

# Microsoft® SharePoint® Server 2010

---

An Introduction

**Dell™ Solutions Engineering**

*Ravikanth Chaganti and Kevin Guinn*

May 2010



## Executive Summary

Microsoft® SharePoint® Server 2010 is the latest version of Microsoft's platform for web-based collaboration, portals, and content management. The product has undergone several changes compared to its predecessor, Microsoft Office SharePoint Server (MOSS) 2007. A wide array of out-of-the-box functionality is available, but SharePoint also provides a programming platform that enables significant customization so that its services can be integrated into existing business processes.

This paper provides an overview of common business use cases for which SharePoint is a good fit, introduces the services and functionality within SharePoint Server 2010, highlights key changes from MOSS 2007, and helps readers understand the terminology and concepts needed for planning or deploying a SharePoint 2010 farm. The new service application architecture is outlined in detail, and key services and related products are cataloged. Finally, the paper outlines some considerations for choosing a strategy to upgrade from MOSS 2007 to SharePoint Server 2010.

Three companion papers, available from [www.dell.com/sharepoint](http://www.dell.com/sharepoint), provide specific reference architectures and additional details related to designing and building small, medium, or large SharePoint 2010 farms. In conjunction with the appropriate reference architecture paper, this paper can help administrators determine how to implement SharePoint Server 2010 in their environment.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

*Dell*, the *DELL* logo, and the *DELL* badge, *PowerEdge*, *EqualLogic*, *PowerConnect*, and *PowerVault* are trademarks of Dell Inc. *Microsoft*, *Windows*, *Windows Server*, *SharePoint*, *SQL Server*, *Excel*, *Outlook*, *PowerPoint*, *Visio*, *InfoPath*, *OneNote*, and *Active Directory* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *Intel*, *Xeon*, *Pentium*, and *Celeron* are registered trademarks and *Core* is a trademark of Intel Corporation in the U.S. and other countries.

# Table of Contents

Executive Summary .....	ii
Introduction .....	2
Microsoft Office and SharePoint Client Applications .....	2
Integrating with Other Microsoft Server Applications .....	2
Server Hardware Architecture.....	3
Common SharePoint Use Cases.....	4
Portal.....	4
Social Media .....	4
Collaboration.....	4
Enterprise Search.....	5
FAST Search Server 2010 for SharePoint .....	5
Business Process Enablement.....	5
Publishing / Content Management.....	5
Business Intelligence.....	5
SharePoint Containment Hierarchy.....	6
Key Enhancements in SharePoint 2010 .....	8
SharePoint 2010 Deployment Changes .....	8
Windows PowerShell support .....	8
Central Administration Improvements .....	9
Backup and Restore .....	9
SharePoint 2010 Health and Monitoring .....	10
Unified Logging Service (ULS).....	11
Logging Database .....	11
Remote BLOB Storage.....	11
Shared Service Applications .....	12
Search Service Architectural Changes.....	14
Managed Metadata Service.....	15
Claims-based Authentication .....	15
Support for Additional Web Browsers and Mobile Devices.....	16
Closely-Related Products .....	16
Upgrading to SharePoint Server 2010 .....	17
Summary.....	18

## Introduction

The term “SharePoint” is broadly used to describe a family of products and technologies that interact with Microsoft SQL Server® and Internet Information Server (IIS) to provide a Web-based engine and a platform for deploying a wide range of business services. The most common solutions deployed using this platform are collaborative sites, content management systems, and Web portals. SharePoint solutions are usually deployed in a farm environment that provides scalability by distributing database, application, and presentation roles across different groups of servers.

There are many out-of-the-box templates for common types of sites and libraries, but SharePoint also serves as a development platform and allows for significant amounts of customization. This paper provides an overview of SharePoint Server 2010, and outlines several key enhancements over its predecessor, Microsoft Office SharePoint Server (MOSS) 2007.

SharePoint Foundation 2010 provides the core engine, services, and facilities for creating and using templates. This core functionality employs Microsoft programming and Web platforms, such as ASP.NET; it can be enhanced and extended by developing custom templates, WebParts, and code. SharePoint Server 2010 significantly extends the functionality and capability beyond the core services offered by SharePoint Foundation. There are different editions of SharePoint Server 2010; some of the service applications and functionality discussed in this paper will only be available in Enterprise Edition.

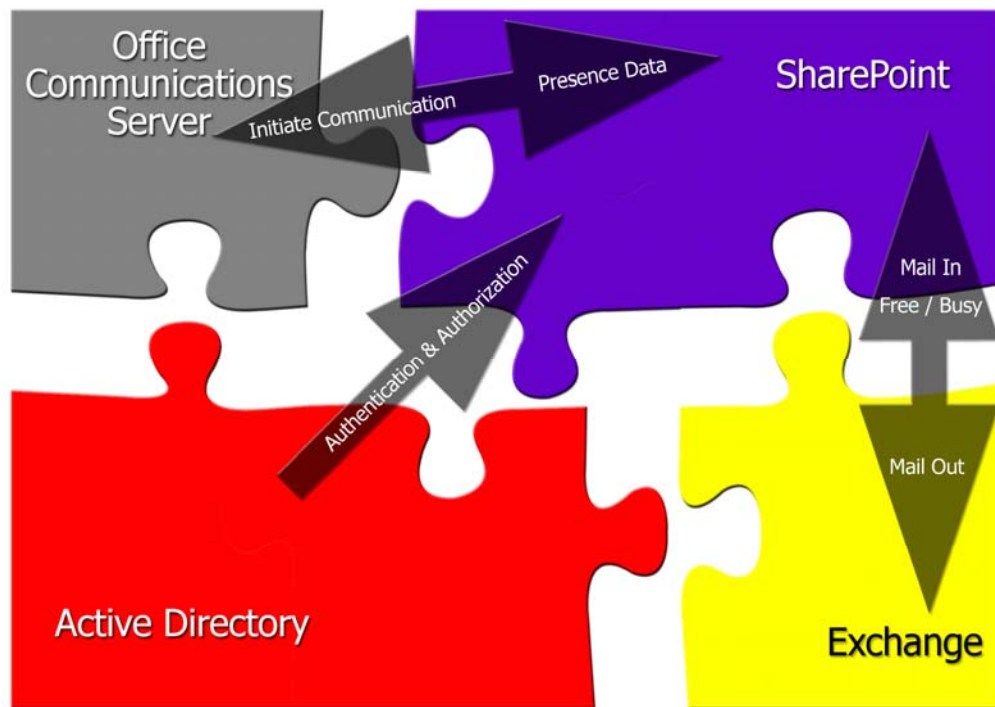
## Microsoft Office and SharePoint Client Applications

SharePoint offers integration and ease-of-use benefits when used in conjunction with Microsoft Office client applications. For example, documents stored in a SharePoint library can be directly opened from Microsoft Word®, PowerPoint®, or Excel®. Also, from within Microsoft Outlook®, users can subscribe to and display list items from a SharePoint site or RSS feeds provided by a SharePoint-powered blog. Users can also create blog entries directly from Word. This integration makes data stored in a SharePoint infrastructure more accessible to end users.

SharePoint Designer 2010 is a client application that simplifies the creation of custom sites, templates, applications and WebParts. SharePoint Workspace 2010 enables synchronization with SharePoint document libraries and provides the ability to access, view, and edit documents when working offline.

## Integrating with Other Microsoft Server Applications

Installing a server farm for SharePoint Server 2010 requires the inclusion of certain infrastructure and services to fully exploit all SharePoint features and functionality. For example, Active Directory® (AD) is a pre-requisite, because it provides authentication and authorization among the servers in a farm and can be used to import user profile information from AD into SharePoint. If your farm is heavily used, adding additional directory servers may be necessary to handle the authentication traffic.



*Figure 1: Integrating SharePoint into an Enterprise Infrastructure*

Exchange Server can be used for the mail-out and mail-in connections for SharePoint. These connections enable features such as e-mail notification of changes within a SharePoint collaboration site and the ability to create blog entries by sending an e-mail. In addition, if Outlook Web Access (OWA) is configured in the Exchange environment, then data stored in Exchange - such as a user or group calendar, task list, or e-mail items - can be directly displayed within a page on a SharePoint site.

Similarly, Office Communications Server (OCS) enables user presence information to be displayed on SharePoint pages. For example, the familiar “gumball” from Office Communicator is displayed next to user names on the SharePoint page, providing the ability to view free/busy data, initiate instant messaging conversations, send e-mail, or even initiate a call with another user. To fully exploit this integrated functionality, Exchange Server and Office Communications Server must be properly configured and a SharePoint end-user must be logged in to both Office Communicator and Outlook.

## Server Hardware Architecture

SharePoint Server 2010 is exclusively offered for the 64-bit extended (x64) architecture. All host servers in the farm, and the Windows Server operating systems installed thereon, must use the x64 architecture. Compared to the legacy 32-bit “x86” architecture, x64 enables SharePoint to more easily access, and more efficiently use, host server memory - particularly when the host is configured with more than four gigabytes (4 GB) of RAM. Note that an x64 version of Windows Server 2008 with SP2 or Windows Server 2008 R2 is also required. For a complete list of hardware and software requirements for SharePoint 2010, see [http://technet.microsoft.com/en-us/library/cc262485\(office.14\).aspx](http://technet.microsoft.com/en-us/library/cc262485(office.14).aspx).

## Common SharePoint Use Cases

SharePoint can be used as the basis for web sites that address many common business needs. A farm may be designed to address one or more of these scenarios, each of which may be exposed as public internet sites, private intranet sites, or extranet sites that allow interaction between employees and authorized outside parties.

### Portal

Web portals consolidate links and information from many different sources, and make it accessible from a central location. They serve as a starting place that allows site users to locate the information or services that they require. An intranet portal, for example, may share company-wide news and information and link to sites that are managed by individual departments. It may also provide access to enterprise applications, reporting systems, business data, or other online resources.

### Social Media

Social media generally refers to a set of technologies that enable user-generated content and community participation. Key examples that are enabled by SharePoint include blogs, wikis, and content tagging and ranking. Blogs allow experts to share information and promote discussion with and commentary from their audience. Wikis provide allow groups to generate content through in an open collaborative environment. Tagging enables users to identify content that is meaningful to them and facilitates search relevance, and ranking provides quick feedback about its usefulness to the audience.



*Figure 2: Social Tagging*

In addition to these technologies, SharePoint Server 2010 provides a facility to search for people based on their roles and areas of expertise. Users can self-identify areas of expertise, or can be identified and tagged by their peers and colleagues. SharePoint can also be configured to suggest colleagues and connections to users. By importing user profile data from Active Directory, it is also possible to customize people search to allow navigation of complex organizational structures. These technologies are intended to help people identify experts and make connections so that they can work together more efficiently.

### Collaboration

In most organizations, people work in teams with shared goals and responsibilities. While social media provides one avenue for collaboration, it is not necessarily well-suited for every type of team activity. SharePoint provides several features that allow teams to share ideas and work together. Document libraries provide a repository for the team and, through a check-out mechanism and version control, help manage changes and revisions. Document and meeting workspaces allow team members to contribute ideas and foster discussion in a structured environment. Announcements and discussions foster awareness and participation among the team. Shared calendars and task lists make it easy to track team objectives and project milestones.

## Enterprise Search

SharePoint allows users to search for content within the farm and from external sources, which makes it possible for users to quickly locate the information that they need. Content and documents are indexed, so that key words and important metadata can be used as search criteria. For supported document types, the index is built using not only the document title and properties, but also a full-text scan of their contents. Crawler rules can be customized and optimized to help improve search relevance, therefore making it easier for users to access the data they desire.

SharePoint Server 2010 Search allows indexing of metadata such as author, department, creation and modification dates, and many other document properties. The search indexing functionality can be extended with filters that enable full indexing of PDF files, and other content types that are not natively provided. One key enhancement is the introduction of faceted search filters, also known as refinements, to the default "SharePoint Search Center" interface.

### *FAST Search Server 2010 for SharePoint*

FAST Search Server 2010 for SharePoint is a separate product that can be deployed to offer additional enterprise search capabilities. This technology provides unique natural-language analysis and content processing capabilities and provides its own FAST Query Language (FQL) that can be used to develop search-based applications. FAST enables extraction of metadata from documents and other unstructured data within a SharePoint farm, and can be integrated with the Business Data Connectivity services to allow searching among structured data stored in external databases (such as those that power many business applications). FAST search provides a toolkit for faceted search and results refinement operators. It extends best bets, which are administrator-defined, managed search results by providing a visual rendering (picture or video) on the results page.

FAST provides a development platform for search-based applications and can enable customization of results sets and refinement facets on a per-user or per-group basis. The FAST search service uses a scale-out model that can enable it to be successfully employed in scenarios that feature very large amounts of content or are heavily utilized (e.g. in terms of the number of search queries per second).

## Business Process Enablement

Business processes often play a key role in determining how information is approved, published, and accessed. SharePoint provides workflows, which can be used to enable approval, review, publishing, and information retention processes. When workflows are configured, extended, or customized correctly, they can help enforce policies that underly these business processes.

Visio® services and InfoPath® services are available for SharePoint. Visio services can be used to help visually design custom workflows, and InfoPath services enable incorporation of electronic forms to drive business processes.

## Publishing / Content Management

By combining business process workflows, data retention policies, and the document library facilities within SharePoint, it is possible to build a content management system that works with a company's information management and retrieval policies. Such systems also benefit from version control, search, and reporting functions offered by SharePoint.

## Business Intelligence

Other services can be combined to enable SharePoint to be used as a platform for Business Intelligence, giving business stakeholders access to key information and analytical tools to make

informed decisions. PerformancePoint services enable reporting and creation of key performance indicator (KPI) dashboards for content and data that is accessible by SharePoint. PowerPivot for SharePoint can be used in conjunction with PowerPivot for Excel to enable users to perform complex, faceted, mining and analysis of data across multiple sources. The SharePoint Business Data Connectivity service enables integration of external content, such as a SQL Server or Oracle database, an Enterprise Resource Planning (ERP) system, or a Customer Relationship Management (CRM) system. Once connections have been established, SharePoint search and reporting functions can be extended to include the external data.

## SharePoint Containment Hierarchy

When designing and maintaining a SharePoint solution, it is important to understand the various levels at which information is organized and contained. The containers within a SharePoint infrastructure are outlined in Table 1. The most granular individual items are located at the bottom of the table, and the level of aggregation increases as you progress to the top of the table. These containers provide physical and logical boundaries to consider while designing and deploying a SharePoint infrastructure.

*Table 1: SharePoint Containment Hierarchy*

Container	Description
SharePoint Farm	A set of servers that collectively provide the databases, applications, and Web services that comprise a SharePoint solution.
SharePoint Servers	Individual servers that run the operating system and application software required to perform roles or provide services for the SharePoint farm. Examples include a Web front-end server, an application server, and a database server.
IIS Application Pool	A container that is configured within Internet Information Server (IIS) to constrain a defined set of content to operate within a defined set of system processes. Application Pools provide logical barriers that protect against the threat of a compromised site being used as a vector to attack other sites hosted on the same Web server.
IIS Web Application	An IIS Web site with a unique domain name that is created and used by SharePoint products. Three Web applications must be configured: Central Administration, Shared Services Provider (SSP), and content. Additional Web applications may be useful for providing content isolation or for establishing distinct management or SLA boundaries within the farm.
SharePoint Database	Individual SQL Server databases that are used to store information about or data from within a SharePoint farm. The core databases used by SharePoint are Configuration, Administration, SSP, Search, and Content. Depending on its architecture and needs, a farm may feature multiple SSP, Search, and Content databases.
Site Collection	A set of sites that feature the same owners and administrative settings (such as content types or quotas). A site collection features a top-level Web site and may also contain several sub-sites. Generally, all of the sites within a site collection share a common navigational design. One content database can host multiple site collections, but data from a given site collection must reside in the same content database. Similarly, one or more site collections may be configured within the same Web application.
Site	A set of Web pages stored within a site collection that deliver common features or content to users. Sites may be structured, such as a top-level portal site, or may be ad hoc, such as team sites for collaboration. MOSS provides templates for several types of sites, such as blogs, wikis, team sites, and portals.



Container	Description
List	A means of collecting, storing, and organizing data within a site. Some common examples include document collections, calendars, and tasks.
Item	An individual data object within a list. Some common examples include document or image files, contacts, and calendar entries.

Figure 2 represents the relationship between a site collection and the sites, lists, and items of which it is comprised. The items are organized into lists, which are - in turn - organized into sites within the site collection. This site collection is configured within a single Web application, and all of its data and items are stored in a content database. A SharePoint farm can house many separate site collections, and its information architecture plays a critical role in determining the relationships between the Web applications and other higher-level entities outlined in Table 1.



*Figure 3: Relationships between Entities within a Site Collection*

## Key Enhancements in SharePoint 2010

Microsoft SharePoint 2010 builds upon the legacy of MOSS 2007 to provide a rich set of improvements in collaboration, information sharing, and document management. SharePoint 2010 adds several new features to enable these improvements in the overall product. It is important to understand these features and how you can make use of them within your SharePoint 2010 deployment.

### SharePoint 2010 Deployment Changes

To take advantage of the performance improvements in 64-bit hardware platform and make it a standard going forward, SharePoint 2010 is available only for 64-bit extended architecture. Because of this, you will require 64-bit hardware platform, operating system and SQL server software to deploy SharePoint 2010 in your organization.

There are quite a few prerequisite components<sup>1</sup> you need to install before deploying SharePoint 2010. To facilitate this process, SharePoint deployment now includes a prerequisites installer that can be used to install the necessary components such as Internet Information Server, the .Net framework, etc. The prerequisite installer tries to download the required components from the Internet and install them on the local computer. If the farm servers do not have Internet access, manually download all the required components to a local folder and direct the prerequisite installer to install these components from the specified local path. The latter method is also helpful when installing SharePoint onto several farm servers.

Another addition to the deployment process is the farm passphrase. Farm passphrase is a password used to encrypt all farm communications. You will be prompted to create this passphrase when deploying the first SharePoint server in your farm. This passphrase is required to add or remove farm servers. Administrators can also use Active Directory group policies to prevent unauthorized installs of SharePoint software on network servers.

### Windows PowerShell support

Windows PowerShell is now the preferred command-line interface for SharePoint 2010 administration<sup>2</sup>. STSADM has been retained for backward compatibility with existing scripts and tools, but it is officially deprecated in SharePoint 2010. PowerShell cmdlets can be used to achieve certain advanced management and configuration tasks that cannot be performed using central administration.

After SharePoint 2010 is installed, you can access SharePoint 2010 PowerShell cmdlets by accessing the SharePoint 2010 management console. This console, by default, loads Microsoft.SharePoint.PowerShell snap-in and provides access to all the SharePoint 2010 cmdlets. To get a list of all cmdlets, you can use the built-in Get-Command cmdlet. Type the following at the SharePoint 2010 management console and press Enter:

```
Get-Command -Noun SP*
```

---

<sup>1</sup> See "Installing Software Prerequisites" at [http://technet.microsoft.com/en-us/library/cc262485\(office.14\).aspx#section4](http://technet.microsoft.com/en-us/library/cc262485(office.14).aspx#section4)

<sup>2</sup> See "Windows PowerShell for SharePoint Server 2010 Reference" at [http://technet.microsoft.com/en-us/library/ee890108\(office.14\).aspx](http://technet.microsoft.com/en-us/library/ee890108(office.14).aspx)

Administrators can even use Windows PowerShell to script the entire SharePoint 2010 deployment process. While this is not a trivial task, doing so can help ensure that all servers in a farm are configured in a consistent manner.

## Central Administration Improvements

Central administration has been redesigned completely with features like a ribbon interface that is similar to Office 2007 and 2010 products. The ribbon is context-aware, and makes accessing the subtasks within any selected functional area more intuitive. The central administration site looks similar to Windows control panel, where different tasks and activities are grouped into categories. The new home page also provides convenient access to most common tasks by grouping them together into categories such as application management, monitoring, security, backup, and so on.



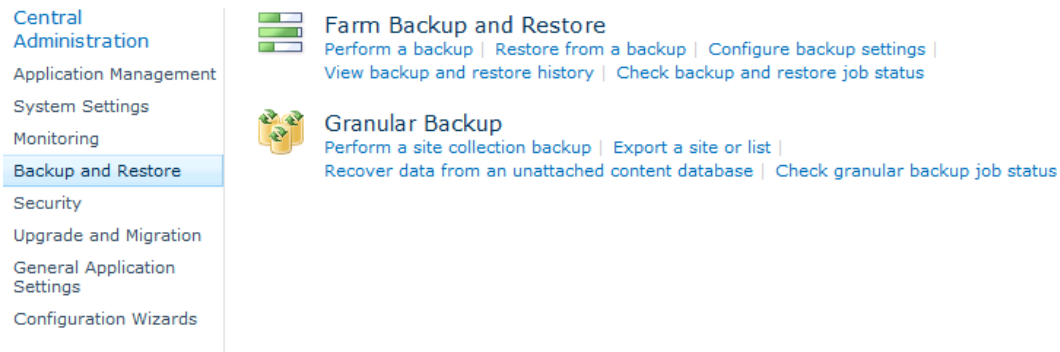
*Figure 4: SharePoint Server 2010 Central Administration*

A new farm configuration wizard interface has been added to provide step-by-step guidance to configure the SharePoint 2010 farm. ISVs and third party developers can add additional wizards to provide similar guidance in configuring other aspects of the SharePoint 2010 product.

## Backup and Restore

Microsoft added a few very important enhancements to SharePoint backup and recovery. SharePoint 2010 provides several levels of granularity for performing backup and restore. This includes farm configuration, site collections, sub sites and lists.

The central administration site provides a simple interface to perform full backup and restore, perform granular backup or configure backup settings. In earlier versions of SharePoint, administrators built a secondary farm to perform granular recovery of SharePoint content. This is no longer required with SharePoint 2010. In SharePoint 2010, administrators can perform content restore from unattached content databases. Using unattached content database feature, you can connect to a content DB restored to any SQL server in the network and use SharePoint 2010 to browse the contents of the database. An administrator will be able to restore content at a very granular level.



*Figure 5: Backup and Restore Options in Central Administration*

Similar to any other administrative function within SharePoint 2010, you can automate or script your backups using Windows PowerShell. Every backup and restore activity that can be performed using central administration can also be done at the command-line using PowerShell. Windows PowerShell also provides a few more additional backup/restore options such as file compression and SQL snapshots which are not available through central administration webpage.

## SharePoint 2010 Health and Monitoring

SharePoint 2010 includes new tools to monitor health and track performance of a SharePoint farm. This includes features such as unified logging, developer dashboard, usage database, SharePoint maintenance engine, and out-of-box usage reports. Most of these improvements are done to provide a more proactive way of monitoring and reporting. Administrators can also use PowerShell cmdlets to monitor the health of a SharePoint farm.

The integrated Health Analyzer identifies common problems and provides prescriptive guidance to help administrators resolve them. The tool includes several default rules for Security, Performance, Configuration, and Availability. Each rule can be independently enabled and have its scope set to determine the servers that are checked. If a rule is enabled, then SharePoint 2010 will periodically check the conditions associated with the rule, based on a defined schedule. Certain rules can even be set to allow SharePoint automatically correct violations.



*Figure 6: Health and Reporting in Central Administration*

Developer dashboard enables an in-depth look at page performance in a SharePoint farm. Though central administration, this feature is available to both developers and administrators. Developers can

use this to measure how the custom code is performing under peak/non-peak load and to optimize the code if required. This feature is turned off by default, but can be enabled using PowerShell.

### *Unified Logging Service (ULS)*

SharePoint 2010 provides a few improvements to the way trace logs are created and maintained. Trace logs contain information about the raised exceptions and regular usage of a SharePoint product. Unless the location has been changed by the administrator, trace logs are stored in the

`%commonProgramFiles%\Microsoft Shared\Web Server Extensions\14\Logs\` directory.

The default retention period for these log files has been extended to 14 days, but administrators can use central administration or PowerShell to change this setting. To limit log file storage requirements, a disk space quota can also be applied.

### *Logging Database*

In SharePoint 2010, the Unified Logging Service also implements a database to store the trace log information. This logging database can be configured to capture page requests, webpart rendering details, performance counters, timer jobs, feature usage, search crawl statistics, etc. This database is used by the web analytics feature to show the usage patterns of a SharePoint 2010 deployment. If many counters are enabled, this database can grow in to very large size in the production environment; it is recommended to allocate enough database storage to store and archive logs. The size of this database will depend on the number of farm servers, the number of configured counters, and the log retention policies.

The logging database stores event information, but does not provide a way to directly perform an action based on triggers. System Center Management Pack for ULS can be used to surface the events in SCOM. The logging database is also the only SharePoint 2010 database that allows direct interaction, DBAs and 3<sup>rd</sup>-party products are allowed to read from and write to this database. This level of access can make it possible to monitor and react to events.

### *Event Logging Flood Protection*

Event flooding can be controlled by using the event throttling feature of ULS. Using this feature, only a few occurrences of the event are logged, thereby saving log space and providing a cleaner look at the log files. You can use central administration monitoring options to configure event logging flood protection.

## **Remote BLOB Storage**

By default, a SharePoint farm uses content databases on SQL Server as its only data source. Farm content generally includes documents and pictures, which can each be on the order of several megabytes in size. Files and other unstructured data have historically been stored as binary large objects (BLOBs) within the database.

While there have been many improvements in the way that BLOBs are handled within a SQL Server database<sup>3</sup>, there are limitations. The VARBINARY(MAX) data type is used for native BLOB storage, and can accommodate files up to two gigabytes (2147483648 bytes) in size. In rare circumstances, this may not be adequate for all of the files that will be stored in the farm. Also, the performance of writing BLOBs to and retrieving BLOBs from the database is generally not as good as the performance offered by file systems, which can degrade users' perception of a SharePoint solution. To address this, a

---

<sup>3</sup> See "Managing Unstructured Data with SQL Server 2008" at <http://www.microsoft.com/sqlserver/2008/en/us/wp-sql-2008-manage-unstructured.aspx>

Remote BLOB Storage (RBS) feature pack is available for SQL Server 2008 and SQL Server 2008 R2; it provides libraries and APIs that allow SQL Server to interact with BLOBs stored outside of the database.

SQL Server 2008 and SQL Server 2008 R2 provide a data type called FILESTREAM that allows BLOBs to be stored within an NTFS volume, but managed by the database. This feature cannot be directly used by SharePoint<sup>4</sup>, but there is an RBS provider for FILESTREAM that can be used with local SharePoint databases (such as those often employed by SharePoint Foundation 2010); this local database requirement makes the FILESTREAM provider for RBS unsuitable for most farm-based architectures.

For farm deployments with separate database servers, RBS providers other than FILESTREAM may still be used. For example, a provider for a dedicated external store, such as an EMC<sup>2</sup> Centerra system, that makes use of the SQL Server RBS API calls may be employed. These RBS providers allow third-party products to leverage RBS to store and retrieve files that reside outside of the database.

Using Remote BLOB Storage requires configuration of both SQL Server 2008 (or SQL Server 2008 R2) and SharePoint Server 2010. These systems are expected to offer higher performance for reading and writing large files, but will generally require special attention when administering sites or performing backup/restore operations.

### Shared Service Applications

Shared service applications in SharePoint 2010 provide a new service application architecture that is more flexible than shared service providers (SSP) in MOSS 2007. SharePoint 2010 provides an extensible way to create custom service applications and deploy them in a SharePoint farm. Extensibility allows custom and third-party service applications to leverage the same APIs and provide a similar administrative experience as out-of-box service applications.

In MOSS 2007, the concept of Shared Service providers was introduced to enable sharing of several services across multiple server farms, web applications and site collections. The major drawback of SSP architecture was every web application in a SSP bore the overhead of all services even when the web application did not consume any of those services. Also, there was no way to configure or customize a particular service for a web application without creating a new SSP. Unlike MOSS 2007, SharePoint 2010 farms can both publish and consume services. The new shared service application architecture in SharePoint 2010 eliminates these limitations.

With the new service application architecture, administrators can configure a set of services to be associated with a web application and a completely different set of services associated with a second web application. This is done by creating custom proxy groups in SharePoint 2010. Hence the web applications can be configured to use only the services that are needed by the application. This architecture provides flexibility to configure these services differently for different web application needs. Web applications are also allowed to consume services from more than one group.

It is also possible to share a single service across multiple sites and farms within SharePoint 2010. You can also run the services that need to be shared across farms in a central SharePoint farm and access these services from other farms. However, this flexibility is not available with all service applications. Service applications can be shared between farms in two methods. One is sharing the service application only and second method is to share both application and data. The latter is the default

---

<sup>4</sup> See Overview of Remote BLOB Storage (SharePoint Server 2010) at [http://technet.microsoft.com/en-us/library/ee748649\(office.14\).aspx](http://technet.microsoft.com/en-us/library/ee748649(office.14).aspx)

sharing behavior in SharePoint 2010. However, service applications can be configured to isolate data by using partitions. Partitioning allows multiple organizations to be hosted on the same SharePoint farm, with their data isolated from one another.

SharePoint 2010 adds a great number of new out-of-box service applications along with those available in MOSS 2007. The following table lists different service applications available in SharePoint 2010.

*Table 2: SharePoint 2010 Service Applications*

Service Application	Description	Can be Shared across farms?	Enterprise Edition Only
Search service	Crawls content, performs indexing and responds to search queries.	Yes	
Managed Meta Data	Provides access to the managed taxonomy, keywords, social tags and content type publishing across site collections.	Yes	
Business Data Connectivity	Enables access to data stored in external sources, such as Line-of-Business applications or other databases.	Yes	
Access Services	Allows interacting with Access databases from a browser.	No	X
Excel Services	Allows interacting with Excel spreadsheets in a browser.	No	X
User Profile (People)	Enables creation and sharing of user profiles, organization profiles and social tags. This service is required for MySites and depends on Managed meta data service. Enterprise search service can access the people service data while serving search content.	Yes	
Secure Store Service	Provides capability to store data securely and associate it to a specific identity or group of identities. This also allows for single sign-on functionality.	Yes	
State Service	Temporarily stores user session data (therefore maintaining the state of a user session) for SharePoint Server components.	Yes	
Usage and Health Data Collection	Collects usage and health data across the farm, and provides reporting for this data.	No	
Visio Service	Enables interaction with Visio diagrams in a web browser.	No	X
Web Analytics Service	Enable rich insights in to web usage patterns by processing and analyzing web analytics data.	Yes	
Word Automation Service	Provides framework for performing automated document conversions.	No	
PerformancePoint Service	Provides flexible, easy-to-use tools for building dashboards, scorecards, reports, and key performance indicators (KPIs), and hence enables everyone across an organization make informed business decisions that align with companywide objectives and strategy. Scorecards, dashboards, and KPIs help drive accountability.	No	X

### *Search Service Architectural Changes*

SharePoint 2010 adds several improvements to the out of the box search service. There are also architecture changes that provide a more scalable and highly available search service. In addition to the SharePoint 2010 search component, Microsoft offers FAST Search Server 2010 for SharePoint as a separate product.

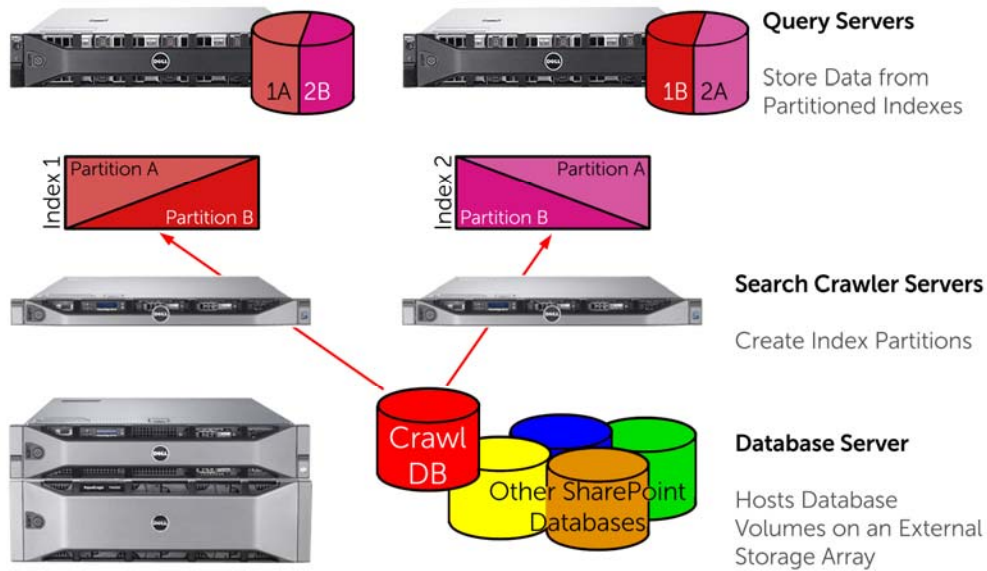
From an end-user perspective, the search service enables great enhancements to out of box user experience. SharePoint 2010 provides refinements, which enable users to filter the search results. These refinements include social tags, taxonomy, and metadata.

If FAST search is available, users can also suggest that content can be promoted or demoted in the results based on relevance and usefulness. Also, when metadata or social tags match the content of a document, the document will be promoted and rank higher in search results.

The new search service application architecture in SharePoint 2010 includes greater redundancy. The new design provides flexibility and allows the query and crawler roles to be scaled-out separately on an as-needed basis. Search crawlers are now stateless, they do not store a copy of the index. The index does, however, still propagate and is stored locally on the query servers.

As shown in Figure 7, the search index on the query server can now be divided into index partitions. Index partitions contain only a part of the complete search index. So, in the following example, two index servers are configured to crawl the SharePoint content and place index partitions on two separate query servers. In this case, redundancy for the index content is provided by spreading the index partitions from each crawler onto multiple query servers. When a query server receives a search query, it forwards the request to all peer query servers to process. The server processing the query then merges the results from across all index partitions to provide search results to the end user. Hence, this new scale-out architecture can help balance the search query processing load and improve query performance. Depending on how partitions are configured and assigned, it may be possible for the loss of a query server to cause users to receive partial results.





*Figure 7: Partitioned Search Indexes*

Administrators can also provision multiple crawl databases to improve the crawl performance. These crawl databases store information about hosts and content that is already indexed or pending index processing. When there are multiple crawl databases, the search administration database is used to keep track of content which has been crawled. There is native support for SQL mirroring for the crawl databases; crawl servers employ a connection string that includes the primary and mirror databases. If the database mirror is configured in high-safety mode and uses a witness, its failover capability is extended to the crawl database.

In SharePoint 2010, the search service application can be shared across multiple farms. This means that you can create a central search farm and publish search services to other farms that require search services.

### *Managed Metadata Service*

The Managed Metadata Service application in SharePoint 2010 provides organizations a method to share metadata taxonomies and terms across SharePoint site collections and web applications. This also enables sharing of content types across web applications. The metadata taxonomy and terms are stored in a database. Administrators can publish this information so that it is accessible to other web applications and site collections. These shared content types will be available in the content type gallery of all site collections using the metadata service.

### **Claims-based Authentication**

SharePoint 2010 now supports two modes of authentication. One is the classic authentication that includes NTLM, Kerberos and anonymous authentication. The second mode of authentication is the claims based authentication. Claims based authentication in SharePoint 2010 is designed to support multiple identity providers, integrate with identity providers via existing standards, and to support user identity delegation to other service applications across SharePoint farms and external authentication providers such as Live ID, ADFS, etc.

Claims-based authentication provides a common way for applications to acquire identity information from users inside their organization, in other organizations, and on the Internet. This identity information is stored in a security token that contains one or more claims about the user. A claim can be thought of as a set of information about a user. This may include name, email address, etc. This stays with the user throughout the session.

## Support for Additional Web Browsers and Mobile Devices

With MOSS 2007, only Internet Explorer (IE) provided an integrated experience for managing documents and pictures within a SharePoint site. SharePoint Server 2010 now offers nearly the same complete and rich experience for Mozilla Firefox users and for Apple Safari users on Macintosh. Safari for Windows, and other browsers such as Google Chrome and Opera allow content to be downloaded, but do not offer the integrated check-out, check-in, or direct editing within preferred applications. Some editing and management features require compliance with current web standards and employ ActiveX controls; such features will only work properly with IE 8 (or higher) and Firefox 3.5 (or higher) running on a Windows system.

SharePoint Server 2010 supports several micro-browsers, which are often employed on smartphones and other mobile devices. Mobile layouts with simplified navigation are also available for most sites. These features make it possible for a mobile workforce to access their content from practically anywhere.

## Closely-Related Products

Many service applications are included with SharePoint Server 2010 (see Table 2), but additional functionality can be provided to a SharePoint farm by incorporating services that are sold as separate products. Several such products that are designed to work closely with SharePoint Server are outlined in Table 3.

*Table 3: Products Closely Associated with SharePoint Server 2010*

Product	Description
FAST Search Server 2010 for SharePoint	Greatly enhances the scalability and performance of search operations. Makes it possible to perform content processing and property extraction, and offers advanced customization and programmability for search-based applications.
Office Web Apps	Provides multi-platform viewing, creation, and editing capabilities for Microsoft Word, PowerPoint, Excel, and OneNote® documents from within a web browser. The web-based user interface is similar to that offered by the corresponding Office 2010 client applications, and the resulting files are fully compatible.
Project Server 2010	Hosts Project Web Access instances that offer scheduling and certain calculations for Microsoft Project data, and exposes web services interfaces that allow interacting with these data.
SQL Server PowerPivot for SharePoint	Extends SharePoint Server and Excel Services to work with large workbooks that have been created with PowerPivot for Excel. These PowerPivot workbooks can benefit from SharePoint document management and publishing, collaboration, and server-side processing features.

## Upgrading to SharePoint Server 2010

MOSS 2007 Service Pack 2 is the minimum code level that can be directly upgraded to SharePoint Server 2010. As outlined in Table 4, there are two available upgrade methods.

*Table 4: SharePoint 2010 Upgrade Methods*

Upgrade Method	Description
In-Place Upgrade	Existing hardware is used, and all of the farm content and settings are updated in a single unified process. Many customizations are available after the upgrade, but often require manual changes to function as intended. <b>NOTE:</b> The upgrade requires taking the farm servers offline, and proceeds as a single continuous operation. Down time must be allocated to allow the entire upgrade process to proceed.
Database Attach Upgrade	The content is upgraded when an underlying content database is attached to a separate SharePoint 2010 farm. Farm services and settings are not migrated or updated, and must be configured manually. Content databases can be upgraded in any order, and multiple databases can be upgraded concurrently - often resulting in faster upgrade times. This method also makes it possible to merge content from several farms into a single SharePoint 2010 farm. <b>NOTE:</b> Settings and customizations must be manually applied to the new farm. Customizations that were not applied may lead to a poor user experience or degrade functionality. Direct access to the database servers is required, and the upgrade plan should account for the time required to copy the database across the network.

Which upgrade method<sup>5</sup> is best suited for a given SharePoint farm depends on many factors; an upgrade strategy should be planned carefully, and ideally be tested outside of the production environment. One helpful tool is the Pre-Upgrade Check utility. This utility scans your environment, and generates a report that highlights any warnings or red flags that need to be addressed before completing an upgrade. The tool can be accessed by running the following command:

```
stsadm -o preupgradecheck
```

---

<sup>5</sup> For more information, see "Upgrading to SharePoint Server 2010" at [http://technet.microsoft.com/en-us/library/cc303420\(office.14\).aspx](http://technet.microsoft.com/en-us/library/cc303420(office.14).aspx)

## Summary

SharePoint Server 2010 introduces a new service application architecture that enables improved scalability and availability for key services. These service applications can be combined to enable use cases in which users to interact with one another, find the information that they need, and and extract meaningful insights from company data.

While there is a great deal of out-of-the-box functionality in SharePoint Server 2010, additional value can be extracted through customization and integration with existing business processes. This paper is intended to introduce the SharePoint vocabulary, and to help administrators determine how SharePoint Server 2010 may be used in their environment. The small, medium, and large SharePoint Server 2010 farm reference architecture papers on [www.dell.com/sharepoint](http://www.dell.com/sharepoint) provide additional information that assists in determining a farm topology and selecting the underlying hardware. Dell also offers services for designing and deploying farms, and for customizing SharePoint and integrating it with existing infrastructures and business processes.