

# Microsoft® SharePoint® Server 2010

---

Designing and Implementing a Medium Farm

Dell™ Solutions Engineering

*Ravikanth Chaganti and Kevin Guinn*

May 2010



## Executive Summary

A SharePoint® Server 2010 farm hosts service applications that provide many different functions for its users. These functions are made available through Web servers, and store their configuration information and data in SQL Server databases. These interactions mean that a classical three-tier model is appropriate for SharePoint Server 2010. This paper provides a basic reference architecture that maps SharePoint roles to this model, and then extends this architecture to provide high-availability and redundancy at each tier.

The paper discusses key criteria and decision points for determining what server and storage hardware are appropriate for a medium-sized SharePoint Server 2010 farm. It offers guidelines and rules of thumb for determining how much storage space should be allocated, and when separate storage volumes are desirable. The intent of this paper is to enable administrators to understand important planning criteria, and to begin mapping their plan to a specific farm topology and well-suited hardware components. Dell™ offers more direct and specific guidance through consulting and other services.

A companion paper, *SharePoint Server 2010: An Introduction*, is available from [www.dell.com/sharepoint](http://www.dell.com/sharepoint). It offers an overview of SharePoint Server 2010, and provides common concepts and definitions that form a good basis for understanding the reference architectures presented in this paper.

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.

May 2010

## Table of Contents

- Executive Summary ..... 2
- Introduction ..... 4
- Designing and Building a Medium Farm ..... 4
  - Medium Farm Topology ..... 4
  - Database Considerations ..... 5
    - System Hardware Architecture ..... 6
    - Operating System and Database Editions ..... 6
    - Selecting a PowerEdge Server for SQL Server ..... 7
    - Database Storage Configuration ..... 7
    - Estimating Storage Capacity ..... 9
  - Considerations for WFE and Application Servers ..... 11
    - Selecting PowerEdge Servers for the Web Front-End and Application Servers ..... 12
    - Operating System and SharePoint Editions ..... 12
- Building High Availability and Redundancy in a SharePoint Farm ..... 13
  - High Availability for the Database Server ..... 14
    - Failover Clustering ..... 14
    - SQL Database Mirroring ..... 15
  - High Availability for Web Front-End Servers ..... 16
  - Redundancy for Application Server Roles ..... 17
- Summary ..... 18

## Introduction

Microsoft SharePoint Server 2010 builds on the capabilities that were offered in Microsoft Office SharePoint Server 2007 to provide a rich platform for collaboration, information sharing, and document management. SharePoint 2010 adds several new features, and introduces important architectural changes and product improvements. For a complete overview of SharePoint Server 2010, its features, and key dependencies, please see *SharePoint Server 2010: An Introduction*, a companion to this whitepaper that is available from .

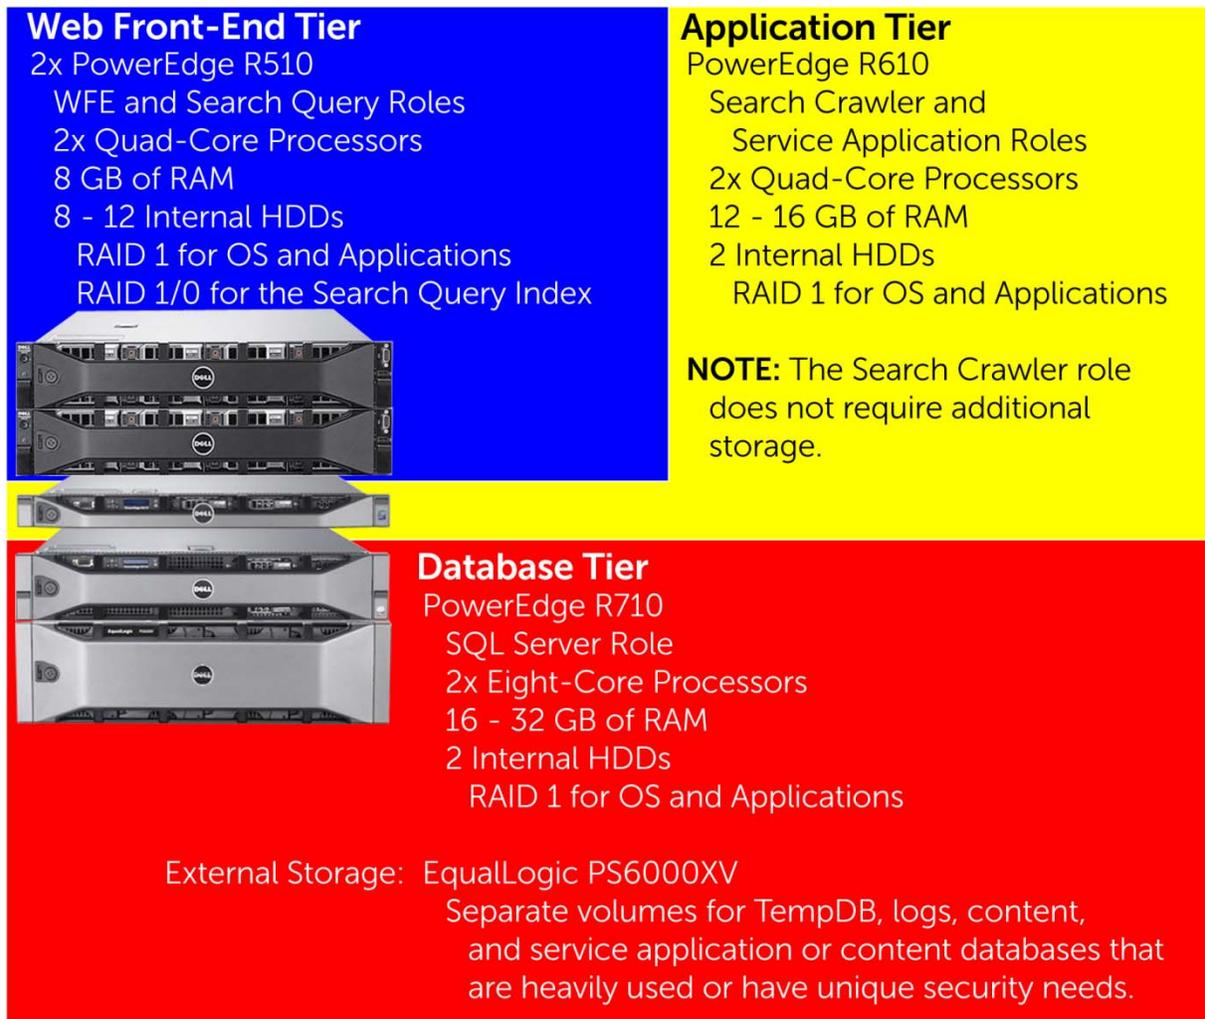
The remainder of this paper defines reference architectures that illustrate how a medium-sized SharePoint Server 2010 farm may be designed. There are many factors, such as information isolation policies, or the need to provide Enterprise Search capabilities for content outside of the farm, that can introduce the need for additional servers or more complex designs.

## Designing and Building a Medium Farm

Organizations planning to implement a SharePoint 2010 solution must plan for increasing workloads, anticipate changing user requirements, and determine how to provide availability for SharePoint sites and services. A server farm deployment model helps ensure that the solution infrastructure is scalable, flexible, and resilient to hardware failures. The farm model uses a dedicated database server, employs one or more application servers, and generally distributes Web front-end server roles across multiple hosts.

### Medium Farm Topology

A typical SharePoint medium server farm consists of three tiers: Web front-end, application, and database. Dedicated servers are used to host each tier to provide process isolation and to allow for future growth. Figure 1 depicts such a farm, which generally can be expected to provide services for up to 5000 users and to host up to a couple of terabytes of data. These numbers are broad guidelines, and will vary based on how the farm is configured, what services it provides, and how users interact with the farm. For example, two servers running Web front-end and search query roles are depicted in the figure. In environments with a highly active user population that performs many searches and other activities, the second server may be necessary. In less-heavily utilized environments, a single server may be sufficient to handle the load, and the second server could be used to provide additional availability for these farm roles. The following section provides detailed explanations and recommendations for allocating SharePoint 2010 server roles into this farm configuration. Additional information about providing high availability for various farm roles appears later in this paper.



*Figure 1: Medium Farm Topology*

## Database Considerations

SharePoint 2010 brings significant changes to the database tier of a SharePoint farm. With the new shared service application model, there are many additional SharePoint databases. In contrast to the “Shared Services Provider” and its single database that was employed in MOSS 2007, each service application now stores its content and configuration details in its own database. The database server remains an important aspect of a SharePoint 2010 farm; it is used to store SharePoint configuration details, to host user data and farm content, and to enable search functionality. To successfully address all of these needs, the database server must have adequate processing power and system memory to provide a good user experience and meet overall farm service level agreements (SLAs).

Certain SharePoint databases are generally more active than others and should be placed on their own volumes. To allow for future growth and scalability, it is recommended that SharePoint databases be placed on external storage arrays that are well suited to data-intensive applications with mixed I/O characteristics. SharePoint content databases are not easy to strictly categorize into traditional

database workload patterns—many user interactions, such as displaying list items, will tend to produce lightweight transactions, and others, such as uploading a file to a document library, will require moving a significant amount of data.

The database server queries and updates the database in response to requests from front-end Web servers. When configuring a medium server farm, install the Web and database server roles on separate machines to simplify planning, to increase load balancing and redundancy, and to facilitate the process of adding additional servers to the farm as needed. The SQL Server host should be configured before configuring SharePoint roles for any of the other farm servers; it must be available before the configuration database, which stores information about the farm servers and their roles, can be created.

If the farm will service requests from Internet users, consider logically isolating the database server from the Internet to reduce its security footprint. Although Web servers must be publicly accessible in this scenario, the database server can be located behind a firewall. In fact, the database server can be on a separate network segment and configured to only allow queries from specific hosts.

SharePoint relies heavily on the database server, perhaps more than any other server farm role, because it stores all configuration data and content. If there are any bottlenecks in the database server, performance will be impacted across the entire server farm. To achieve high performance from the SharePoint farm, administrators should monitor and maintain the database server.

#### *System Hardware Architecture*

SharePoint 2010 is available only for 64-bit extended (x64) architecture. This allows the SQL Server host to use RAM beyond 4 GB efficiently, which can help increase the number of connections and transactions that the database server can process. To fulfill this requirement, the host server must employ x64 hardware and run an x64 edition of a supported Windows Server® operating system. Currently-shipping Dell PowerEdge servers with Intel Xeon or AMD Opteron processors, such as those featured in this paper, employ x64 hardware and may be successfully used in a SharePoint farm.

#### *Operating System and Database Editions*

Choosing among the various editions of Windows Server 2008 with SP2 or Windows Server 2008 R2, and SQL Server 2008<sup>1</sup> or SQL Server 2008 R2 is primarily a function of which system specifications and software features are important in the intended environment. The most important considerations for the operating system are listed in Table 1, and those for the database software are listed in Table 2. The high-availability features listed in these tables—such as failover clustering and database mirroring—are examined later in this paper.

*Table 1: Critical Factors for Selecting an Operating System Edition*

	Windows Server 2008 SP2 or 2008 R2 x64 Standard Edition	Windows Server 2008 SP2 or 2008 R2 x64 Enterprise Edition
Support for Failover Clustering	No	Yes, up to 16 nodes
Maximum Physical RAM	32GB	2TB

---

<sup>1</sup> SQL Server 2008 with Service Pack 1 and Cumulative Update 2 is the minimum required version of SQL Server 2008. SQL Server 2005 with Service Pack 3 and Cumulative Update 3 is also supported.

*Table 2: Key Factors for Selecting a SQL Server Edition*

	SQL Server 2008 or 2008 R2 x64 Standard Edition	SQL Server 2008 or 2008 R2 x64 Enterprise Edition
Support for Failover Clustering	2 nodes	Max number of nodes supported by OS
Database Mirroring	High-safety mode only	All high performance and high safety modes
SQL Server Reporting Services	Available	Available
SQL Server Analysis Services	Available	Available
Database Compression	Not Available	Available
Transparent Data Encryption	Not Available	Available
Resource Governor	Not Available	Available
Remote BLOB Storage	Not Available	Available

Consider the following example: a single database server with content databases on an external storage array provides sufficient scalability for a company's medium farm. However, there is a business need to enhance the security of data stored in the SharePoint content database by deploying transparent data encryption and to provide higher levels of service uptime using either failover clustering or database mirroring.

The database encryption requirement would drive the selection of SQL Server Enterprise Edition. If failover clustering were ultimately chosen as the database availability solution, then the hosts in the SQL Server cluster would also each require the Enterprise Edition of Windows Server.

#### *Selecting a PowerEdge Server for SQL Server*

The SQL Server host will generally require more CPU and system memory resources than other farm servers. For a typical medium farm, Dell recommends a minimum of eight processing CPU cores and 16 GB of RAM for the database server. If the farm will be used as a platform for Business Intelligence (BI), and will make extensive use of SQL Server Reporting Services and Analysis Services, then increasing the number of CPU cores is strongly recommended. For deployment scenarios where a high number of transactions is likely, where the corpus of data exceeds 2 TB, or where many service applications will be hosted, increasing the minimum RAM configuration to 32 GB is advised. In extreme situations, heavily-used farms with high activity, many service applications, a large corpus of data, and extensive use of search may make multiple SQL instances advisable to facilitate memory allocation and tuning.

Dell offers many modular (blade) servers and rack-mountable servers that meet or exceed these requirements. When external storage arrays, such as the PS6000XV, are used to house the SharePoint databases, the primary considerations for the SQL Server host are processing, memory, and I/O capability. Servers such as the M610 and R610 provide two processor sockets, each of which can accommodate an eight-core processor and offer maximum memory configurations well beyond the above recommendations. If more I/O capacity is needed, the M710 and R710 provide additional expansion slots that can accommodate storage or network controllers. For BI environments, the M910, and R910 provide four CPU sockets (up to 32 processing cores) and offer additional memory capacity.

#### *Database Storage Configuration*

By default, without any additional configuration, SharePoint will store all of its data in a single content database, and a medium farm can reasonably be expected to store up to a few terabytes (TB) of data. In MOSS 2007, the general rule-of-thumb was to limit the size of an individual SharePoint content database to 100 GB. SharePoint Server 2010 improves the content storage architecture and can

mitigate some performance problems associated with large lists by allowing throttling. To improve content isolation, reduce backup / recovery times, and generally make data management easier, it is still recommended to plan and implement an information architecture that includes multiple moderately-sized content databases and restricts the number of site collections that are associated with each individual content database. Similarly, to ensure the best farm performance, it is recommended that general-purpose content databases be limited to 200 GB or smaller<sup>2</sup>. That said, it is still strongly recommended that each site collection is limited to 100 GB, unless it is the only site collection within its content database.

As mentioned previously, most SharePoint 2010 service applications employ their own database. The I/O patterns associated with these databases will vary based on how often the associated services are accessed, and in what capacity. An excellent overview of all SharePoint 2010 databases, with recommendations for planning and maintenance is available on Microsoft® Technet<sup>3</sup>.

Another new database in SharePoint Server 2010 is the usage and health data collection database, also known as the logging database. A SharePoint timer job collects, aggregates and collates ULS logs from each farm server, and stores them in the logging database. For planning purposes, it is best to assume that this database will be large, heavily used, and write-intensive.

Table 3 outlines some of the database entities that should normally have dedicated volumes, and recommends appropriate RAID configurations. Physical disks that operate at 15k rpm are generally preferred for all transaction logs and for more-heavily-used SQL Server data volumes.

*Table 3: Database Storage Recommendations*

Function	Preferred RAID Level	Notes
Transaction Logs	RAID 1/0	Transaction logs volumes should be optimized for writes, and should allow ample room to allow for log growth.
Content Databases	RAID 1/0 for content databases that host site collections focused on collaboration or that will feature extensive write activity.  RAID 5 for content databases that host site collections focused on portals or document repositories that feature predominantly read activity.	Consider dedicated content Databases for any site collection that is expected to have an unusually high level of activity.  Use separate databases for site collections whose access patterns differ greatly, and optimize volumes by access pattern.  Provision multiple data files for each configuration database.

<sup>2</sup> In specific read-intensive content and records management scenarios, databases up to 1 TB are supported.

<sup>3</sup> See “Database types and descriptions (SharePoint Server 2010)” at [http://technet.microsoft.com/en-us/library/cc678868\(office.14\).aspx](http://technet.microsoft.com/en-us/library/cc678868(office.14).aspx)

Function	Preferred RAID Level	Notes
TempDB	RAID 1/0	TempDB use can be influenced by SQL Server factors outside of SharePoint (such as rebuilding indexes).  It is best to treat TempDB as if it will be heavily utilized.
Logging Database	RAID1/0	This database should be considered to be heavily utilized.
Search Crawler database	RAID 5 or RAID 1/0	Read-heavy, but poor write performance could increase the time required to complete a crawl.
Search Property database	RAID 1/0	Optimize for good read and write performance. Although the typical access pattern favors writes, poor read performance can impact search query times.

### *Estimating Storage Capacity*

Estimating the total data storage space required for a SharePoint 2010 farm can be challenging, and is highly dependent on how SharePoint 2010 will be used within the organization. A departmental or company-wide portal may demand significantly more storage be allocated versus a team site or an individual MySite; however, there are likely to be many more of these smaller sites hosted within the farm. In addition, document libraries will generally require more storage than site features such as blogs, wikis, and other types of lists. If the versioning, recycle bin, or audit features of SharePoint will be used, then additional space must also be allocated for this data.

Table 4 outlines a guideline<sup>4</sup> and provides some examples for estimating the required amount of storage for the content database. Specifying a lower fill factor provides additional space in the database that can be used to store document versions and site recycle bins, or can be used to fulfill audit requirements. This “unused space” also serves as a buffer that can accommodate growth of the content stored in the farm.

---

<sup>4</sup> At the time this paper was written, SharePoint Server 2010 had not yet been released, and updated formulas for estimating content database size based on document and list parameters were not yet available. This information is expected to be released in the future and made available on [http://technet.microsoft.com/en-us/library/cc298801\(office.14\).aspx](http://technet.microsoft.com/en-us/library/cc298801(office.14).aspx)

*Table 4: Estimating Content Database Size Requirements*

	Size of content to be stored	Database overhead	Fill factor (for growth and versioning)	Minimum disk space to allocate
Rule of thumb	X	20% (0.2 * X)	50 - 70% recommended	~1.7 * X (70% fill factor) ~2.4 * X (50% fill factor)
150 GB content	150 GB	30 GB	50%	360 GB
350 GB content	350 GB	70 GB	70%	600 GB
1 TB content	1 TB	205 GB	60%	1.92 TB

**NOTE:** In Table 4, the examples all yielded recommended disk allocation sizes in excess of 200 GB. In these situations, it is recommended that site collections in the farm be split among several content databases, and that each individual general-use content database is no larger than 200 GB and each site collection is no larger than 100 GB.

**NOTE:** Employing a Remote BLOB Store (RBS) will shift a significant amount of document and file storage to the RBS provider, and therefore reduce the size of associated content databases. Other list items are largely character-based, and will continue to be stored directly in the content databases.

In addition to the content databases, the SQL Server host will also require space for housing the configuration, central administration, search (administration, crawl, and property), and service application databases. The configuration, central administration, and search administration are generally small (on the order of 1-2 GB). The search crawl and property databases are based on the amount of content in the farm. Based on initial Microsoft testing and analysis<sup>5</sup>, assume the following sizes for these databases:

- Crawl: 4.6% of the total combined size of all content databases
- Property: 1.5% of the total combined size of all content databases

In SharePoint 2010, trace logs & usage information are stored in the logging database. The default retention policy for this data is 14 days, but it can be configured using PowerShell. Longer retention periods will require additional storage space. Depending on how Usage and Health Data Collection is configured, this database can store very granular information to enable web analytics and usage monitoring; this information is collected and aggregated from every server in the SharePoint 2010 farm. For these reasons, the logging database can grow to a large size, and may do so within in a short time if many different facets of health data are collected. This database should generally be housed on a dedicated volume, and should be monitored carefully if a large amount of data will be collected for an extended period of time.

Employing quotas to limit the size of individual sites and establishing governance policies to manage content and control the number of sites in the farm can help control the total space required by a given farm. Regardless, it is important to plan for capacity to increase over time, and to build a flexible and scalable infrastructure that will enable the farm to grow.

---

<sup>5</sup> See "Storage and SQL Server capacity planning and configuration (SharePoint Server 2010)" at [http://technet.microsoft.com/en-us/library/cc298801\(office.14\).aspx](http://technet.microsoft.com/en-us/library/cc298801(office.14).aspx)

The Dell EqualLogic™ PS6000XV is well suited for storing typical SharePoint databases. Multiple volumes can be configured to support these databases without the need to determine which individual disk spindles should be allocated to which individual volume. If the farm grows and more IOPS are required to meet the needs of the farm, additional PS6000XV arrays can be added to an EqualLogic group without having to reconfigure or rebuild the existing volumes. These arrays also provide hardware-based snapshots, replication, and other advanced storage features that administrators can use as part of a data protection strategy for SharePoint. For BI environments that aggregate large amounts of data, the PS6010XV - with its 10 gigabit connectivity and associated increased bandwidth (in terms of MB/s) - could improve performance.

## Considerations for WFE and Application Servers

In a SharePoint farm, Web front-end servers render HTML to requesting clients and route requests to other services in the farm. Web servers also host the SharePoint Web application in Internet Information Services (IIS), and these servers send queries and updates to database servers in response to user requests. As the front-end Web server is the only server role that is directly exposed to the end user, install it behind a firewall or take appropriate steps to secure the server in the network. All SharePoint features, site definitions, rendering, and localization files are stored on the front-end Web server. SharePoint also installs a number of Windows services on front-end Web machines that are required to communicate with other servers in the farm and to perform administrative tasks.

When deploying a front-end Web server, it is generally a good idea to deploy the “Complete” installation type to ensure easier future scalability. In a typical medium farm deployment, deploy more than one Web front-end server for load balancing and redundancy purposes.

Managing a server farm is easier when multiple servers host the same service. It allows administrators to perform a rolling update of both software and hardware without affecting the availability of SharePoint services. Also, it enables them to add additional servers and reallocate roles as the farm grows to meet changing business needs.

Unlike MOSS 2007, in SharePoint 2010 index and query roles can be installed on multiple hosts in the farm. SharePoint 2010 also includes a crawler role that can be hosted on multiple hosts for load balancing and redundancy.

To improve search performance, deploy query server roles on the Web front-end servers so that all SharePoint search requests are processed and returned to the Web front-end component efficiently. The Query component can be installed on multiple servers, as can the Web front-end role. In SharePoint 2010, the search index that is stored on the query servers can be divided into index partitions, and each index partition can be assigned to one or more query servers in the farm. An index partition holds a part of the overall search index. For example, if there are three query servers in the farm, three index partitions could each be assigned to a single query server. This way, each query server holds one third of the overall search index and therefore requires less disk space. Such a configuration, however, would not provide high-availability and could cause partial results to be returned to users. To provide availability, each index partition can be assigned to more than one query server.

Therefore, the rule of thumb for disk size for query servers is based on the size of the crawl database, and the number of index partitions that are assigned to each query server. If a query server will store the entire index, allocate sufficient space for approximately 2.5 times the size of the crawl database (roughly 12% of the size of the crawled content). If there are multiple index partitions, then this

storage allocation can be reduced based on the ratio of assigned partitions versus total index partitions. While this rule of thumb provides a baseline, the index size will vary based on the definitions provided by the content source. For best results, store the search index on its own RAID 1/0 volume on each query server. Read performance helps provide query results more quickly, and write performance helps reduce the time required to complete a crawl operation.

In SharePoint 2010, crawler is stateless. This means that the crawler component does not store a copy of the index as it did in MOSS 2007. Therefore, there are not any special storage requirements for the search crawler server. Crawling content and building index partitions can be a CPU- and memory-intensive operation; if the crawler servers host other service applications, administrators should generally schedule crawls to occur at off-peak times.

### *Selecting PowerEdge Servers for the Web Front-End and Application Servers*

For a host that provides the Web front-end and query components, consider a server with at least four processing cores and 8 GB of RAM. The Web front-end component can be network I/O intensive, so administrators should take precautions to handle potential bottlenecks. NIC teaming software can be used to not only provide redundancy, but also to distribute network activity between multiple NICs on the server. Provisioning the query component requires that the host server be able to hold copies of all assigned index partitions.

The PowerEdge™ R510 is a two-socket server that can be configured with eight or even twelve internal disks, making it very well suited for systems that host the query component. If the farm is being deployed using modular servers, the M710 offers four hard disk drives, which helps make it capable of storing indexes. If a high volume of search activity is expected, it may be preferable to provide external storage volumes for the query component. Without the requirement for internal storage, systems such as the M610, R410, and R610 become good fits for these roles. Because the crawler and service applications do not require local storage, these servers are also well-suited for the application tier.

### *Operating System and SharePoint Editions*

It is uncommon for a SharePoint application server or Web front-end server to benefit from the enhancements in Windows Server® 2008 or Windows Server 2008 R2 Enterprise Edition that are not available in Standard Edition. At the application tier and Web front-end tier, failover clustering is not employed for availability; instead, roles within the farm are configured to run on multiple servers. This means that, in most cases, Windows Server 2008 or Windows Server 2008 R2 Standard Edition is sufficient for application and Web front-end servers.

For SharePoint 2010, the biggest differentiator between Standard Edition and Enterprise Edition is related to the features and services that are used either to develop solutions for Business Intelligence and Business Analytics or to offer enhanced services for specific Microsoft client applications. These features are enabled by the Business Data Connectivity, Excel Services, Visio Services, etc. If these features are (or will be) needed in the farm, select SharePoint 2010 Enterprise Edition; otherwise, Standard Edition is likely to be sufficient.

## Building High Availability and Redundancy in a SharePoint Farm

A business critical application hosted on a SharePoint Server solution that is not designed for high availability can lead to unacceptable amounts of service downtime and potentially even revenue loss. To ensure that hardware failures do not affect the business, farm architects must consider availability as a part of the solution design. Overall capacity and performance requirements will affect how to implement availability within the SharePoint farm. To determine whether to implement high availability within a solution, look at the expected uptime, the number of users affected if service becomes unavailable, and the productivity and business losses that will accumulate until service is restored. If the solution requires high levels of uptime to prevent productivity loss when the service is offline, consider high availability options for each of the server roles in the SharePoint solution. Typically, high availability is among the first considerations in determining the number of server computers required to deploy a server farm.

Within the scope of this paper, availability refers to service or server role availability only, and does not include high availability of data through backup and restore or data protection mechanisms. To mitigate any unexpected data loss, farm designers must evaluate various disaster and data recovery options as a part of SharePoint solution design. SharePoint 2010 offers out of the box backup and restore capabilities. Native backup and restore can be performed using the SharePoint central administration user interface or using Windows PowerShell. SharePoint 2010 offers granular backup and restore options, including restore from an unattached SQL database. For more information, see Backup and Recovery (SharePoint Server 2010) at [http://technet.microsoft.com/en-us/library/ee662536\(office.14\).aspx](http://technet.microsoft.com/en-us/library/ee662536(office.14).aspx). Alternatively, use Microsoft SQL Server 2008 backup and restore methods, or other mechanisms such as Microsoft Data Protection Manager 2010.

## High Availability for the Database Server

Microsoft SQL Server 2008 provides features to ensure high availability. Table 5 summarizes and contrasts database mirroring and failover clustering, which are the two most-commonly used techniques.

*Table 5: Comparison of Failover Clustering and Database Mirroring*

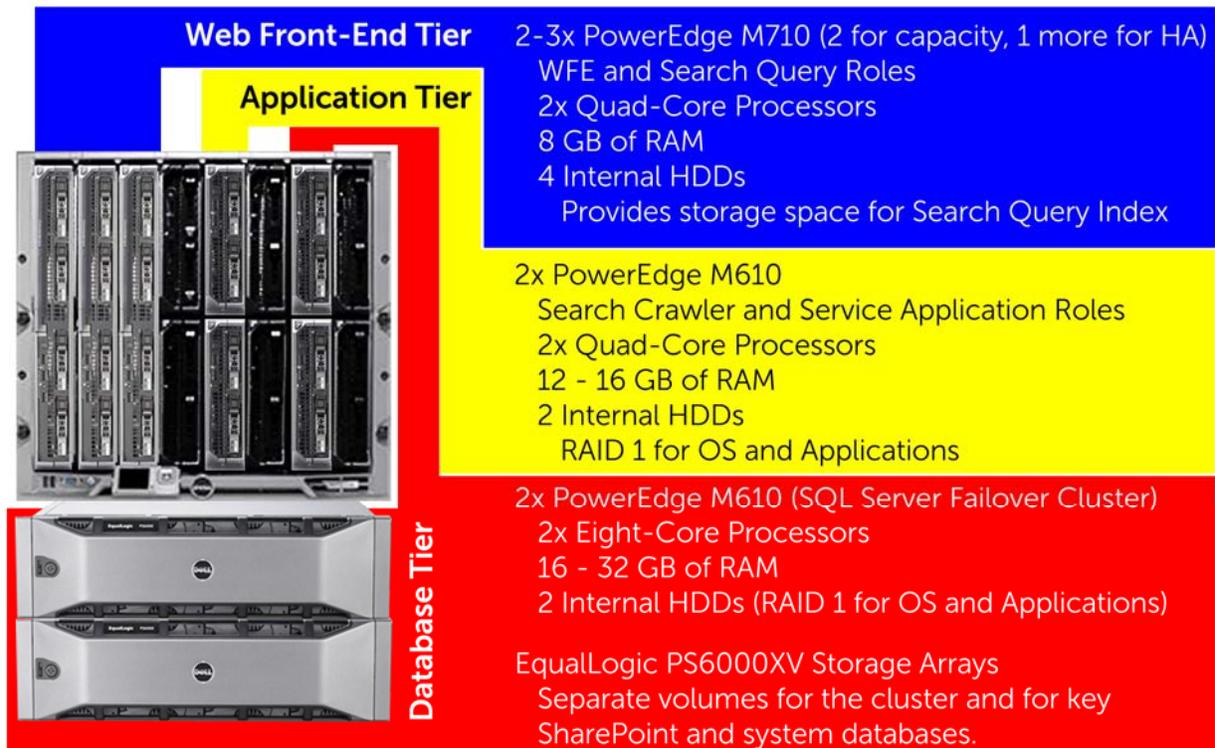
	SQL Failover Cluster	SQL Database Mirroring
<b>Protection Provided</b>	Data is stored on a single shared storage array and employs RAID, hardware redundancy, and multipath I/O for protection. The clustered SQL Server instance can run on any node in the cluster, providing tolerance for a wide range of hardware and software faults on the host servers. The design goal for a failover cluster is to militate against a single point of failure that might cause the database to remain offline.	Both the primary SQL Server and its mirroring partner have their own data copy, and each node's storage subsystem features RAID and hardware redundancy. Many SharePoint 2010 databases and service applications allow specification of mirror-aware connection strings that help facilitate automated failover.  The fact that there are independent database servers and separate copies of the data provides protection against faults that may adversely impact one of the database host servers or its storage array.
<b>Recovery Method</b>	When there is a failover event, there will be some downtime as SQL Server services are started on the alternate cluster node and the database steps through the redo logs to ensure consistency. Once this is complete, services will resume without intervention.	SharePoint 2010 provides native support to SQL Server mirroring for many of the SharePoint 2010 databases. While creating web applications and service applications, specify a failover database server; in case of primary SQL server failure, SharePoint automatically switches to a secondary. The default time for switching to a failover server is 15 seconds.
<b>Performance Impact</b>	The overhead associated with running a clustered instance versus a standalone instance of SQL Server is generally considered negligible.	Logs are compressed and transmitted to the database mirroring partner. This adds some additional CPU, network, and I/O load to the primary SQL Server system.
<b>Operating System and SQL Server Editions</b>	<b>Operating System:</b> Enterprise Edition is required for failover clusters <b>SQL Server:</b> Standard Edition supports just two cluster nodes. Enterprise Edition supports as many nodes as the operating system.	<b>Operating System:</b> Standard Edition is generally sufficient. <b>SQL Server:</b> Standard Edition provides high-safety (synchronous) mirroring. Enterprise Edition adds support for high performance (asynchronous) mirroring. To enable automated failover high-safety mode and a witness server are necessary.

### *Failover Clustering*

In a failover clustering scenario, the database is placed on an external shared-storage array, such as a Dell EqualLogic PS series array. SQL Server 2008 failover clustering works in a shared-nothing model; this means that each node in the cluster owns and manages its own resources. In the case of a node failure, services running on the failed node are moved to an active node.

Figure 2 shows a medium farm that is configured to provide high-availability at every tier. An additional load-balanced Web front-end server has been provisioned; service applications, including the

search crawl and query components have been configured to run on multiple servers. Finally, this configuration features a failover cluster that hosts the SQL Server databases for the SharePoint farm.



*Figure 2: Highly-Available Medium Farm with a SQL Server Failover Cluster*

Microsoft SQL Server 2008 failover clustering is active-passive with respect to a clustered SQL Server instance. A given instance of SQL Server operates on a single cluster node at any given point in time, but additional cluster nodes may host other services or instances. If the node hosting a specific SQL instance fails, all of the requisite resources and services will fail-over to a surviving cluster node. Ideally, at least one cluster node should be configured to be fully passive; its role in the cluster is to provide capacity to continue operations in the event that another node fails.

### *SQL Database Mirroring*

In SQL Server 2008, database availability can also be achieved using the database mirroring feature. This feature is implemented on a per-database basis. Database mirroring works by compressing transactions logs from the primary database, and shipping them to the mirror database on a separate server. The logs are applied to the mirror database, which allows both copies to reflect the same updates, inserts, and other changes to the database. SharePoint 2010 generates connection strings that help enable automatic failover to the mirrored database instance for many of its databases. For this to work best, configure mirrors in high-safety mode and employ a witness server; other types of mirror configurations do not allow for automatic failover. Such a database mirroring configuration is illustrated in Figure 3, which also features additional servers to provide availability for the Web front-end and application tiers of a medium SharePoint 2010 farm.

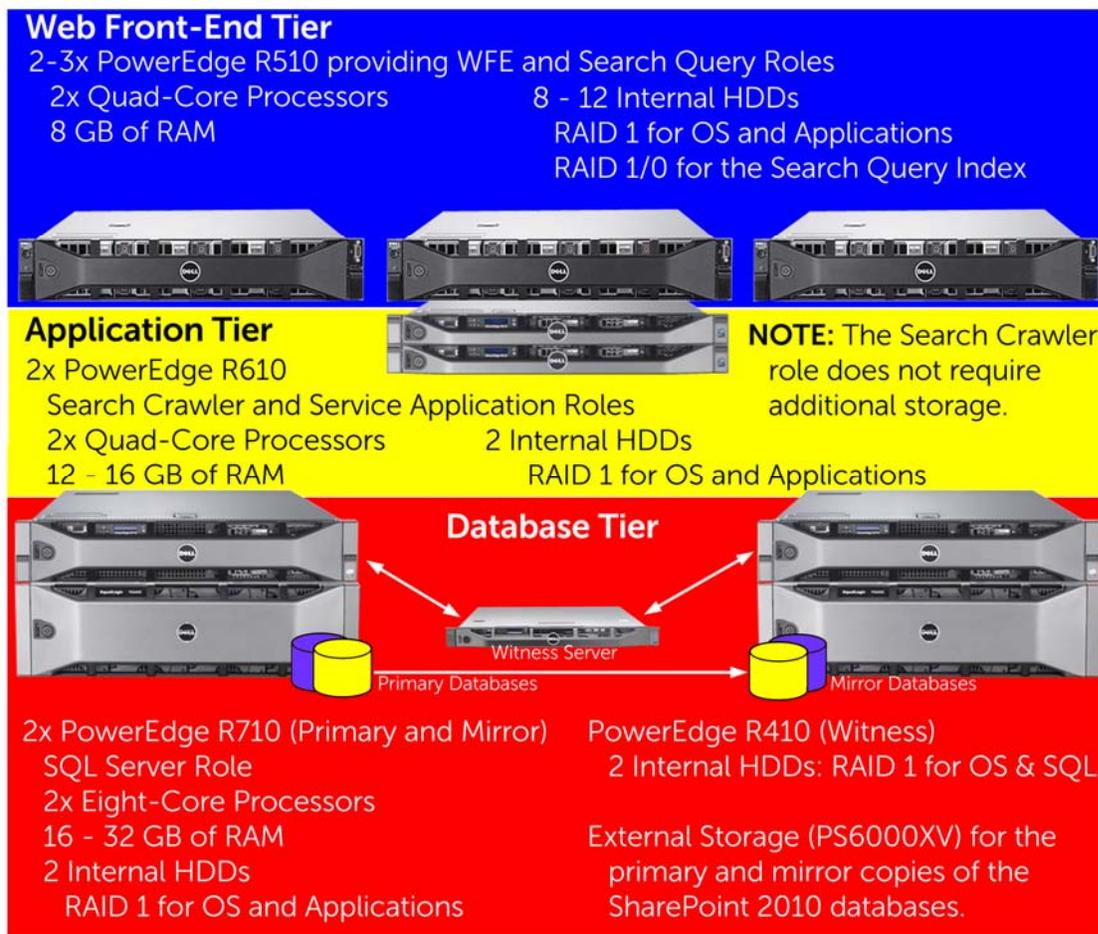


Figure 3: Highly-Available Medium Farm with Database Mirroring

NOTE: The SharePoint logging database does not support mirroring, but in the event of a SQL Server host failure, this database can be manually recreated.

### High Availability for Web Front-End Servers

To increase availability and handle higher Web server workloads, Windows Server 2008 provides built-in software for Network Load Balancing (NLB). In an NLB cluster, multiple servers share a common IP address and distribute Web server workloads among nodes. If advanced load balancing heuristics, detailed monitoring, or statefulness for clients are required, consider a hardware load balancer.

Multiple Web front-end servers are shown in Figure 2 and Figure 3. In each of these sample farm architectures, a third server has been added for redundancy. If the number or concurrency of user requests in a particular farm is low, then two WFE servers may be sufficient to provide availability (e.g. one server is capable of handling the workload). However, in a high-demand environment, two or more servers will be needed to handle user requests; an additional WFE server allows service levels to be maintained in the event of a failure. The query server role is also configured on all of the Web front-end servers to ensure availability.

The query component can be distributed across several servers. To ensure redundancy, each configured crawler index partition should be assigned to more than one query server. For example, consider a

farm with three crawl index partitions (A, B, and C), and three servers hosting the query component (Q1, Q2, and Q3). To provide redundancy for the query role, each index partition should be mapped to at least two query servers. One such configuration is as follows:

- Query server Q1 is assigned index partitions A and B
- Query server Q2 is assigned index partitions B and C
- Query server Q3 is assigned index partitions A and C

In the above configuration, the redundant index partition assignments ensure that the loss of a single query server does not provide incomplete search results. If server Q1 were to fail, results from partition A would be handled by server Q3, and results from partition B would be returned by server Q2. The reference architectures illustrated in Figure 2 and Figure 3 can be configured in this manner to provide redundancy not only for the load-balanced Web front-end, but also for the query component.

## Redundancy for Application Server Roles

The application tier of a SharePoint farm hosts multiple components, such as the search crawler and other service applications (e.g. Excel Services, Word Automation, Managed Metadata, etc.). One major improvement is that the crawler component in SharePoint 2010 can be made highly available. This can be achieved by introducing more than one server to host this function in to the SharePoint farm, and providing “N+1” hosts to create redundancy (i.e. if one crawler can complete a full crawl in the desired time, a second crawler can be added to provide availability for this function).

In the architectures illustrated in Figure 2 and Figure 3, high availability was provided for the crawler component by provisioning multiple servers at the application tier, and hosting the crawler component on more than one host. This same technique applies for most of the service applications as well.

**NOTE:** For farms that need to crawl a large corpus of content, it is also possible to provision more than one search crawler database. Each crawler database is associated with a unique search administration component, but can be shared among several crawler components. By protecting the crawl databases, provisioning multiple crawlers, generating index partitions, and assigning the index partitions to more than one query server, a SharePoint 2010 farm can provide end-to-end high availability for the search function.

## Summary

A companion paper, *SharePoint Server 2010: An Introduction*, is available from [www.dell.com/sharepoint](http://www.dell.com/sharepoint). That paper provides a high-level overview of SharePoint Server 2010, and introduces key concepts and terminology that are helpful for understanding the reference architectures presented in this paper.

A medium SharePoint 2010 farm consists of several servers, each of which is provisioned with different SharePoint components. Three-tier architectures are recommended - featuring a presentation tier that provides Web and search query functions, an application tier that provides search indexing and service application functions, and a database tier that hosts the SQL Server databases for the farm.

Because farms of this scale often contain important data or provide necessary business functions, the farm architecture should be designed with high availability and redundancy in mind. For the SQL Server database, the alternatives include failover clustering and database mirroring. SharePoint Server 2010 provides enhanced support for database mirroring, natively building SQL connection strings that are mirror-aware. At the application tier, SharePoint 2010 has enhanced its search architecture, and can now be more-easily configured to offer redundant search crawlers. Similarly, most service applications and the search query component can be hosted on more than one server. Finally, a load-balancing solution is recommended for the Web front-end servers.

Dell modular and rack-mounted servers are well-suited for all SharePoint 2010 farm roles. Some of these servers have been highlighted in the reference architectures presented in this paper, along with external storage arrays that can house the SharePoint databases. Because SharePoint is both a content repository and a development platform, each farm is unique. Dell also offers services to help design, deploy, optimize and customize SharePoint Server 2010.