Technical Report

NetApp SnapManager 6.1 for Microsoft SharePoint on Data ONTAP 8.1 Operating in Cluster-Mode Best Practices Guide

Amarnath Rampratap, Rob Barker, NetApp
April 2012 | TR-4028

Abstract

This document discusses planning considerations and best practices for deploying Microsoft® SharePoint® Server 2010 on NetApp® storage systems. It also covers best practices for the NetApp enterprise content management solution for SharePoint—NetApp SnapManager® 6.1 for SharePoint.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Scope

This report documents best practices for SnapManager for SharePoint (SMSP) with NetApp Data ONTAP® 8.1 software operating in Cluster-Mode. Cluster-Mode is one of the main features of the latest release of SnapManager 6.1 for SharePoint. All features from the previous releases of SMSP are part of this release as well. This release also continues to support management of Data ONTAP operating in 7-Mode storage systems. For the best practices for SnapManager for SharePoint using 7-Mode storage systems, refer to TR-3887: SnapManager 6.1 for SharePoint with Data ONTAP 8.1 Operating in 7-Mode.

2 Introduction

This document discusses deployment of NetApp SnapManager for SharePoint on Data ONTAP 8.1 operating in Cluster-Mode. Data ONTAP 8.1 operating in Cluster-Mode allows nondisruptive operations during storage infrastructure maintenance and upgrades using volume move and intercluster asynchronous volume replication (replication between volumes hosted on different clusters).

Figure 1 illustrates the architecture of Data ONTAP 8.1 operating in Cluster-Mode.

Figure 1) Architecture of Data ONTAP 8.1 operating in Cluster-Mode.

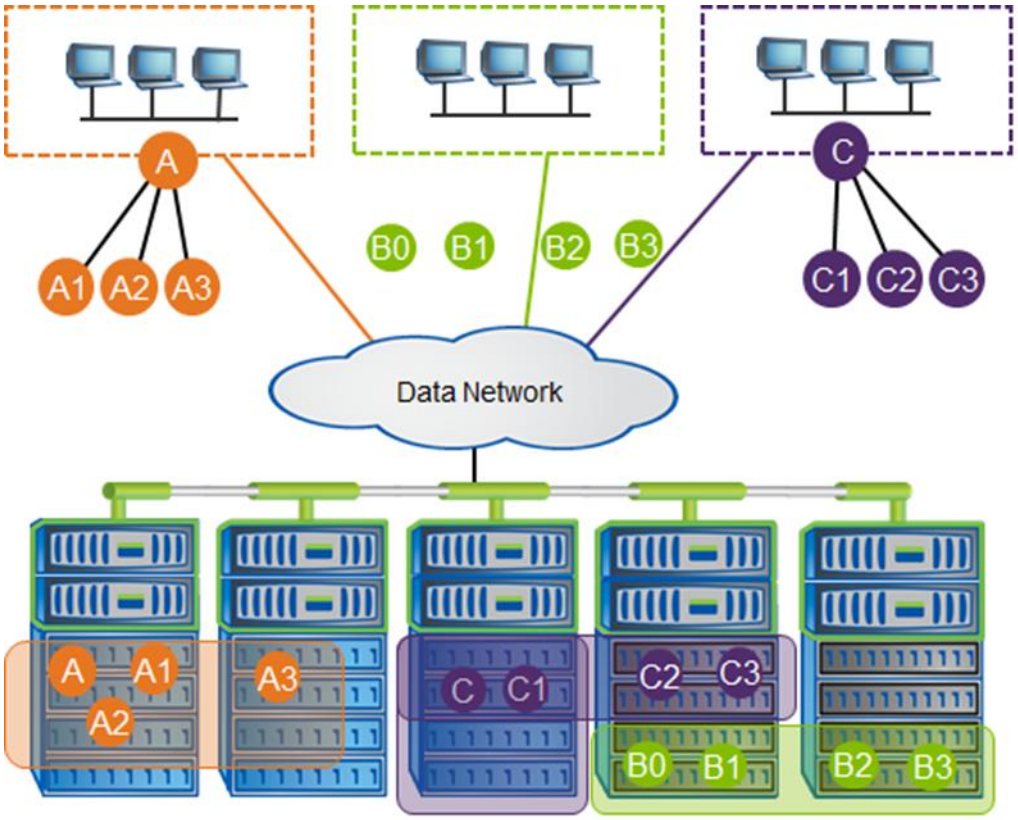

With the release of Data ONTAP 8.1 operating in Cluster-Mode, the solution provides enterprise-ready, unified scale-out storage. This solution is the basis for large virtualized shared storage infrastructures and, most importantly, it is built on the solid foundation of Data ONTAP.

The following sections discuss Data ONTAP 8.1 operating in Cluster-Mode and the benefits that Cluster-Mode provides to SMSP and SharePoint Server.

3 Benefits of Cluster-Mode

The major benefits of Data ONTAP operating in Cluster-Mode include:

- Scalability
- High availability
- Flexibility
- Multi-tenancy

3.1 Scalability

All storage controllers have physical limits to their expandability. Expandability is limited by the number of CPUs, memory slots, and space for disk shelves that define the maximum capacity and performance of the controller. If more storage or performance capacity is required, you might be able to add CPUs and memory or install additional disk shelves. Ultimately, however, the controller has no additional space for hardware. In this case, the only option to increase storage or performance capacity is to acquire another controller. One way to do this is to scale up.

Each additional controller is a completely independent management entity that doesn't provide any shared storage resources. If the original controller must be replaced by a newer and larger controller, data migration is required to transfer data from the old controller to the new one. This is time consuming and disruptive and likely requires configuration changes on all the attached host systems. If the new controller coexists with the old controller, there are now two storage controllers that must be managed individually, but there are no built-in tools to balance or reassign workloads across them.

The situation worsens as the number of controllers increases. With scaling up, the operational burden increases as the environment grows. The result is an unbalanced environment that is difficult to manage.

Technology refresh cycles require substantial planning in advance, lengthy outages, and configuration changes that can introduce risk into the system. By contrast, Data ONTAP operating in Cluster-Mode provides a scale-out strategy, which means that as the storage environment grows, additional controllers are seamlessly added to the resource pool that resides on a shared storage infrastructure. Host and client connections, as well as datastores, can move nondisruptively anywhere in the resource pool, and existing workloads can be easily balanced over the available resources. New workloads can be easily deployed. Technology refreshes, such as replacing disk shelves and adding or completely replacing storage controllers, are accomplished while the environment remains online and serving data.

Data ONTAP operating in Cluster-Mode helps with the rapid and seamless deployment of new storage. It is built for continuous scale-out operations with no downtime, and it scales out transparently to a compute farm with an unchanged namespace.

3.2 High Availability

Data ONTAP 8.1 operating in Cluster-Mode is architected with the high-availability (HA) requirements of today's businesses in mind. Recovery capability is provided by a pair of nodes, or storage systems, called an HA pair. The HA pair is redundantly configured to serve data in case one of the nodes fails. Operations such as volume movement, node failover, and switch failure can be performed without disrupting the storage or applications.

Meet the needs of SharePoint Server data growth and the increased and changing SharePoint Server application workload by adding more controllers or storage without disrupting current applications. In virtualized SharePoint environments especially, add Web front ends (WFEs) or media servers to scale out the environment. This increases the SharePoint Server application uptime and failover protection during hardware infrastructure maintenance and software upgrades using a highly available storage back end.

Storage failover protection is a core attribute of Cluster-Mode. HA pairs of controllers are the building blocks that form the storage cluster. This architecture enables transparent controller clustering and

failover capability in which a failed storage controller causes its partner node to take over its disk arrays, volumes, and running services to provide continuous operation.

Cluster-Mode can scale up linearly as the number of nodes increases. To attain high availability, the layers in the setup also have the following redundant features:

- Linear throughput scaling to multi-GB/sec
- Linear scale performance for single volume with striping
- Linear scale read performance with load-sharing mirrors and collective read/write performance in a single namespace

Table 1) Cluster-Mode components and features.

Component	Feature
Sessions from the host to the storage subsystem	Data ONTAP DSM 3.5
Host Ethernet port	NIC teaming (supported systems only)
Server	Windows® failover cluster manager
Ethernet	Two Ethernet switches
Fibre Channel (FC)	Two fabric switches
Storage systems	Two-node Data ONTAP Cluster-Mode system (minimum)
Logical interfaces (LIFs) in the storage systems	Multiple ports mapped to the same LIF

3.3 Flexibility

Data ONTAP 8.1 operating in Cluster-Mode provides ample flexibility for storage operations by accommodating any mix of FC and SATA drives. Cluster-Mode allows nondisruptive data movement between tiers. There is transparent access to volumes on any node from any node and a connection of many volumes into a single namespace. Cluster-Mode allows you to move volumes between nodes transparently.

3.4 Multi-Tenancy

A cluster is composed of physical hardware, including storage controllers with attached disk shelves, network interface cards (NICs), and optional Flash Cache cards. Together, these components create a physical resource pool that is virtualized as logical cluster resources to provide data access. Abstracting and virtualizing physical assets into logical resources provides flexibility and potential multi-tenancy in Data ONTAP, as well as the data motion capability that is at the heart of nondisruptive operations.

Figure 2) Cluster-Mode architecture.

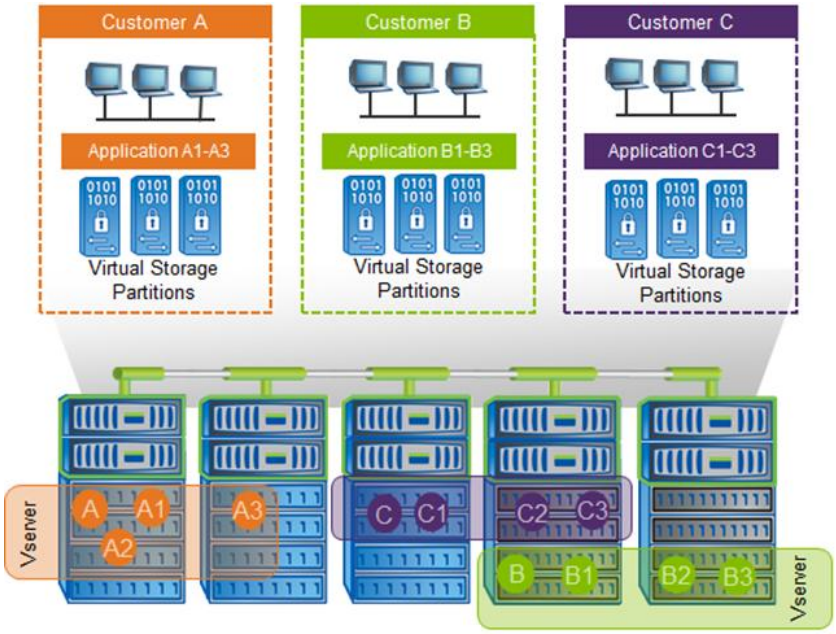

Physical Cluster Components

Although there might be different types of storage controllers, they are, by default, all considered the same in the cluster configuration; they are all presented and managed as cluster nodes. Individual disks are managed by defining them into aggregates, which are groups of disks of a particular type that are protected against failures using NetApp RAID-DP® technology. This is similar to how NetApp Data ONTAP 7G and NetApp Data ONTAP operating in 7-Mode function.

NICs and host bus adapters (HBAs) provide physical ports, such as Ethernet and FC, for connection to management and data networks. The physical components are visible only to cluster administrators and not directly to the applications and hosts that use the cluster. The physical components constitute a pool of resources from which the logical cluster resources are constructed. Applications and hosts only access data through virtual servers that contain volumes and logical interfaces.

Logical Cluster Components

The primary logical cluster component is the Vserver. Data ONTAP supports from one to hundreds of Vservers in a single cluster. Each Vserver enables one or more storage area network (SAN) and network-attached storage (NAS) access protocols and contains at least one volume and at least one LIF. The administration of each Vserver can also be delegated, if desired, so that separate administrators can be responsible for provisioning volumes and other Vserver operations. This is particularly appropriate for multi-tenanted environments or in environments in which workload separation is desired.

Data ONTAP operating in Cluster-Mode facilitates multi-tenancy at the storage layer by segregating storage entities such as aggregates, LIFs, LUNs, and volumes and containing them in a Vserver. Because each Vserver operates in its own namespace, each unit/customer mapped to a Vserver is completely isolated. Each Vserver supports role-based access control (RBAC), and specific protocols such as NFS, CIFS, iSCSI, FC, and FCoE can be assigned to it.

4 Cluster-Mode Terminology

This section defines common Cluster-Mode terminology.

- **Cluster.** In Cluster-Mode, a cluster is a group of connected nodes, or storage systems, that share a global namespace. The cluster can be managed as a single Vserver or multiple virtual servers, which enhances performance and reliability and provides scalability benefits as well.
- **Command-line interface (CLI).** The Cluster-Mode CLI provides a command-based mechanism that is similar to the UNIX® `tcsh` shell in that it provides tab completion, advanced queries, and patterns and wildcards in UNIX style.
- **Epsilon.** Epsilon is an extra partial weight configured to one node. It does not determine the master, but it helps to form a majority. The epsilon node is epsilon for the entire cluster and not just for the individual replicated database (RDB) units. It is manually configurable, but can change automatically. To add an epsilon to a node, use the `cluster modify` command with `–node` and `–epsilon` parameters.
- **HA pair.** This is a pair of nodes, or storage systems, redundantly configured to serve data for each other if one of the two nodes fails.
- **High availability.** This is the recovery capability provided by a pair of nodes, or storage systems, called an HA pair, which are redundantly configured to serve data for each other if one of the two nodes fails.
- **LIF.** A LIF is a logical interface that is mapped to a physical port. A physical port can have up to eight LIFs in Cluster-Mode. A LIF is required to access a Vserver. The three types of LIFs include cluster management, node management, and data management for Vserver.
- **Network.** The cluster network connects nodes to form a cluster, and the data network connects the cluster to the client.
- **Quorum.** A quorum is formed when a majority of the eligible nodes in a cluster are healthy and in contact with one another. There is one quorum per RDB ring at any given time. The node with the lowest SiteID that is online is elected master, while the rest of the nodes are secondary RDB members.
- **RDB.** RDB is a replicated database that stores and maintains the data that manages the cluster. The operations in an RDB are transactional in nature. There are four RDBs, including VLDB, VifMgr, Management, and SpinAuth. This is the key to maintaining high-performance consistency in a distributed environment. Each RDB unit has its own replication ring.
- **Ring.** A ring is made up of one master, which is a read/write database, and other read-only databases. The writes go to the master and are then replicated to others in the ring through the cluster network.
- **SFO.** An SFO is a storage failover. When two nodes are connected it makes an SFO pair. The SFO pair must be of the same controller model. It can be enabled from either node.
- **VIF.** A VIF is a virtual interface, as opposed to network ports, which are physical. There are three types of VIFs and ports: management, cluster, and data. Four physical ports can be grouped into a single VIF.
- **VLDB.** A VLDB is a volume location database. A VLDB contains information about an index of which D-blade owns a volume and serves an aggregate. VLDB content is cached on each N-blade to speed up the data path.
- **Vserver.** A Vserver is a secure, virtualized storage server assigned to a single tenant.
- **Web interface.** The Cluster-Mode Web interface provides a model with which to interact using a Web browser. Access the Web interface through System Manager 2.0 in the Advanced section.

5 Cluster-Mode Limitations and Recommendations

This section describes the limitations of Cluster-Mode and provides recommendations for managing them.

5.1 Limitations

Although Cluster-Mode provides many new features, some current limitations exist, including:

- Large environments require multiple clusters.
- Clusters don't stretch beyond the data center in multisite configurations.
- There are limits to the supported scale of clusters.
- There is no IPv6 support.
- Cluster-Mode–based storage systems don't support NetApp SnapVault® technology.
- There is no support for file-based thin provisioning.
- There is no support for VMDK on NFS and VMFS datastores residing on Cluster-Mode systems.
- There are limits to the different types of hardware and software used in a cluster.
- Storing the RBAC configuration file in the storage system root volume of the Vserver is not supported.
- Users cannot establish a NetApp SnapMirror® relationship between a 7-Mode source volume and a Cluster-Mode destination volume.

5.2 Recommendations

NetApp recommends the following actions or settings to enhance Cluster-Mode operations.

- Configure user credentials for all cluster server and Vserver management end points from the storage system accessed by NetApp SnapDrive® for Windows, SnapManager for SharePoint, and SnapManager for Microsoft SQL Server®.
- For a highly available connection to the storage system, NetApp requires installation of the supported version of multipathing software, such as Data ONTAP DSM 3.5 for Windows MPIO.
- Asymmetric Logical Unit Access (ALUA) support is available on all Cluster-Mode configurations with Data ONTAP DSM, as well as with Microsoft DSM.
- To perform LUN management tasks, there must be a minimum of one iSCSI LIF and one management LIF per Vserver.
- Add the Cluster-Mode IP address and credentials (cluster server and Vserver credentials) along with the management LIF to SnapDrive > Transport Protocols Settings > Storage Systems.
- Map the Vserver and cluster server IP address in DNS, or explicitly specify it in the Windows host file, located in `C:\WINDOWS\system32\drivers\etc`.
- The cutover window defined for a SAN volume must not exceed the expected timeout value on the host side. During the volume cutover phase in a volume move, all input/output (I/O) access is queued and requests are blocked to the source volume. SnapDrive sets a timeout value of 120 seconds on the host during the volume move. Also, when a SAN volume is moved, ALUA is used to optimize access to the volume.
- Each node must have a data LIF for optimized access to the volume.
- Snapshot™ copy scheduling can be performed only if cluster server credentials are provided. It is the responsibility of the cluster server administrator to manage space allocation on the Vservers.
- Virtual guest machines residing on ESX hosts are supported. However, this support is restricted to RDM LUNs only. Cluster-Mode does not currently support VMDK on NFS or VMFS datastores residing on Cluster-Mode systems.
- Provide cluster server credentials in the SDW storage settings to receive NetApp AutoSupport™ tool alerts with Cluster-Mode.

- Intercluster replication requires at least one intercluster LIF per node. The intercluster LIF can be assigned to a data port or a dedicated intercluster port.
- Configure the Windows firewall to allow SnapDrive services for Windows communications.

6 Supported Software Versions for Windows

Table 2 displays the supported software versions for both 7-Mode and Cluster-Mode in a Windows environment.

Table 2) 7-Mode and Cluster-Mode software versions.

7-Mode	Cluster-Mode
Data ONTAP 8.1 or earlier	Data ONTAP 8.1
DSM 3.4 or earlier	DSM 3.5
SnapDrive 6.3 for Windows or earlier	SnapDrive 6.4 for Windows
Windows Host Utility Kit 5.0	No longer required for Data ONTAP DSM. The DSM 3.5 installer includes the MBRAlign tool and the `LinuxGuestConfig.iso` in the installation package. If Microsoft DSM is installed, then Windows Host Utilities 6.0 must be installed.
SnapManager for Hyper-V™ 1.0 P1	SnapManager for Hyper-V 1.0 P1 (no change)
SnapManager for SQL Server 5.2 or older	SnapManager for SQL Server 5.2
SnapManager for SharePoint 6.1 or older	SnapManager for SharePoint 6.1
ESXi 5.0 and previous releases	ESXi 5.0 and previous releases

7 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 roles of the three tiers 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 topology is ideal for a user base of less than 100. This farm consists of a single server handling the requirements of the Web front end (WFE), database server, and application server.

Figure 3) Single-server farm.

In the small-farm environment, NetApp still recommends deploying the SnapManager for SharePoint Media Server on a separate machine so that the box hosting SharePoint has room to serve user requests.

Two-Server Farm

This farm is equipped with one database server and one Web server that performs all the application services. For high availability, NetApp recommends a clustered or mirrored database server. This farm can cater to the requirements of a 100- to 10,000-user base.

Figure 4) Two-server farm.

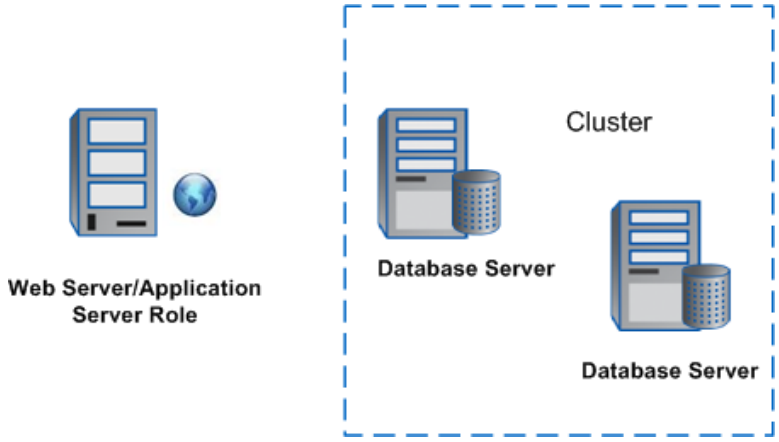

Two-Tier Small Farm

This environment is adequate for a user base of 10,000 to 20,000 with low service usage. One Web server performs the tasks of the application server and one functions as the database server (clustered or mirrored).

Figure 5) Two-tier small farm.

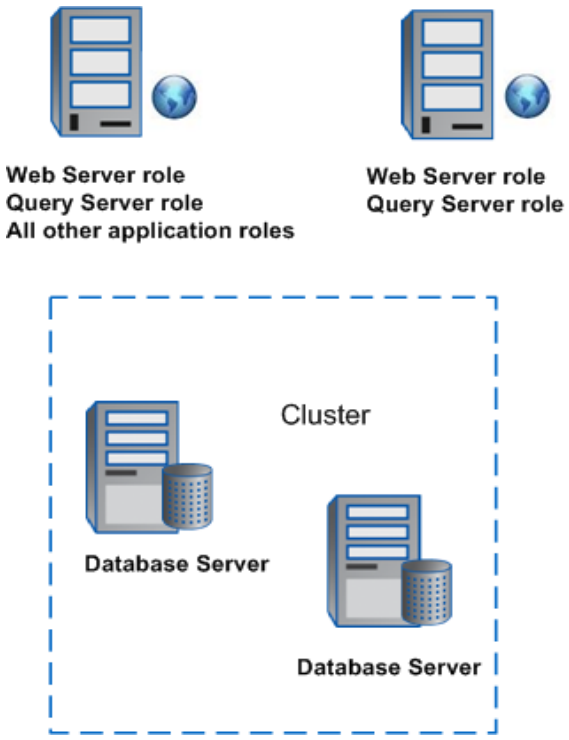

Three-Tier Small Farm

A three-tier small farm is similar to a two-tier small farm. A dedicated application server is added to handle the requirements of these 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 may be considered a large search database. Though it is optional, it is a good practice to have a database cluster or a mirror.

Figure 6) Three-tier small farm.

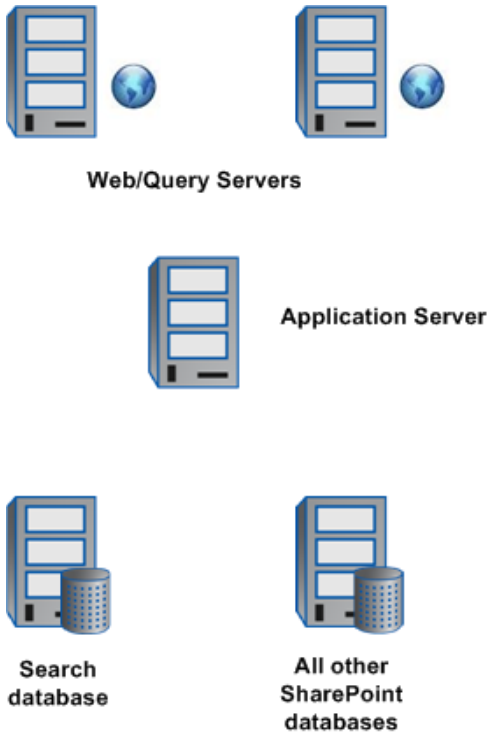

There has been a significant architectural change in SharePoint Server 2010 from its predecessor, SharePoint Server 2007, primarily in its performance and scalability. One of the major architectural changes in SharePoint 2010 is the shared service architecture. To learn more about the service architecture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

Figure 7 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 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. For more details on the features available in different SKUs, see http://sharepoint.microsoft.com/en-us/buy/Pages/Editions-Comparison.aspx.

Figure 7) 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.

8 SharePoint Server 2010 Planning Considerations

Planning is required for SharePoint Server 2010 implementation so that all features are optimized.

8.1 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 learn more about the hardware and software requirements of the server, refer to the TechNet article http://technet.microsoft.com/en-us/library/cc262485.aspx.

8.2 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 must sustain. This involves factors such as 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 the 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 targets vary widely depending on the volatility and criticality of the data stored within the SharePoint farm.
- **RTO** is an acronym for **recovery time objective.** RPO drives the way that SharePoint data protection should be approached prior to a disaster. RTO denotes the timeline within which postdisaster farm and data recovery must be completed. If the data that is stored within SharePoint is highly critical to business operations, then RTOs generally trend toward hours, minutes, or, perhaps, 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 effects 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.
- Three hundred 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 sites 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 back end, because this allows more flexible backup and recovery schedules. Having a single monolithic SQL Server can have implications for RTO/RPO times for data recovery as well as for the backup window.
- The SharePoint 2010 I/O requirement is 0.25 IOPS per GB of content stored. NetApp recommends 2 IOPS per GB for optimal performance.

For SharePoint 2010 Sp1 and later, the content database can hold data up to 4TB. However, there are limitations and constraints for content databases that exceed 4TB. Databases that exceed 4TB are referred to as document archiving and/or for records management use.

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

Table 3 lists the databases that are created as part of the SharePoint deployment, based on the product version and edition.

Table 3) Database details.

Product Version and Edition	Databases
SharePoint Foundation 2010	• 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)
Additional databases for SharePoint Server 2010 Standard Edition	Search service application: • Search administration • Crawl (one or more) • Properties (one or more) User Profile service application: • Profile • Synchronization • Social tagging Web analytics service application: • 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	• 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 default naming for content databases from SharePoint because it contains the Global Unique Identifier (GUID) as part of the name. The reason for this is that the maximum length for the database name is 128 characters. NetApp recommends naming the content database based on the content it contains.

Calculating Content Database Size

Verify 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 (IOPS) while maintaining low latency and queuing on

all disks. SharePoint 2010 requires disk subsystem performance of 0.25 IOPS per GB; NetApp recommends 2 IOPS per GB 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.

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 approximates 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 tax the size of the content database. Cache size can be configured up to 100GB. A 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, and then configure the Office Web apps cache to use that database.

Calculating Index and Crawl Database Size

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

Use the following formulas to estimate the index size.

TotalIndexsize = ContentDBSum * 0.35

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

QueryIndexsize = TotalIndexsize/ (number of index partition)

Next, calculate the crawl database size:

TotalCrawlDBSize = ContentDBSum * 0.046

For more information about sizing the SharePoint deployment, download “SharePoint Server Caches Performance” from the Microsoft Download Center: 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 move the content from a 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 the SQL Server 2008 Enterprise version.

8.3 SharePoint Upgrade Considerations

If your 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 a 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 preupgrade checker to find potential issues.**
  The preupgrade 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 because it can be difficult 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 upgrading.
- 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. Before performing the database attach upgrade, keep in mind that it could increase the size of the database. Third-party tools are available to perform the upgrade and can help to avoid the capacity issues and reduce the complexity of 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.

9 Storage Planning for Microsoft SharePoint Server 2010

9.1 NetApp Storage Software and Tools

Because the deployment of SharePoint is burdensome to storage resources, 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 on 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
  - 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 space-efficient NetApp backup capabilities
  - Streamlines data management and automates 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 8 depicts the different components that are required for the SnapManager for SharePoint solution.

Figure 8) Different components required for SnapManager for SharePoint.

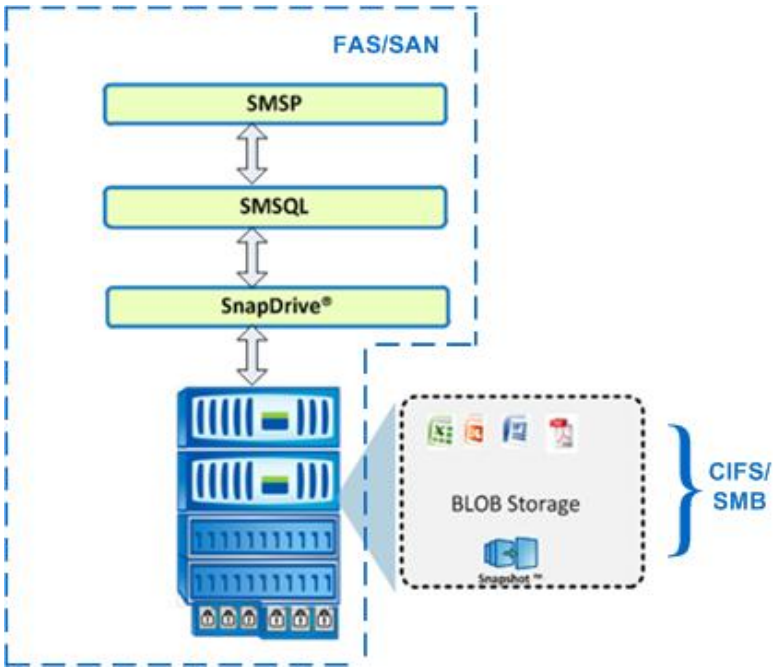

9.2 Principles of Designing Storage for Microsoft SharePoint 2010

Storage efficiency is the ability to store and manage data in a way that consumes the least storage 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 must take into consideration these technologies.

Figure 9 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 together, the larger the savings become.

Figure 9) NetApp technologies for storage efficiency.

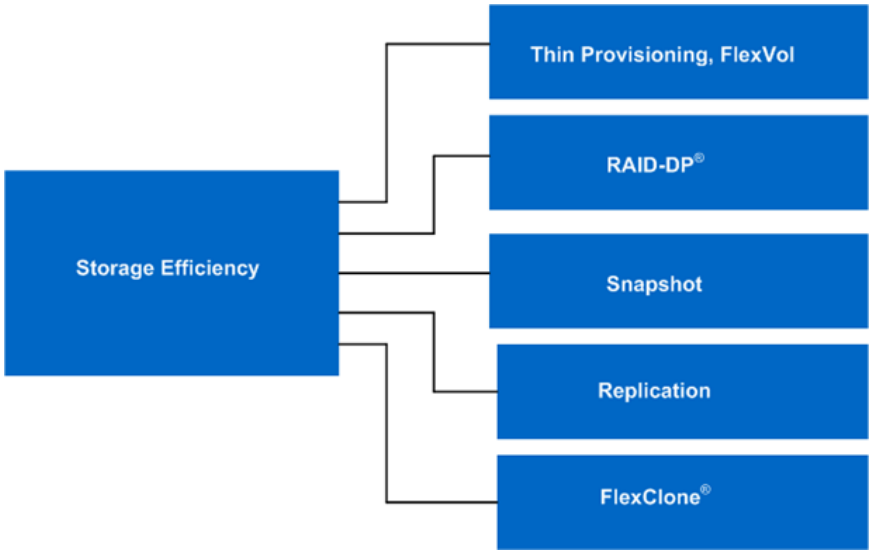

RAID-DP

The performance acceleration provided by the NetApp Write Anywhere File Layout® (WAFL) file system and the double-disk protection provided by NetApp RAID-DP technology make economical and large-capacity 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 now 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 the NetApp implementation of double-parity RAID 6, an extension of the NetApp 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, RAID-DP can 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® technology 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 not 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 on the traditional method of allocating all the blocks up front. Thin provisioning eliminates almost all white space, which helps avoid poor utilization rates. The enabling technology behind NetApp thin provisioning is its 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 can maximize performance; however, it might 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 be improved further by the use of Flash Cache or PAM modules 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 promotes optimal performance of the storage system.

Table 4 details the benefits and trade-offs of two possible aggregate layouts.

Table 4) Aggregate layouts.

Number of Aggregates	Benefit	Trade-off	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 are 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 for separating volumes is to 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

The 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,” this typically enables the amount of space required by the FlexVol volume to be 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 fail if the containing aggregate does not have enough available space. LUN reservation enables the LUN to have space in the volume, but setting guarantee=none does not make certain 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 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 copy 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 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
Enough space must be 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.

A general rule is to keep the user and system databases separate. When you create new MDF and LDF files, presize them to minimize autogrowth 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 the SharePoint database can be configured with either FULL recovery or SIMPLE recovery mode. To learn more about the database and Microsoft’s recommended recovery model, follow this link:

Database type and descriptions: http://technet.microsoft.com/en-us/library/cc974471.aspx.

NetApp recommends setting the file sizes appropriately and growing them manually on a planned basis at times of minimal system use. 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 effect on performance. This impact is minimized but not eliminated through fast file initialization.

Observe the following points about the locations 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 provisioning and FlexVol volumes in the volume planning. The thin provisioning that NetApp recommends includes guaranteeing volumes and LUNs, setting the fractional reserve to zero with volume autosize, and having Snapshot copy 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 is not 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 are 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 must 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 in which testing or development occurs, any situation in which progress is made by locking in incremental improvements, and any situation in which 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.

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

```
volume clone create -vserver smsp_vs -flexclone fc_vol_1 -parent-volume fv2 -
junctionactive

true -foreground true -comment "SnapManager for SharePoint FlexClone"
```

This example command creates a FlexClone volume fc_vol_1 from parent volume fv2 on Vserver vs1 and the job runs as a foreground process.

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, the core NetApp operating architecture. NetApp deduplication combines the benefits of granularity, performance, and resiliency with a significant advantage in the race to address 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 that 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.

There is a limit on the maximum size of the volume for deduplication, because deduplication depends primarily on the amount of system memory, which varies based on the storage platform. In considering using deduplication in SharePoint 2010 environments, be sure to consider this factor when sizing the volume layout.

Best Practice
Replication of deduplicated volumes is supported using SnapMirror, but do not use synchronous SnapMirror to replicate volumes 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 significantly 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 enables individual root-level sites within the SharePoint Web application to be recovered.
- **Web.** This enables subsites under the root-level sites to be recovered.
- **Folder.** This enables document libraries and lists to be recovered.
- **Item.** This enables each item in the SharePoint Web application, including lists, announcements, documents, calendars, events, and so on, to be individually recovered.
- **Item version.** This enables all versions of each item in the SharePoint Web application, including lists, announcements, documents, calendars, events, and so on, to be individually recovered.

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 datastore. Note that the change in the number of objects increases as the level of granularity becomes finer. This means that it is possible to 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 in which SMSP stores the index created as 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:

LUN_SIZE = N * (X+Y) * 0.1 * (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:

VOLUME_SIZE = (LUN_SIZE * 1.1) + [(number of Snapshot copies to keep online) * (0.1 * (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 must 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 must be migrated to a NetApp storage system must first be laid out according to the best practices for keeping the SQL databases on 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 1TB 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 volumes 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 LUNs 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's Snapshot copy can take a long time since it is essentially 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 factor in 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® technology at the volume level for faster restores.
- Compression and deduplication can be used together and are transparent to Microsoft SharePoint.
- Create multiple CIFS shares, one per volume, for each Web application in each farm.
- Due to their different I/O activity patterns, create separate CIFS shares for archived BLOB data and for extended BLOB data. 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. Verify that the security is set to the SMSP user account used to access these shares.
- Each volume designated for CIFS shares should be a minimum of 1GB. There is no defined maximum limit; the maximum depends on the number of volumes created 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 might 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 a 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 optimal 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 its components such as 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 multithreaded and deals with large amounts of data). A 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.

10 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, as well as storage optimization and scalability improvement for SharePoint server deployments.

10.1 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 data center administration costs 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 leverages the EBS/RBS framework from Microsoft in SharePoint 2007 and 2010 and SQL Server 2008
- Offers enhanced command line support for data protection along with integration with Windows PowerShell v2
- Provides content archiving and content migration capabilities from EBS to RBS and from filestream to RBS
- Delivers centralized management of backup and recovery for multiple SharePoint farms
- Can back up all objects and entities of a SharePoint farm
- Integrates with NetApp SnapMirror
- Integrates with the Microsoft Operations Manager (MOM) and System Center Operations Manager (SCOM) event log and NetApp AutoSupport™ (ASUP™) system for streamlined support

10.2 SnapManager 6.1 for SharePoint Architecture

SnapManager for SharePoint Server is a 64-bit application that can protect SharePoint 2007 and SharePoint 2010 databases. It was designed to provide:

- Content archiving and content migration capabilities like migrating from EBS to RBS
- Centralized management of backup and recovery for multiple SharePoint farms
- Ability to back up all objects and entities of a SharePoint farm

Figure 10 illustrates the SnapManager for SharePoint data protection and storage optimization architecture, showing its relationship with the SMSP Manager, Media Server, Control Agent, Archiver/Extender Agents, Member Agents, and so on.

Figure 10) SMSP data protection and storage optimization architecture.

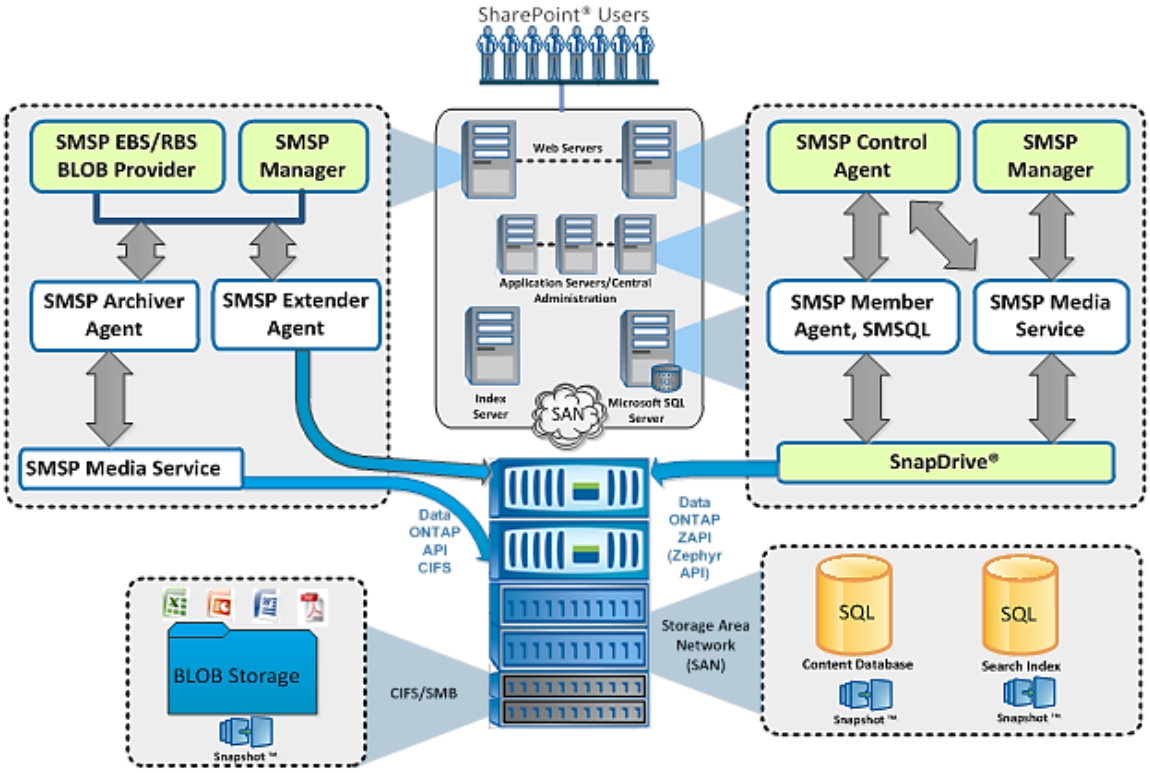

To meet the objectives listed under section 8.2, SMSP uses an agent-based architecture. These agents help provide centralized management, and they automate most backup and recovery tasks.

This section describes the components of SnapManager for SharePoint.

- SMSP Manager
- SMSP Media Server
- SMSP Control Agent
- SMSP Member Agent
- SMSP Archiver, Extender
- SMSP EBS/RBS Provider

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. The 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. The 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 back up 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 is granular and allows the creation of rich business rules around content that is archived and eventually aged out.

NetApp recommends 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 better scalability and, when used with NetApp storage, you can leverage storage size compression and deduplication

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 enables communication between the agents to flow easily, even with complex deployment scenarios such as demilitarized zone (DMZ) deployments in which 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 storage controller, SMB signing must be enabled on the NetApp controller as well as on the WFE servers.

Note: Performance impact on the controller and the WFE server must be considered before implementing SMB signing. There is a performance overhead of 10–15% when 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 5 summarizes the ports used by the SMSP components.

Table 5) Ports used by SMSP components.

Component	Service Name	Default Port	Description
Manager	Web service	8080 (HTTP port for external communication) 8443 (HTTPS port for external communication) 12002 (internal communication) 12011 (access the CLI)	Web server for admin console
	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/Extender

Best Practice

NetApp recommends that the SMSP default port settings 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 the range of 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: Customers using SnapManager for SharePoint with Avepoint DocAve should note that the ports must 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 listed in Table 5 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

SnapDrive for Windows and SMSQL require the following port: 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 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 belong to 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 11 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 may 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 11) Account Manager.

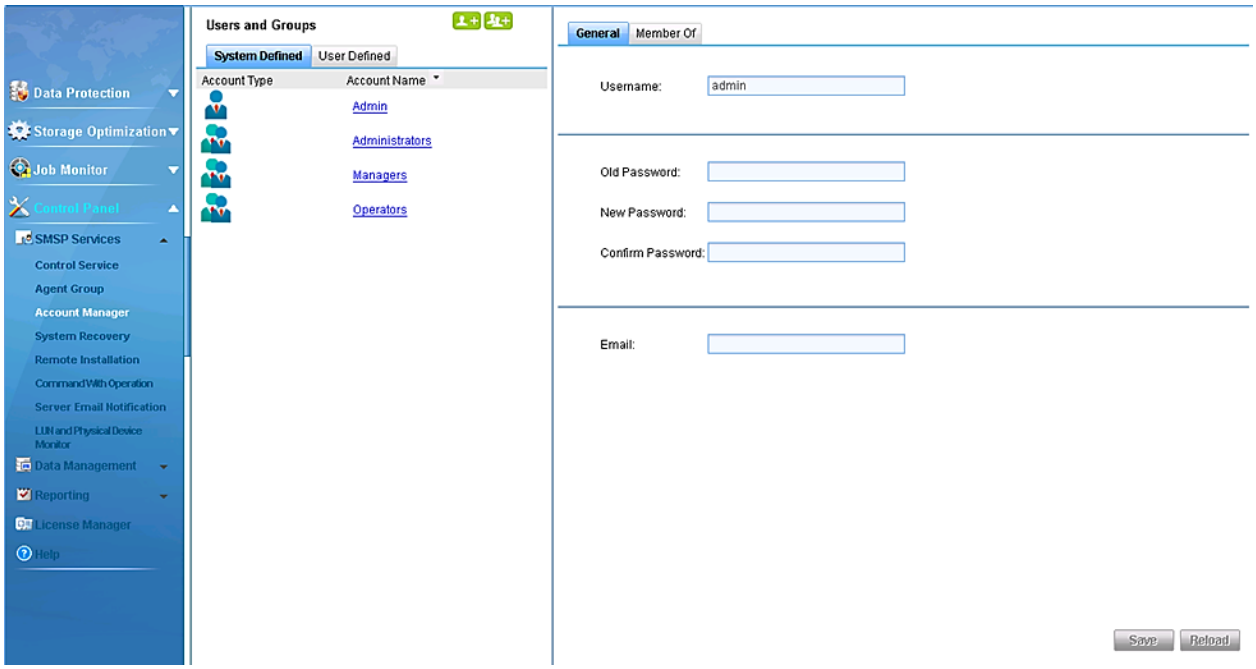

Best Practice

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

10.3 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 migrates 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.

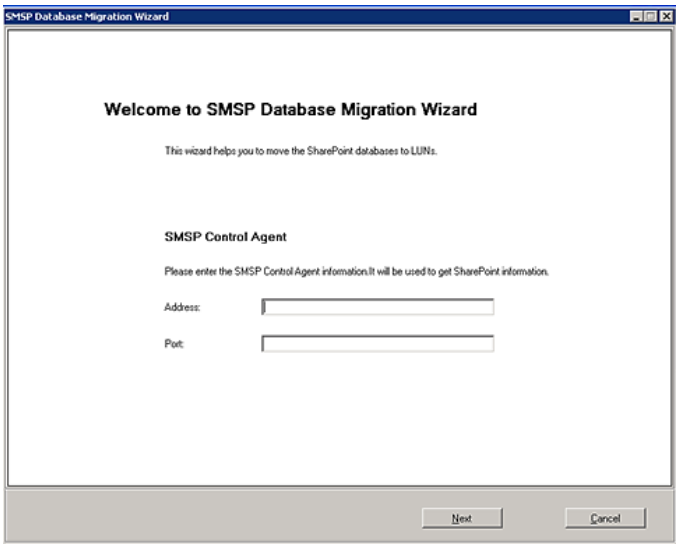

2. Stop all the services listed in Table 6 in the SharePoint farm.

Table 6) SharePoint Server 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
IIS Service	IIS Admin Service
Office SharePoint Server Search	SharePoint Server Search 14
Windows SharePoint Services Search	SharePoint Foundation Search V4

3. 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 you plan 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.

4. Start the services listed in step 2.

Best Practice

The SharePoint VSS service startup type must be manual; do not change it to automatic or start this service manually.

10.4 Upgrading SMSP 6.0 to SMSP 6.1

SMSP 6.1 supports seamless upgrading from the earlier versions 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 coexist with EBS in SharePoint 2010 environments.

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

Best Practices

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 an 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

10.5 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 7 to understand the role of the servers in a SharePoint farm and the SMSP components that should be installed on them.

Table 7) 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

10.6 Security Permission Required in the Configuration of the User Account in SQL

A set of permissions is required on the Microsoft SQL Server for the SnapManager for SharePoint service account.

- 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 that is then applied into this group.
- This account must have SYSADMIN rights.

Figure 12) Login properties—DETDC\administrator.

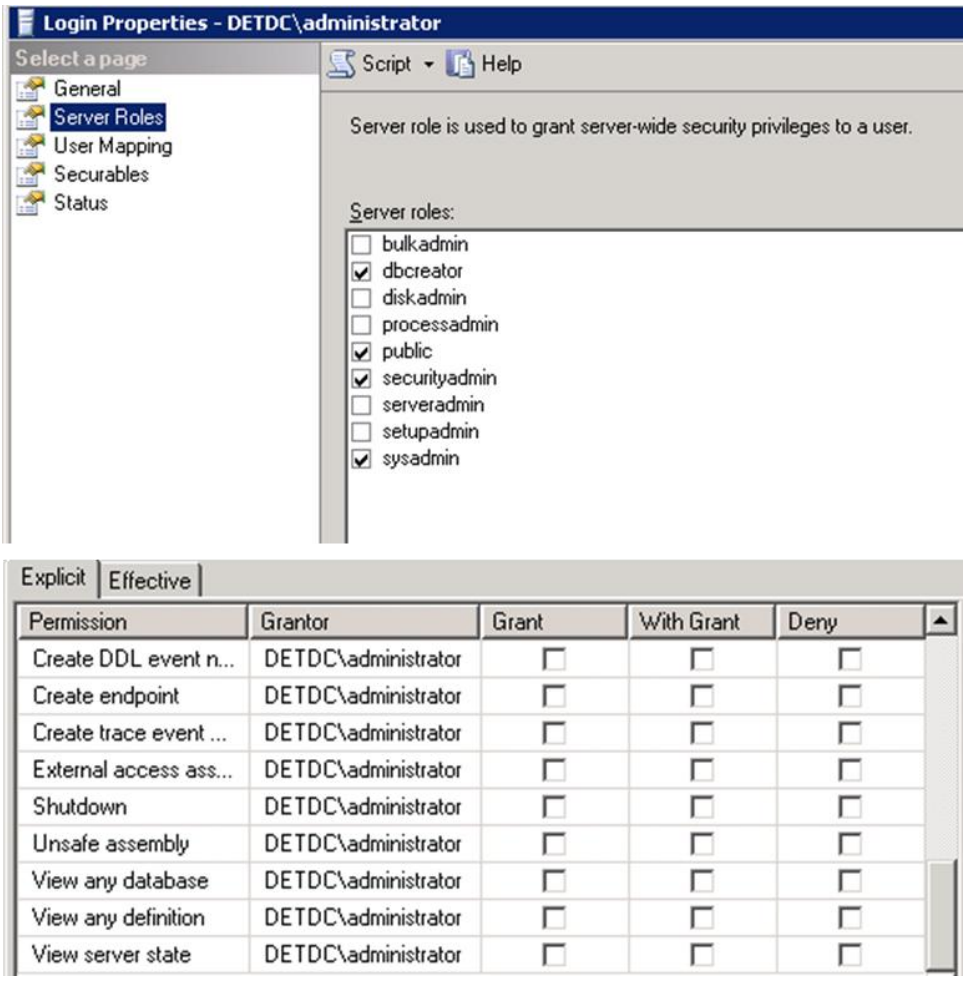

The SMSP account discussed in this section must have local administrator permissions for any Windows Servers that have SMSP agents installed, including:

- SQL servers
- Application servers
- SMSP Manager server
- Web front-end servers

Where group policy is applied to Windows machines with restrictive settings, local administrator permission does not suffice. Edit the local security policy (and/or default domain policy) for these servers and add the following:

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

10.7 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 off-loading to a CIFS share on

lower-cost storage. BLOB offl-oading is done either by 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 2008 R2. To run RBS on a remote server, you must be running SQL Server 2008 R2 Enterprise edition. The advantages of RBS over EBS include much tighter integration with SQL Server 2008 R2 ENT. Microsoft recommends using RBS for any new deployments, because EBS is supported through 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 systems.
- 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 provides 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 both the I/O and processing load on the SQL Server database server.
- If your current SharePoint version is 2007 or earlier, you want to move to SharePoint 2010, and you have no externalized content, 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 preferred option is to use SMSP6.1 and perform one of the EBS-to-RBS direct stub migration options that are available.

Device Manager in SMSP

SMSP Media Service manages data by using the following device types.

- **Backup device.** For backing up job data, this device can be on the local disk, LUN, or UNC path. NetApp recommends using a NetApp LUN.
- **Archived index device.** For the Archiver/Extender index and full text index, this device can be on a LUN or an SMB share on NetApp storage. NetApp recommends using a NetApp LUN.
- **Archived data device.** For Archiver/Extender data, this device can be on a LUN or an SMB share on NetApp storage. NetApp recommends using a 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 in the list. The SMSP Manager checks the status of the physical device every 30 minutes. When no free space is left on the device, if sufficient space exists 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 administrators to configure logical devices for backup jobs and archive jobs. There are five kinds of logical devices.

- Local
- Network
- LUN
- CIFS share
- SnapLock

For each logical device type, multiple physical devices can be defined. Once a physical device is full, data is automatically 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 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 is reused after the checking is done.
- If EBS and RBS coexist in a SharePoint farm, create separate physical devices for them.

10.8 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 a separate backup plan for backup of WFE servers from the database backup. This helps prevent the SMSP backup index from getting bloated, especially if you have large amounts of GAC data.
- Backup maintenance plans are very important because database verification and granular restore index generation during a backup job create overhead in the server and increase the duration of the backup job. 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 because 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 the 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 OnCommand™ Unified Manager is configured for SMSQL.

Best Practices

- To protect the Media Server from failures at the primary site, use SnapMirror to replicate the Media Server backup data to a secondary storage system at a DR location.
- 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 an 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 OnCommand Unified 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.

10.9 Retentions

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

- Backup data retention. The backup data retention policies are a method for 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, along with the number of backup cycles.
- Archived data retention. The archived data retention policies are a method for pruning the old archived data to make room for new archived data. The archived data retention rule specifies the retention period to keep the archived data in the given interval for one Archiver plan.

Archiver should be used in scenarios in which data retention policies are well defined because you can't recover BLOBs that are deleted after the policies have expired.

- 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 six 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 six 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 length.

10.10 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.

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 must be specified.

You should 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 be restored either to a site or to a site collection, and a list or library should 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 must 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 6.1 for SharePoint Installation and Administration Guide.

10.11 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 centralized monitoring. Configure these settings under the “Reporting” module in the SMSP Management GUI, through AutoSupport settings, and 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. They will be available in the NetApp communities with instructions so that the user can monitor all events created by SnapManager products. In SnapManager for SharePoint 6.1 the events are stored under a separate tree, as shown in the following screenshot.

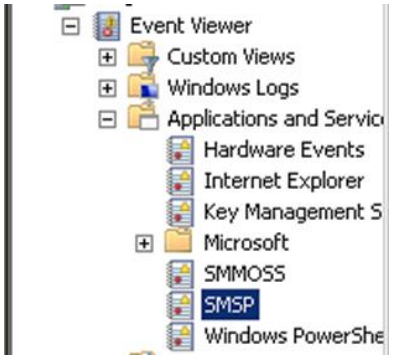

Best Practices

- The AD account specified in the SnapManager configuration tool is used for SCOM integration. This account must be the database owner for the SCOM database.
- The status of OpsMgr SDK service on the SCOM server should be "Started."
- The user account entered in the SnapManager configuration tool must be the IIS (MOM Site) pool user or local administrator.

10.12 Troubleshooting

SnapManager for SharePoint Log Manager 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 in a log file called SMSP.

Figure 13) SnapManager for SharePoint Log Manager.

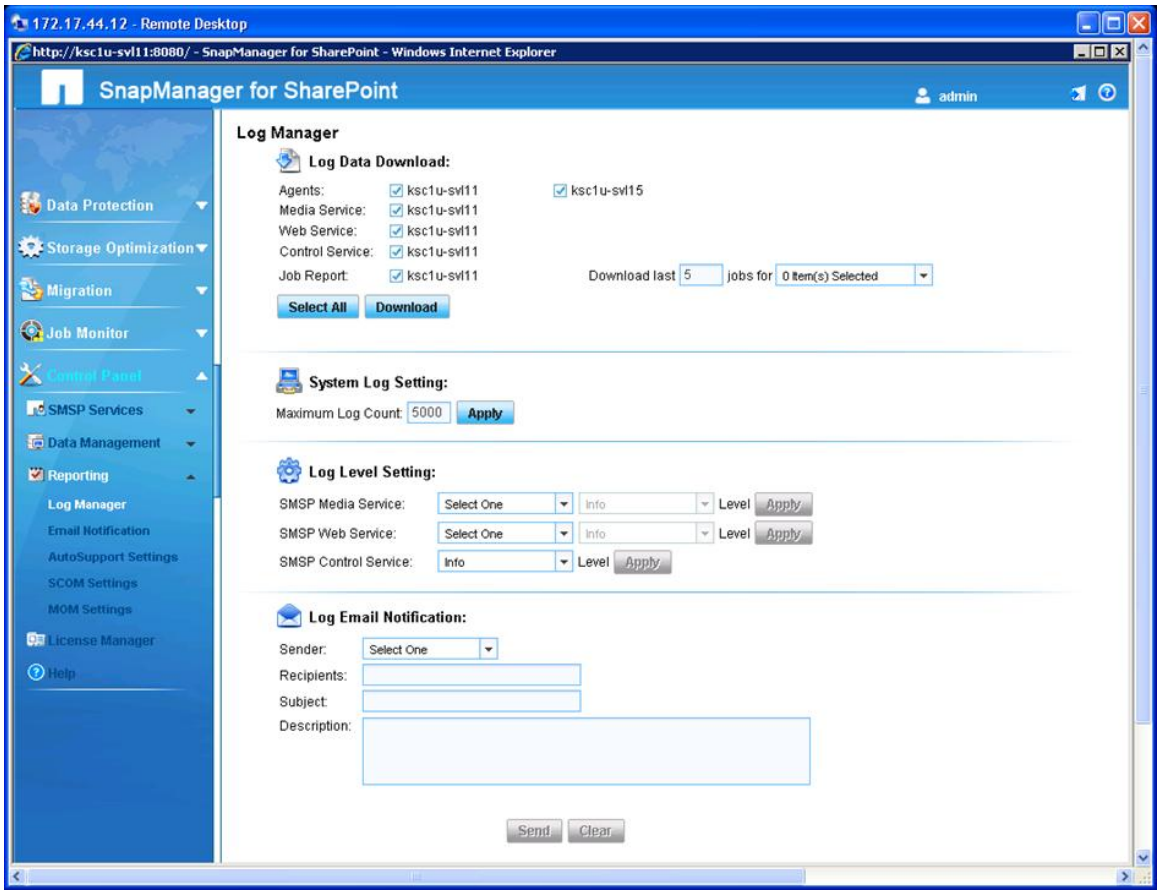

Log Data Download is used to centrally download logs from multiple SnapManager for SharePoint components, including the control, media, and Web services, as well as agents. Select all agents or services for which you require logs and click Download and save the consolidated .zip file to your specified 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 contacting Technical Support, NetApp recommends that you create debug logs. Debug logging can be enabled by setting the logging in Log Manager.

To select from various levels of logging, 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.

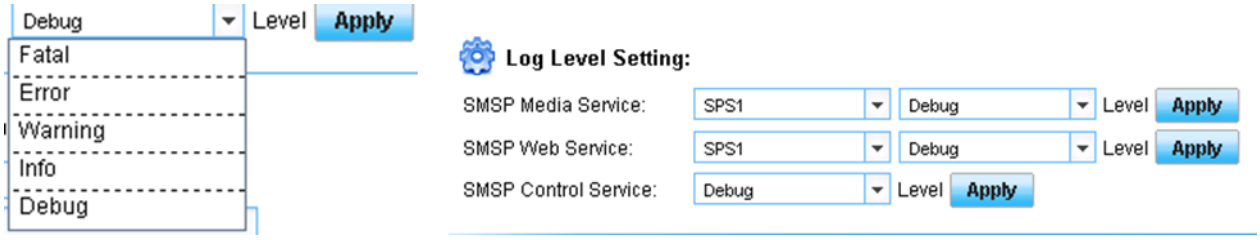

After the selection is complete your debug level logging for the different services is set.

11 Performance

For peak performance in terms of 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 in which the individual content size is small enough (for example, less than 500KB), keep all of the content in the SQL database itself.
- For larger objects, deploy EBS or RBS to keep the BLOBs in the external SMB shares on cheaper SATA disk storage.
- Use Flash Cache (PAM modules) on the controller to optimize performance.

To maximize performance in a 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, have adequate memory size and CPU power.
- Both the archived data device and the archived index device must be on a CIFS share.

Note: To learn more about EBS performance statistics for NetApp storage systems, refer to TR3874: Performance Report of SharePoint 2007 EBS Provider for NetApp Storage System.

Best Practice

As a best practice, create logical devices with the storage name in lowercase. For details, see https://kb.netapp.com/support/index?page=content&id=2015280.

To maintain a 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.

12 High Availability

12.1 Application Resiliency

High availability (HA) is an important requirement for business-critical applications like SharePoint when downtime can adversely affect business. SharePoint farm setup can involve load-balanced Web front-end servers, application servers for central administration, search, Excel services, and so on. Choosing the appropriate solution requires careful consideration when 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 might 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 a 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-availability mirroring 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. SQL Server 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 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 SnapManager for SharePoint on a cluster is available in the SnapManager 6.1 for SharePoint Installation and Administration Guide.

Based on the preceding discussion, we realize that some components can be made redundant in SharePoint. Table 8 describes the server roles in a SharePoint Server 2010 environment and their redundancy.

Table 8) Server roles in the 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 hardware or software network load balancing.

Server Roles and Redundancy Strategy in a Farm Server Role	Redundancy
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.

12.2 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 most fully 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 14 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 combined, the more resiliency and high availability can be achieved.

Figure 14) NetApp technologies for storage resiliency.

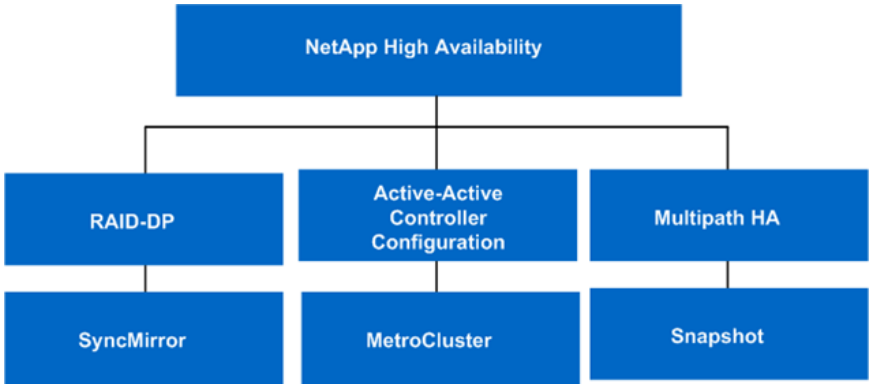

NetApp RAID-DP

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 are:

- It protects against double disk failures or uncorrectable bit read errors from the disk while in reconstruction mode.
- There is no impact on required capacity because RAID-DP groups can be double the size of RAID 4 groups.
- RAID-DP enables nondisruptive upgrade of disk firmware, which results in zero downtime.

Best Practices

NetApp recommends using RAID-DP 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 SnapMirror

SnapMirror in Cluster-Mode provides asynchronous volume-level replication based on a configured replication update interval. SnapMirror replicates Snapshot copies from a source volume to a partner destination volume, thus replicating source object data to destination objects at regularly 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.

Active-Active Controller Configuration

Active-active controller configuration (previously known as clustered failover), eliminates 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 effect 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 a smaller performance effect on SharePoint 2010 workloads (assuming that the storage processor CPU or 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 the performance impact on SharePoint 2010 databases, but the reconstruction process takes longer.

NetApp Snapshot Copies

NetApp Snapshot copies are backups of how a volume looks at a particular point in time. Snapshot copies work almost instantaneously, with very little effect 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.

13 Virtualization

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

During the planning of virtualization it is necessary to evaluate and decide between the virtualization technology and the differentiating factors of multiple vendors, specifically Microsoft Windows Server 2008 R2 or Hyper-V Server 2008 R2 or the VMware ESX virtualization stack.

Here are some of the best practices for a virtual deployment of SharePoint.

- **Use hardware-assisted virtualization.** Hardware-assisted virtualization (HAV) helps improve the performance of virtualization software and application response times. It has been observed that HAV offers a 5% or greater increase in throughput. Hardware-assisted virtualization is provided by the Intel® Virtualization Technology processor family and the AMD Virtualization processor family.
- **Enable hyperthreading on processors that supports this technology.** Hyperthreading is a technology that is available with processors that support symmetric multithreading. This technology provides two threads for each processor core. The net effect is a doubling of logical processors for each core.
- **Configure nonuniform memory access correctly.** Nonuniform 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.** The following can help you obtain the optimum performance from the host server in a virtualized environment.
  - Dedicate the host computer to Hyper-V. Do not run additional applications on the host.
  - Install and configure only the required roles and services.

  - Do not store VM data on system drives. The root partition should only contain the OS and Hyper-V.
  - 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 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(WS.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. NetApp recommends fixed-size disks. Fixed-size disks maximize overall performance and reduce the risk of exceeding physical storage space.
- **Configure general virtual machine settings for performance and stability.** A few points can help obtain optimal performance and stability from the virtualized environment. Prevent boot swarm by using time delay if VMs are configured for automatic start. Shut down the VM exactly the same way as a physical computer; don't 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 in time. NetApp does not recommend using VM snapshots in the production environment. Microsoft SharePoint Server uses timer job extensively; taking VM snapshots affects time-sensitive operations and can result in data corruption or loss. NetApp offers applications like SnapManager for SharePoint and SnapManager for Hyper-V to remove the limitations by taking block-level Snapshot copies.
- **Design the virtual topology for optimal performance.** In order to gain maximum throughput in the SharePoint farm, design the topology across hosts by mixing farm server roles on each host and monitoring the performance of the WFE on a single host. For example, mix WFE servers with application servers on each host. This measure reduces disk contention because usually these servers do not write to disk at the same time.
- **Do not overload the Hyper-V host.** Make sure that the processor and memory are not oversubscribed, because this has a 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 prevents resource-intensive jobs from running at the same time. For example, if execution of a backup program happens on the physical computer and the virtual machines at the same time, resource contention occurs.

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 these links:

- **Best Practices for Virtualization (SharePoint Server 2010)**

http://technet.microsoft.com/en-us/library/hh295699.aspx

Perform the hardening of the Hyper-V role and best practices to reduce the attack surface of Hyper-V and the recommendations to configure the secure virtual networks. Managing and delegating the virtual machine management securely within the organization are essential too. For more information see the following 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 practices for Hyper-V and SMHV, read through these guides.

- **TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide**

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

- **TR-3702: NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V**

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

- **TR-3749: NetApp and VMware vSphere Storage Best Practices**

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

Table 9 describes the SharePoint Server roles and the virtualization decision that is made by each in the deployment stage.

Table 9) SharePoint Server roles and 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 started with the WFE (Web front end) because this is the role that renders content. The WFE is the most commonly virtualized role in a SharePoint farm because of its smaller memory and disk requirements. The consideration of the WFE for a virtual environment is quite simple; we can provision additional servers for load balancing and fault tolerance. Application and patch levels are the same on all the servers that are provisioned.

Then we have the query server role, which processes search queries. This server requires the propagated copy of the local index. The propagated copy of the index could be large depending on the type of content that is getting indexed. If we have the query role with the WFE, we build a high-availability solution into the search service.

We must also account for application servers that host services like Excel form services or any other application service role. Provision more servers (VMs) as resource requirements for individual applications increase.

Investigation is required to determine whether the index server role should be virtualized, because this server performs the crawl on the content repository, is resource intensive, and needs more memory and processing than the Web server role.

A database server with lower resource usage requirements can be virtualized. However, the read/write requirements of a virtualized database server generally tend to tax the performance of the physical host. This is why it is usually advisable to keep the database on the physical machine.

With the release of SharePoint Server 2010 SP1, additional planning may be needed to manage these databases effectively. For more information, see the following article:

- **SharePoint Server 2010 Capacity Management: Software Boundaries and Limits**

http://technet.microsoft.com/en-us/library/cc262787.aspx

14 Disaster Recovery

Various methods can be used 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 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 10 summarizes the replication models that can be used for Microsoft SharePoint Server 2010.

Table 10) Replication models.

Configuration	Replication Model
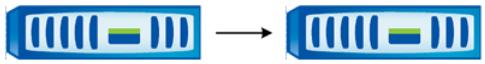	Unidirectional replication
	Bidirectional replication
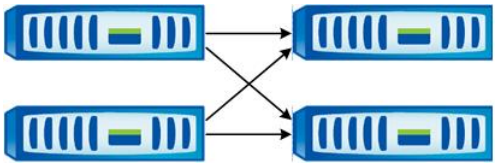	Multidirectional replication

Setting Up SnapMirror Relationships

After the cluster peer relationship has been successfully created between the two clusters, create the intercluster SnapMirror relationship. Run this command to check the cluster peer:

```
SMSP::> cluster peer show
Peer Cluster Name         Cluster Serial Number Availability
------------------------- --------------------- --------------
APPCL                     1-80-000011           Available
SMSP                      1-80-000011           Available
2 entries were displayed.
```

Use the hostname of the source and destination clusters to create SnapMirror relationships in Cluster-Mode. The hostnames of the source and destination clusters are exchanged and stored while creating the cluster peer relationship. If the hostname resolution is not configured on the clusters, still create SnapMirror relationships using only the names of the clusters;

Complete the following requirements before creating the intercluster SnapMirror relationship.

- Configure the source and destination nodes for intercluster networking.
- Configure the source and destination clusters in a peer relationship.
- Create a destination NetApp Vserver; volumes cannot exist in Cluster-Mode without a Vserver.
- Verify that the source and destination Vservers have the same language type.
- Create a destination volume with a type of data protection (DP) and a size equal to or greater than that of the source volume.
- Assign a schedule to the SnapMirror relationship to perform periodic updates. If any of the existing schedules are not adequate, a new schedule entry must be created.

After all of these requirements are completed, create a SnapMirror relationship:

```
OBAV::> snapmirror create -source-path  SMSP://sql/testsqlvol -destination-path
OBAV://test1/sdtest_DP -type DP -tries 8 -throttle unlimited -vserver test1
[Job 1962] Job succeeded: SnapMirror: done
```

After the creation of the SnapMirror relationship, review the relationship and initialize it in the destination folder to confirm the progress of the replication:

```
OBAV::> snapmirror show
Source              Destination  Mirror        Relationship  Total
Path          Type  Path         State         Status        Progress   Healthy
------------ ---- ------------ ------------- -------------- ---------- -------
APPCL://infraserver/mirrorprim
              DP   OBAV://test1/mirrordest
                                 Snapmirrored  Idle           -          true
APPCL://sqlserver/s
              DP   OBAV://test1/dst1
                                 Snapmirrored  Idle           -          true
APPCL://sqlserver/sqldb
              DP   OBAV://test1/sql
                                 Snapmirrored  Idle           -          true
                   OBAV://test1/sqldestination
                                 Snapmirrored  Idle           -          true
SMSP://sql/testsqlvol
              DP   OBAV://test1/sdtest_DP
                                 Uninitialized Transferring   -          -
5 entries were displayed.
```

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 replicated 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 when the backup job finishes.

Note: Both database volumes and BLOB store volumes should be configured in a 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 that allows the backup of system settings used in that SnapManager for SharePoint deployment.

System Backup saves the system settings to a flat file in the location defined under Backup Destination. The default path is `C:\Program Files\NetApp\SnapManager for SharePoint\UserData`.

Best Practices

- Use a nondefault location. For disaster recovery purposes, NetApp advises changing 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 to the DR site.
- As described in "SharePoint Components Backup" in section 10.8, the Media Server backup data should also be protected by moving it 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 it is 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 so that the topology is the same.
- If a separate AD domain is used at the 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.

14.1 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 interrupting 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 11 describes the LUN status on SnapMirror destination systems during different replication states.

Table 11) LUN status on SnapMirror destination systems.

Replication State	DR LUN Status at the SnapMirror Destination
Mirror being asynchronized	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 size of the destination volume to be equal to or greater than that of 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 do not 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 entity 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, affects the applications that use the storage. Therefore, it is important to understand the business requirements for application performance and data protection in order to enable informed decisions about 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 perform all operations that are normally required and do additional processing related to SnapMirror to transfer information to the destination storage. This adds significant CPU overhead 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 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 Servers

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 their performance characteristics, it is important to note that in Cluster-Mode the replication mode must be asynchronous.

Replication Mode	Achievable RPO	Performance Impact
Asynchronous SnapMirror	Minutes to hours	Low

15 Data Protection

Data protection is an important component of the solution, and Data ONTAP 8.1 offers these capabilities using SnapMirror.

15.1 SnapMirror

Data ONTAP 8.1 operating in Cluster-Mode allows two kinds of replication:

- **Intercluster asynchronous volume replication.** Replication between volumes hosted on different clusters that enable disaster recovery replication to a cluster in a remote site. This type of replication requires one intercluster LIF per node, and it must use the data port or the dedicated intercluster LIF.
- **Intracluster replication.** Replication between two Vservers in the same cluster.

Note: You cannot establish a SnapMirror relationship between 7-Mode source volumes and Cluster-Mode destination volumes.

15.2 SnapVault

Data ONTAP 8.1 operating in Cluster-Mode does not have SnapVault capability; therefore, SnapDrive for Windows does not support SnapVault when it is connected to Cluster-Mode systems.

15.3 IPv6

Data ONTAP 8.1 operating in Cluster-Mode does not support IPv6; therefore, SnapDrive does not support IPv6 access to Cluster-Mode systems.

16 SnapManager Dependency on SnapDrive

SMSP 6.1 directly communicates with SnapDrive for Windows for application-consistent Snapshot copies.

Note: The SnapManager suite of products uses SnapDrive for application-consistent Snapshot copies. NetApp recommends that you make sure there are no overlaps between the application-specific Snapshot copies and their respective products before the Snapshot copies are initiated, because this reduces the performance overhead on the storage systems.

17 Conclusion

SnapManager for SharePoint 6.1 and Data ONTAP 8.1 operating in Cluster-Mode provide a solution that allows nondisruptive operations during storage infrastructure maintenance and upgrades, enabling higher availability and enterprise-ready unified scale-out storage. This solution is the basis for a large, virtualized shared storage infrastructure, and, most importantly, it is built on the solid foundation of Data ONTAP.

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

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

Go further, faster®

www.netapp.com