



SharePoint MOSS 2007 Portal Implementation

Sunshine Connections

Table of Contents

Hardware Configuration	1
Baseline Architecture – Medium Server Farm (5 Servers)	1
Infrastructure Requirements and Configuration	1
Network Interfaces	2
DNS	3
Active Directory	3
Physical and Logical Architecture	3
Physical Architecture	3
Logical Structure	3
Software Requirements Web and Application Server	3
Security Account Requirements	4
Planning DNS Names	5
Planning Backup and Restore	5
Disk Space Requirements for Backups	5
Backup Strategies	6
SharePoint Portal Server Database Backups	6
My Sites Backups	6
Windows SharePoint Services Site Backups	7
Front-End Web Server Backups	7
Images	7
Database Configurations	8
Database File Configurations	8
General Database Sizing Guidelines	8
Backup Plan for Sunshine Connections	9
SharePoint Backup Utilities	9

Introduction

The user interface (UI) of Sunshine Connections is implemented using Microsoft Office SharePoint Server (MOSS) 2007 technology. Data from DOE information systems will be rendered in the SharePoint environment via reports and other views contained in web parts. Application functionality may be contained in web parts, or in ASP .Net application components. In either case, most UI components of Sunshine Connections will access web services rather than source data systems to provide a layer of abstraction from underlying data systems to avoid having changes in those data systems force changes in application logic and UI components, and to keep data sources isolated from external internet access. Users are challenged for login credentials by the reverse proxy server (a Microsoft ISA Server), which authenticates the credentials against the Microsoft Active Directory domain controller and gives access to areas of the portal by role. All content in the Sunshine Connections portal is specific to the roles that are currently deployed (Teacher, Principal, and Superintendent).

Additionally, the SharePoint UI provides collaboration functionality such as document management and sharing, shared workspaces, threaded discussions, alerts, surveys, and calendars.

Districts and vendors may develop their own UI components, or add to existing application UIs depending on local requirements and local technology platforms by consuming the same web services that the web parts in Sunshine Connections use. Districts and vendors may also create their own web parts if they are using SharePoint locally or some other Web development technology (ASP .NET, JSP, PHP, etc.). The possibility exists to directly share web parts among districts and the DOE where SharePoint is being used, or provide views of those web parts in the Sunshine Connections portal if they are hosted in other locations.

The first release of Sunshine Connections will contain core FCAT analysis functionality which allows teachers to see longitudinal student performance data for each student assigned to their upcoming fall classes. Additionally, as a "support" function, Sunshine Connections is the web-based portal environment, authentication mechanism and data provider that will enable the rollout of Florida's statewide, pre-populated, online Special Education process Web System which includes functionality for the Individual Education Plan (IEP). Subsequent releases of Sunshine Connections will broaden the audiences and make available the vast repository of educational data housed by the Florida Department of Education.

Hardware Configuration

Baseline Architecture – Medium Server Farm (5 Servers)

The existing hardware configurations for front-end Web and search servers in the Sunshine Connections server farm are:

Front-end Web/search servers (2 servers)

- Dell PowerEdge 1800 Dual Xeon Tower Server, 3.0GHz/1MB, 800MHz Front Side Bus
- Minimum 2 GB main memory
- 4 x 36 GB Ultra320 SCSI disks

- 2 x embedded 10/100/1000 gigabit Ethernet network interface card (NIC) ports
- Redundant power supply and fan

Index management/job server (1 server)

- Dell PowerEdge 2800 Dual Xeon Tower Server, 3.0GHz/1MB, 800MHz Front Side Bus
- Minimum 2 GB main memory
- 8 x 36 GB Ultra320 SCSI disks
- 3 x 10/100 Ethernet NIC ports
- Redundant power supply and fan

SQL Server-based database server (active/passive; 2 servers)

- Take advantage of the existing Sunshine ConnectionsSQL Server infrastructure.

Infrastructure Requirements and Configuration

Figure 2.1 shows the baseline high-availability configuration. There are two separate networks, each connected by a dedicated 100-megabit Ethernet switch. The VLAN1 switch connects the corporate LAN to each of the front-end Web/search servers and handles all incoming HTTP traffic as well as the Network Load Balancing (NLB) heartbeat traffic. The VLAN2 switch connects each front-end Web/search server to the SQL Server-based servers and the job/index management servers. A crossover cable is used between the SQL Server-based servers to carry SQL Server clustering heartbeat data.

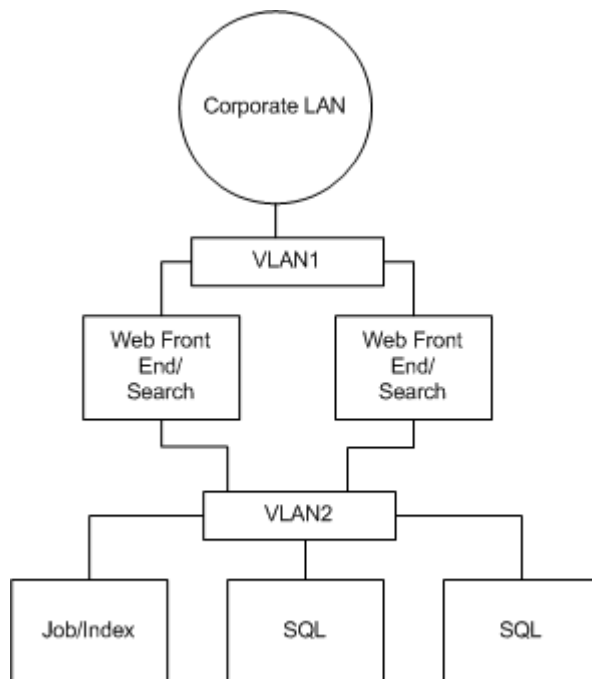


Figure 1.1

Network configuration for a medium server farm

The primary reason for the dedicated public and private networks is to ensure that end-user operations are not impeded by database backups, index propagations, or other large back-end data transfers.

Network Interfaces

Each front-end Web server is equipped with at least two network interface cards (NICs). Figure 1.1 shows how each interface card interacts with the network topology.

The Ethernet switches should be, at minimum, 100-megabit, full-duplex switches. If you anticipate growth beyond the baseline, five server, high-availability server configuration, you should use gigabit (1000base-T) switches and NICs. Consider also that as the server farm scales, the needs of the VLAN2 network may grow beyond the capacity of a 4-port or 8-port switch. By selecting a modular network switch capable of supporting at least twelve full-duplex ports, you will be able to support continued growth.

It is vitally important that there be a domain controller on each of these subnets. In the case of the Microsoft Solution Accelerator for Intranets, the domain controller also serves as the default gateway for each VLAN.

DNS

Each front-end Web server should have access to the corporate name server. During the construction and configuration of the servers, you should have the ability to create unique names for the corporate portal, as well as each divisional portal and site collection. Make sure to reserve meaningful names for these portals.

Active Directory

This architecture assumes that a Windows 2000 Server or Windows Server 2003 Active Directory domain controller is active and available to the runtime environment. Because this solution targets intranets, which are typically hosted inside the corporate LAN alongside domain controllers, a requirement for Active Directory access is not described. SQL Server, SharePoint Portal Server 2003, and Windows SharePoint Services require both local and domain accounts for proper operations.

Physical and Logical Architecture

Physical Architecture

- Two load-balanced front-end Web servers and search servers
- One index server
- Two SQL Server 2005 servers (Clustered)

Table 3.1 is a worksheet that lists the servers in the server farm.

Table 3.1 Sample Server Worksheet

Server role	Server name	Notes
Front-end Web	SC-02	Hardware load balanced
Front-end Web	SC-03	Hardware load balanced
Central Admin	SC-01	
Search Server	SC-04	
Staging/Backup	SC-05	

SQL Server 2005	SC-SQL01	Two servers are required, Clustered Environment
-----------------	----------	--

Logical Structure

The requirements for this deployment are as follows:

- A corporate portal for all employees of Sunshine Connections that provides search, notifications, alerts, audiences, profiles and My Sites.
- A team site collection for each department within Eq, starting with the IT Department.

Software requirements Web and Application Server

- Microsoft Windows Server 2003 R2 (Standard, Enterprise, Datacenter, or Web Edition) with Service Pack 1 (SP1)
- Microsoft .Net Framework 2.0
- Microsoft .Net Framework 3.0
- The Web server and application server computers must be configured as Web servers running Microsoft Internet Information Services (IIS) in IIS 6.0 worker process isolation mode.
- Each of the servers must be using the NTFS file system. Windows Server 2003 includes a conversion utility (Convert.exe) that you can use to convert an existing file allocation table (FAT) volume to NTFS without losing data if necessary.

Security Account Requirements

To install SharePoint Server 2007 in a server farm environment, at-least 2 accounts are required:

- A. A user account that you can use to install Office SharePoint Server 2007 and run the SharePoint Products and Technologies Configuration Wizard. This account must be:
- A domain user account.
 - To A member of the Administrators group on each of your front-end servers.
 - A member of the SQL Server Logins, which grants login access to your SQL Server instance.
 - A member of the SQL Server Database Creator server role, which grants permission to create and alter databases.
 - A member of the SQL Server Security Administrators server role, which grants permission to manage server logins.

- B. A unique domain user account that you can specify as the Office SharePoint Server 2007 service account. This user account is used to access your SharePoint configuration database. It also acts as the application pool identity for the SharePoint Central Administration application pool and it is the account under which the Windows SharePoint Services Timer service runs. The SharePoint Products and Technologies Configuration Wizard adds this account to the SQL Server Logins, the SQL Server Database Creator server role, and the SQL Server Security Administrators server role. It is recommended that you follow the principle of least privilege and do not make this user account a member of any particular security group on your front-end servers or your back-end servers.

Note - Best practice: Always Install MOSS using a domain account that is a local administrator on every machine in the farm. Never install MOSS as a local administrator on the SharePoint server. The Moss install will automatically makes the installation account the administrator for shared services. Using a local account can cause you to be denied access to the farm latter.

After preparing your database and the servers in your farm, run Setup and then run the SharePoint Products and Technologies Configuration Wizard on all your farm servers. Do this on all farm servers before going on to create a Shared Services Provider (SSP).

It is recommended that you run Setup on all the servers that will be in the farm before you configure the farm. You can add servers to the farm at this point, or after you have created and configured an SSP. You can add servers after you have created and configured an SSP to add redundancy, such as additional load-balanced Web servers or additional query servers. It is recommended that you run Setup and the configuration wizard on all your application servers before you create and configure the SSP.

Planning DNS Names

In this prescribed architecture, Sunshine Connections will need two DNS names, one for the portal and one for the associated team sites. For example, for the Sunshine Connections portal, the DNS name for the portal is *DocNet* and the DNS name for the team site collection is *DocNetteams*.

The two DNS names point to the same IP address for the front-end Web server cluster. If the team sites are ever moved to a separate server farm, the team site DNS entry points to the IP address of the separate server farm.

Planning Backup and Restore

Disk Space Requirements for Backups

The amount of disk space required for backup operations depends on three main factors:

- The number of users in your customer's organization and the amount of data and content that is generated and stored in the solution.
- The amount of content stored in document libraries, My Sites, Windows SharePoint Services site collections, lists, and discussions (otherwise known as the *site corpus*).

- Specific backup processes and procedures, such as how long backups are kept on the disk before they are moved to a tape drive.

You should assume the amount of content will increase over time as usage increases.

Consider the following factors when estimating the amount of disk space required for backing up the databases:

- The number of My Sites and the amount of content in the My Sites
- The number of site collections, subsites, and all the content in them
- The number of documents in the document libraries
- The number of lists, news items, and discussions
- The number of alerts
- Windows SharePoint Services full text catalogs
- SharePoint Portal Server content indexes
- Customized site and list templates
- Web Parts
- Custom assemblies
- User profiles
- Gatherer logs
- Usage analysis data
- Audiences

You can help control the amount of disk space used by setting quotas. Quotas can be set on each virtual server for each site collection and My Site. Quotas can help your customer monitor storage space, and locks can stop sites from exceeding their storage quota limits.

You can create multiple sets of quotas, called quota templates, and use them in different areas of the server farm. A quota template can be specified as the default quota template for a virtual server. Any new site collection created under that virtual server automatically uses the values in the default quota template, or your customer can assign a different quota to a specific site collection. For more information, see "Managing Site Quotas and Locks" in the *SharePoint Portal Server Administrators Guide*, accessible from the **Start** menu.

Backup Strategies

The following topics provide an overview of recommended backup methodologies and the estimated amount of disk space required to support backup operations. For more information, see Chapter 2 of the Operations Guide.

SharePoint Portal Server Database Backups

The SharePoint Portal Server databases and the Windows SharePoint Services databases are backed up by using the native SharePoint Portal Server Backup and Restore Utility. The utility backs up SQL Server 2005 databases for each portal. It creates an XML-based manifest file that lists all of the individual backup files for SharePoint Portal Server and Windows SharePoint Services databases.

The size of the SharePoint Portal Server databases depends on your customer's business requirements and usage.

My Sites Backups

By default, My Sites content is contained in the site database of the corporate portal. The database for the portal is backed up on a regular basis and the individual My Sites are backed up in a separate process using the SharePoint Portal Server stsadm.exe command line utility. This process backs up the same data twice, but it allows for backing up and restoring individual My Sites as opposed to restoring the entire site database.

The amount of disk space to back up My Sites depends on the following factors:

- The number of My Sites
- The amount of content in each My Site
- The amount of time that My Site backups are kept on disk before they are moved to a tape drive

You can estimate the maximum amount of disk space required for backing up the individual My Sites by using the following calculation:

The number of My Sites × disk quota = maximum disk space required

For example, if you have 1000 users, each with a My Site and 100 MB storage limit set on each My Site, the maximum amount of disk space required for the My Site backup is approximately 100 GB. This also means that the divisional portal site database, which also contains the My Site data, is equally impacted by the extra 100 GB of data. Keep in mind that most My Sites will not use disk space in amounts approaching the quota. Therefore, the average disk space required for My Sites will be much lower than the quota. Set the quota high enough to accommodate the occasional large My Site. Setting the quota too low could result in support calls from users needing more space.

Windows SharePoint Services Site Backups

The Windows SharePoint Services site collection databases are backed up by the SharePoint Portal Backup and Restore Utility in the same manner as the SharePoint Portal Server databases. For more information, see the preceding section "SharePoint Portal Server Database Backups." In addition, there are alternate prescribed methods for backing up and restoring site collections and My Site data at a more granular level. This means that the backup of some data is duplicated, and this increases the required hard disk space.

The site collections are also backed up individually by using the stsadm.exe command line utility, so that individual site collections can be recovered without restoring the entire site collection database. In the Solution Accelerator for Intranets model with the separate Windows SharePoint Services server farm, the site databases cannot be backed up by using the SharePoint Portal Backup and Restore Utility. In the separate Windows SharePoint Services server farm model, only stsadm.exe backups can be performed.

The amount of disk space required to back up each site collection depends on the amount of content within the site collection and how long the backups are kept on the disk before being moved to a tape drive. You can estimate the maximum amount of hard disk space required for site collection backups by using the following formula:

The number of site collections × disk quota = maximum disk space required

For example, if you have 10 site collections with the quotas set for 100 MB, you need 1000 MB of disk space for backing up site collections. This also means that the Windows SharePoint Services databases containing the site collections are impacted by 1000 MB.

Front-End Web Server Backups

Front-end Web server data includes custom assemblies, configuration files, add-in software, custom templates, the IIS 6.0 metabase, and the Inetpub directory. The size of a front-end Web server backup depends on your customer and how much customization is done on the front-end Web servers. An estimate based on the medium server farm solution consisting of two front-end Web servers is approximately 50 MB for each server. For more information, see the "Backing up and Restoring Front-End Web Servers" section in Chapter 2 of the Operations Guide.

You can estimate the maximum amount of hard disk space required for front-end Web server backups by using the following formula:

The number of front-end Web servers × 50 MB = estimated disk space required

For example, if you have 2 front-end Web servers at 50 MB, each, you need 100 MB of disk space to back up the Web servers.

You must also plan adequate space for backing up the IIS log files. You can estimate the space required by using the following formula:

The number of days of logging × daily log file size = estimated disk space required

For example, if you want to preserve 30 days of logs with an average daily log file size of 1 GB, you need 30 GB.

Images

You must plan for the size of the disk images used for imaging servers on a server farm. The size depends on the backup software used to create the images and the number of images that you need to create. Some image software utilities can compress more than others, but not as quickly. Allow about 8 GB of storage space per image.

Database Configurations

Database File Configurations

You must understand the trade-offs associated with the various SQL Server Recovery models and AutoGrow options. Leaving the database set to Simple recovery mode will restrict your disaster recovery options to only Full backups. So, if you are only doing backups once a week on Sunday and you have a failure on Weds, you will only be able to restore the database its state on Sunday. The transactions between Sunday and Weds will be lost because the log is being truncated immediately after the transaction is committed to the database. Full portal backup is the only method supported by SharePoint Backup and Restore utility.

General Database Sizing Guidelines

Set the database file to fixed size (estimate how much space you will need for all your data), set AutoGrow to 50%. If you have a small initial database size and SQL Server has to AutoGrow by 10% all the time just to keep up with the data being inserted, you will run into some serious fragmentation (and performance) problems.

Set log file to 20-30% of data file, don't AutoGrow. If you are implementing Simple recovery or Full recovery with an incremental backup daily, you should have no problem with your log file setup this way.

Backup Plan for Sunshine Connections

SharePoint Backup Utilities

There are two important backup utilities delivered with the SharePoint Portal Server software.

Tool	Purpose
Stsadm command-line tool	Back up the server farm, Web applications, and any or all of the content databases in your server farm
SharePoint Central Administration Web site	This tool allows you to perform backups and restorations from the user interface

A complete Backup Plan may use both utilities depending on the desired result. The recommended approach for Sunshine Connections includes making a full backup of the portal, backing up each site collection, and then moving the backup files to tape.

Full Portal Backup

For complete disaster recovery, the entire portal should be backed up using the stsadmin.exe tool. For the Pilot, a shared resource was created to store all of the backup files for the farm. It is on the staging server and is called \Backups. The command line below would execute a complete portal backup and place the backup files into \\{server_name}\Backups:

```
STSADM.exe -o backup -directory \\{Server}\{Share_Name}\ -backupmethod full
```

The command line above should be included in a script file located in the root of the administrative server..

Move to Tape

After backup up the portal and the site collections to a shared resource, the backup files should be moved to tape for safe keeping. This allows you to overwrite the backup files each night to save disk space while keeping a history of portal backups equal to the degree of your tape rotation.

Full Portal Restore or Single Site Collection Restore

Can be accomplished by restored by using the SharePoint Central Administration Web Site or the stsadm command-line tool.