision
Web Role	Ideal
Query Role	Ideal
Application Role	Ideal
Index Role	Evaluate
Database Role	Evaluate

Most of the virtualized SharePoint environments have started with the WFE (Web front end). This is the role which renders content. The WFE is the most commonly virtualized role in a SharePoint farm because of its smaller memory and disk requirements. The consideration for WFE for virtual environment is quite simple we can provision additional server for load balancing and fault tolerance. Application and patch levels the same on all the servers that are provisioned.

Then we have the query server role which would process search queries. This server requires the propagated copy of the local index. The propagated copy of the index could be large based on the type of content that is getting indexed. If we have query role with WFE we will build a high available solution into the search service.

Application servers which would host the services like excel form services or any other application service role. Provision more servers (VM) as resources requirements for individual application increase.

Index server role needs consideration for virtualization as this server is performing the crawl on the content repository and it is resource intensive and needs more memory and processing than the web server role.

Database server role could be used in the virtual environment with lower resource usage requirements. However, from my personal experience, I have seen that being the database server there are a lot of read/write that happen on this server and putting it on the virtual environment will be taxing on the physical host. So, it is better to keep it on the physical machine.

With the release of SharePoint Server 2010 SP1 There has to be additional planning that has to be considered to effectively manage these databases that could be found in this link.

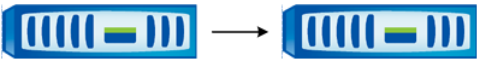
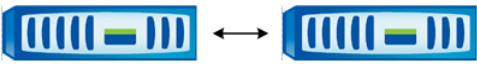
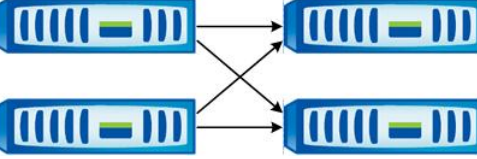
- **SharePoint Server 2010 capacity management: Software boundaries and limits**
<http://technet.microsoft.com/en-us/library/cc262787.aspx>

11 DISASTER RECOVERY

There are numerous ways to augment data availability in the event of hardware, software, or even site failures. Mirroring offers data availability and minimizes downtime. NetApp SnapMirror technology performs block-level mirroring of the data volumes to the destination asynchronously. 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. SnapMirror is a key component in implementing enterprise data protection strategies. If a disaster occurs at a source site, mission-critical data can be accessed from a mirror on the NetApp storage deployed at a remote facility for uninterrupted data availability.

Table 8 summarizes the replication models that can be used for Microsoft SharePoint Server 2010.

Table 8) Replication models.

Configuration	Replication Model
	Unidirectional replication
	Bidirectional replication
	Multidirectional replication

SETTING UP SNAPMIRROR RELATIONSHIPS

SnapMirror requires a source volume on the primary site and a destination volume on the remote or recovery site for mirroring. The destination volume should be greater than or equal to the size of the source volume.

On the source storage controller, use the `options.snapmirror.access` command to specify the hostnames of the destination storage system that are allowed to copy data from the source storage system.

Example: `options snapmirror.access host=<destination storage system>`

The destination volumes should be restricted to allow SnapMirror to access them. Use the `vol restrict` command to perform this action.

Example: `vol restrict <vol_name>`

Note: Use the `snapmirror status` command to monitor the status of the SnapMirror operations on the specified volumes.

SnapMirror.conf

The `snapmirror.conf` file is a core configuration of all SnapMirror operations. The `/etc/snapmirror.conf` file defines the relationship between the source and the destination, the schedule intervals used by the destination to copy data, and the arguments that control SnapMirror when copying data. The file resides on the destination NetApp storage systems.

Note: There is a limit of 600 entries in the `/etc/snapmirror.conf` file for each individual storage system. In a storage cluster, this limit applies to the cluster pair. Older versions of Data ONTAP (earlier than 6.4.3) had a limit of 210 entries per storage system.

Use the `/etc/snapmirror.conf` file to specify the relationship and the schedule interval that determine the type of replication method.

SYNCHRONOUS SNAPMIRROR

Synchronous SnapMirror provides block-level replication at the storage system by synchronously replicating each transaction to the destination as it is written at the source. Before achieving synchronous replication mode, SnapMirror starts by creating a baseline Snapshot copy and replicating the copy to the destination. When the baseline transfer of the data from the source volume is complete, the asynchronous SnapMirror transitions into synchronous mode. Synchronous SnapMirror then uses consistency point forwarding. This method makes sure that writes of data from memory to disk storage are transferred from the primary system to the secondary system. For more information about consistency point, refer to [Data ONTAP Administration Guide](#).

To configure synchronous SnapMirror between two volumes:

Add this entry to the `/etc/snapmirror.conf` file:

```
<src_storage>:<src_vol> <dest_storage>:<dest_vol> - sync
```

This entry specifies that `<Source_Vol>` on `<Source_Storage>` should be synchronously replicated to `<Dest_Vol>` on `<Dest_Storage>`.

ASYNCHRONOUS SNAPMIRROR

SnapMirror replicates Snapshot copies from a source volume to a partner destination volume, thus replicating source object data to destination objects at regular scheduled intervals. SnapMirror source volumes are writable data objects whose data is to be replicated. The source volumes are the objects that are normally visible, accessible, and writable by the storage system's clients, and the destination volumes are read-only objects.

- **To configure scheduled asynchronous SnapMirror updates between volumes:**

Edit the `/etc/snapmirror.conf` file:

```
<src_storage>:<vol_name> <dest_storage>:<vol_name> - <Minute> <Hour> <week of the month> <day of the week>
```

Example: `<src_storage>:<vol_name> <dest_storage>:<vol_name> - 0-59/30 * * *`

This example specifies that `vol_name` on `src_storage` should be replicated to `vol_name` on `dest` storage every 30 minutes every week of the month and every day of the week.

Best Practices

- To synchronously mirror volumes by using SnapMirror the flexible volume must at least 10GB in size.
- When using SnapManager for SharePoint and using SnapMirror to replicate the SMSP Snapshot copies, no schedule should be specified in the `/etc/snapmirror.conf` file, and the relationship entry should be specified as in the following example:

Example: `<src_storage>:<vol_name> <dest_storage>:<vol_name> - - - - -`

SEMI-SYNCHRONOUS SNAPMIRROR

Semi-synchronous mode offers a middle ground that keeps the source and the destination file systems more closely synchronized than asynchronous mode, but with less application performance impact than synchronous mode. Configuration of semi-synchronous mode is identical to the configuration of synchronous mode. The key difference between synchronous and semi-synchronous modes is that NVLOG forwarding is turned off in semi-synchronous mode. This means that client writes are acknowledged soon after the primary system writes the data to its NVRAM. Since NVLOG forwarding is turned off, the data is now protected at the secondary site only through consistency point (CP) forwarding. The maximum window of time between any two CPs is 10 seconds, so a disaster at the primary site could result in a data loss of approximately 10 seconds. For more information about NVLOG forwarding, see the [Data ONTAP Administration Guide](#).

- **To configure semi-synchronous SnapMirror between two volumes:**

Edit the `/etc/snapmirror.conf` file:

```
<src_storage>:<vol_name> <dest_storage>:<vol_name> - semi-sync
```

Note: NetApp SnapManager for SharePoint integrates and supports both synchronous and asynchronous SnapMirror and does not support semi-synchronous SnapMirror.

SNAPDRIVE FOR WINDOWS AND SNAPMANAGER FOR SHAREPOINT WITH SNAPMIRROR

SnapDrive for Windows and SnapManager for SharePoint are integrated with SnapMirror to trigger a SnapMirror update immediately after a backup job finishes. SnapDrive for Windows is used for managing (creating, deleting, renaming) Snapshot copies on the source volume of SnapMirror. Any changes to Snapshot copies on the source system are immediately updated on the destination system.

SnapManager for SharePoint tightly integrates with SnapDrive for Windows, enabling the creation of application-aware consistent Snapshot copies and replicating the Snapshot copies to the destination site. Both the database volumes and the BLOB storage volumes on the production storage system can be mirrored to a disaster recovery site by using SnapMirror, for immediate availability of the data on the DR site. Any time an SMSP backup job is executed for database LUNs, you can select the Update SnapMirror after Operation option from the Advanced Settings tab in the Backup Plan Builder. For the BLOB storage SMB share, the SnapMirror update is triggered automatically triggered when the backup job finishes.

Note: Both database volumes and BLOB store volumes should be configured in SnapMirror relationship with the destination volumes, before executing any backups.

SnapManager system data consists of all the settings, including plans, policies, device configurations, users, and so on. System Backup is an option in SnapManager for SharePoint in the System Recovery module, which allows the backup of system settings used in that SnapManager for SharePoint deployment. System Backup saves the system settings into a flat file to the location defined under Backup Destination. The default path is `C:\Program Files\NetApp\SnapManager for SharePoint\UserData`.

Best Practices

- Use a non-default location. For disaster recovery purposes, it is best to change the default location to a NetApp LUN with SnapMirror enabled, so that system backup data is automatically replicated to the SnapMirror destination when a system backup job is completed. Otherwise, the backup data must be manually copied over to the DR site.
- As described in "SharePoint Components Backup" section 5.8, "Backup Guidelines," the Media Server backup data should also be protected by moving from its default location, `C:\Program Files\NetApp\SnapManager for`

SharePoint\VaultServer\Media\data-backup, to a NetApp LUN or SMB share. For disaster recovery purposes, this NetApp LUN or SMB share should be SnapMirror enabled, so that it can be replicated to a SnapMirror destination volume automatically.

- During recovery make sure that the LUN mappings are correct on all SnapMirror volumes.

DISASTER RECOVERY GUIDELINES FOR SMSP

- The SharePoint database and search index backups must be available. Typically, these are automatically replicated through SnapMirror, when enabled for it.
- SnapManager for SharePoint backup jobs data and system backup data must be available. If Media Service and system backup locations use SnapMirror enabled LUNs, they should be automatically replicated.
- At the DR site, all server topology should be identical to the production site, which includes both the SharePoint farm topology and the SnapManager for SharePoint topology. There are several ways to achieve this. A separate AD domain for the DR site can be used to keep the same server topology; or, if the DR site does not need to coexist with the production site, disk imaging or virtualization technology can be used to make sure that the topology is the same.
- If a separate AD domain is used at DR site, that domain should have a trusted relationship with the primary site's AD so that users can still access SharePoint content. Disaster recovery can also be performed if the DR site is in the same AD domain as the original location.

For the detailed steps involved in the disaster recovery process, refer to the [SnapManager 6.0 for SharePoint Installation and Administration Guide](#).

11.1 DISASTER RECOVERY TESTING

Because an actual DR test involves downtime for production environments, many customers choose not to perform frequent DR testing even though a DR plan exists. When FlexClone is used with SnapMirror DR volumes, the remote site can be used for DR testing without interruption to production operations and DR replication. Applications can be brought up at the DR site for data consistency. SnapDrive for Windows can be used to create FlexClone volumes from the replicated Snapshot copies and to connect the LUNs to the test nodes at the DR site. SnapManager can be used to test recovery of the Snapshot copies from the cloned volume to simulate and test DR; the clones can be destroyed after the DR testing. Table 9 describes the LUN status on SnapMirror destination systems during different replication statuses.

Table 9) LUN status on SnapMirror destination systems.

Replication State	DR LUN Status at the SnapMirror Destination
Mirror being synchronized	Online, read-only Note: FlexClone volumes can be used to create a writable Snapshot copy of the read-only LUN.
Mirror is broken or split	Online, read-write, unmapped

Best Practice

Use SnapDrive to create a FlexClone volume at the SnapMirror destination. This automates the creation of a clone and connects the LUN within the clone to the destination host.

SPACE GUARANTEES

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

- **Thin Provisioning Aggregates on the Source System**

To thin provision an aggregate on the source, create flexible volumes with a guarantee of None or File so that the volume size is not limited by the aggregate size. The total size of the flexible volumes can be larger than that of the containing aggregate.

- **Thin Provisioning Aggregates on the Destination System**

Use caution when thin provisioning an aggregate at the destination system, because SnapMirror fails when the volume runs out of space.

Starting with Data ONTAP 7.3, it is possible to set guarantees on the SnapMirror destination volume so that the SnapMirror updates never fail on that volume. The default behavior is that the volume guarantees are turned off.

The following example demonstrates space usage with and without volume guarantees on the SnapMirror destination volume:

Example: For a 1TB SnapMirror source volume that is 75% full, the SnapMirror destination volume (or replica) needs 750GB with the guarantee disabled and the full 1TB with the guarantee enabled.

PERFORMANCE IMPACT OF SNAPMIRROR ON SHAREPOINT SERVER

Performance in general is a difficult area to quantify. This section discusses the effects of SnapMirror on individual storage systems and its effect on overall performance. Any synchronous replication method, regardless of the technology used, has an impact on the applications that use the storage. Therefore, it is important to understand the business requirements for application performance and data protection, to enable informed choices between various data protection strategies.

CPU Impact

When a storage system that is running SnapMirror in synchronous or asynchronous mode receives a write request from a client, the storage system must do all operations that are normally required and do additional processing related to SnapMirror to transfer information to the destination storage. This adds significant CPU impact to every write operation. Also, any read or write activity performed by the clients on the storage system over the network typically results in CPU utilization.

When replicating data by using synchronous or asynchronous mode, all the data written to the primary storage system by clients must be transferred to the destination storage system across the network. So in addition to processing the data from the clients, the storage controller CPU must also send the data to the destination storage. This can nearly double the CPU utilization on storage systems with synchronous or asynchronous SnapMirror as compared to the same workload on a storage system without SnapMirror.

NETWORK BANDWIDTH CONSIDERATIONS

All data processed by the source storage system must be replicated to the storage system as it is written. When using SnapMirror in synchronous or asynchronous mode, write throughput on the source storage system cannot exceed the network bandwidth available between the source and the destination storage systems. It is important to consider the network throughput requirements before sizing network bandwidth requirements between the source and the destination storage systems. Roughly calculated, 250GB of data could be transferred in 7.5 hours over a 10Base-T full duplex network, and it might take 1.5 hours on a Gigabit Ethernet link.

OTHER SNAPMIRROR FACTORS THAT ADD TO LATENCY ON SHAREPOINT SERVER

Other sources of latency over the network are networking devices through which the traffic must pass, such as routers, switches, and so on. Each device adds latency as it receives the signal on one interface, processes the signal, and then transmits it through another interface. The amount of latency can be considered small, but it can add up if there are many devices in the network.

Given the different modes of replication available in SnapMirror and the performance characteristics, it is important to select the correct replication mode for your environment.

Table 10) Replication modes.

Replication Mode	Achievable RPO	Performance Impact
Synchronous SnapMirror	Zero or near zero	Medium to high
Semi-synchronous SnapMirror	Near-zero to 1 minute	Low to medium
Asynchronous SnapMirror	Minutes to hours	Low

11.2 NETAPP METROCLUSTER

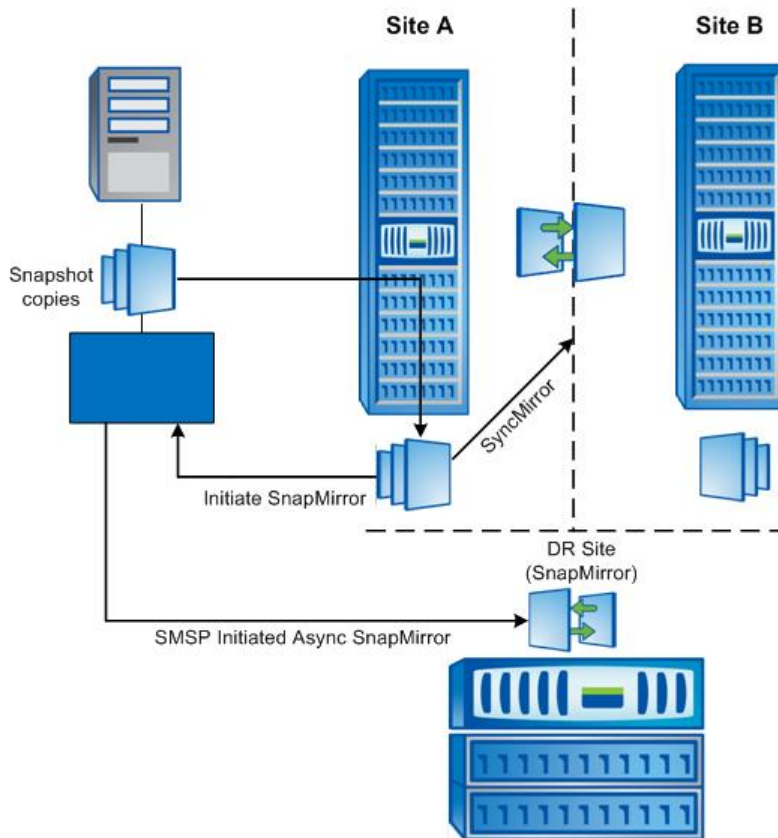
Achieving continuous availability for mission-critical applications like SharePoint is a challenge. The NetApp MetroCluster solution combines array-based clustering with synchronous mirroring to deliver continuous availability and zero data loss with automated failover, from individual component failure to a complete site failover. MetroCluster also reduces planned and unplanned downtime and eliminates repetitive change management activities to reduce the risk of human error and administrative overhead. MetroCluster can be deployed in two modes:

- Stretch MetroCluster
- Fabric MetroCluster

Stretch MetroCluster is a nonswitched, simple active-active storage cluster configuration that can extend up to 500 meters, depending on the speed and the cable type. Fabric MetroCluster uses four Fibre Channel switches in a dual fabric configuration and a separate cluster interconnect card. It can extend up to 100 kilometers between primary and secondary locations. Both stretch and fabric MetroCluster provide array-based synchronous mirroring between cluster nodes by combining disk-mirroring RAID 1 protection with NetApp RAID-DP technology to provide greater resiliency in case of component failure, as well as automated site failover with integration with the ApplianceWatch™ PRO 2.1 tiebreaker solution for MetroCluster.

MetroCluster can be deployed in a single site or a multisite configuration for SharePoint 2010 environments, and asynchronous SnapMirror can be used along with SnapDrive for Windows and SnapManager 6.0 for SharePoint to extend the DR capabilities to a remote site beyond 100 kilometers. Figure 17 shows how MetroCluster and SnapMirror can be deployed with SnapDrive for Windows and SnapManager for SharePoint.

Figure 17) MetroCluster and SnapMirror deployment.



For more information about NetApp MetroCluster design, see the [MetroCluster Implementation and Design Guide](#). For more information about SharePoint behavior in various failure scenarios in a MetroCluster environment, see the [NetApp Solution Guide: Microsoft Mixed Workload on NetApp MetroCluster](#).

12 REFERENCES

- Database types and descriptions (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/cc678868.aspx>
- SharePoint Server 2010 performance and capacity test results and recommendations
www.microsoft.com/downloads/en/details.aspx?FamilyID=fd1eac86-ad47-4865-9378-80040d08ac55&displaylang=en
- Fundamentals of performance
<http://technet.microsoft.com/en-us/library/ff758647.aspx#fundamentals>
- Migrate content into or out of Remote BLOB Storage (RBS) (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/ff628254.aspx>
- Manage Remote BLOB Storage (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/ee748638.aspx>
- Plan for availability (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/cc748824.aspx>
- Availability configuration (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/dd207313.aspx>

- TR 3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations
www.media.netapp.com/documents/tr-3326.pdf
- TR 3874: Performance Report of SharePoint 2007 EBS Provider for NetApp Storage System
www.media.netapp.com/documents/tr-3874.pdf
- TR 3548: Best Practices for MetroCluster Design and Implementation
www.media.netapp.com/documents/tr-3548.pdf
- TR 3804: Microsoft Exchange Server, SQL Server, SharePoint Server Mixed Workload on Microsoft Hyper-V and NetApp Fabric MetroCluster
www.media.netapp.com/documents/tr-3804.pdf
- SharePoint Server 2010 capacity management: Software boundaries and limits
<http://technet.microsoft.com/en-us/library/cc262787.aspx>
- Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)
<http://technet.microsoft.com/en-us/library/cc298801.aspx>

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

Go further, faster®

© 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, ApplianceWatch, ASUP, AutoSupport, Data ONTAP, FlexClone, FlexVol, MetroCluster, NOW, RAID-DP, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapRestore, Snapshot, SnapVault, SyncMirror, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Safari is a registered trademark of Apple Inc. Intel is a registered trademark of Intel Corporation. Microsoft, Active Directory, Excel, Internet Explorer, OneNote, SharePoint, SQL Server, Windows, and Windows Server are registered trademarks and Windows PowerShell is a trademark of Microsoft Corporation. Java is a registered trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3887-0212

