Microsoft® SharePoint® Server 2010 Implementation on Dell™ Active System 800v

A Design and Implementation Guide for SharePoint Server 2010 Collaboration Profile on Active System 800 with VMware vSphere

Dell Global Solutions Engineering

Revision: A01

February 2013

Rev A01

Contents

Tables

Figures

Executive Summary

Traditional IT infrastructures have been transformed with the use of virtualization technologies in the datacenter. Benefits of virtualization, such as consolidation and availability, have enabled creation of large clusters of hardware resources managed with virtualization and related software. These architectures improve the ease of management of underlying hardware resources by administrators and enable coexistence of different tiers of enterprise applications on the same hardware.

Dell™ Active Infrastructure is a family of converged infrastructure solutions that combine servers, storage, networking, and infrastructure management into an integrated and optimized system that provides virtualized resource pools. Active Infrastructure helps IT respond to dynamic business demands, maximize data center efficiency, and strengthen service quality more efficiently. Dell Active System 800 is a pre-integrated virtualized infrastructure system that enables IT and business stakeholders to accelerate application and IT service delivery. Designed for reliability, Active System 800 eliminates single points of failure and incorporates redundancy into infrastructure components.

Critical enterprise applications, including Microsoft® SharePoint® Server 2010, were designed to be run on dedicated physical hardware with best practices around performance and availability detailed in Microsoft and partner documentation. With the introduction of virtualization in the datacenter, Microsoft also provided best practices for moving applications such as SharePoint onto virtualized servers running multiple virtual machines (VMs). However, many approaches focus on dedicating a set of hardware resources for the application. This ensured the appropriate performance and availability of the application during hardware failure scenarios. However, it also leads to the creation of silos of hardware resources dedicated for specific applications, resulting in management complexity for IT administrators.

This guide presents a framework for implementing a virtualized SharePoint 2010 farm in collaboration profile on Active System 800 with VMware® vSphere™ (Active System 800v). This framework builds on the architecture best practices of the Active System and uses VMware virtualization software capabilities to develop rules that allow application VMs to be deployed and managed without administrators worrying about the exact physical servers they might be running on. The design also provides dedicated storage resources for application data, adequate and guaranteed storage performance to meet application requirements.

This guide also provides an instantiation of the application design to support up to 5,000 users using Active System 800v. This implementation was then validated for performance and availability by running the Dell SharePoint Load Generation tool to determine steady-state and degraded-state performance characteristics and ensuring that the virtualization resource scheduling features implemented for the application VMs lead to the application returning to steady-state quickly upon hardware resource failure.

For more information, please refer to the Active System 800v documentation.

Additional information on SharePoint 2010 design and deployment best practices is available from www.dell.com/sharepoint.

Audience and Scope

The content of this guide for SharePoint 2010 on Active System 800v is intended for IT professionals, consultants, and IT administrators who design and implement SharePoint 2010 application farms. This guide describes the design principles and best practices for virtualizing SharePoint Server 2010 application and provides an overview of how Active System 800v pre-integrated platform architecture benefits the application deployment. The reader is expected to have an understanding of VMware vSphere 5.1 and SharePoint Server 2010.

For more information, please refer to the Active System 800v documentation.

For step-by-step instructions for implementing the recommendations in this paper, refer to the References section at the end of the paper.

Introduction

Recent advancements in server, storage, and networking technology have enabled building efficient infrastructure solutions for the private cloud. Dell's latest hardware portfolio is well equipped to take advantage of the technology advancements and help build solutions with converged network infrastructure. Active System 800v is one of such systems that include the latest Dell PowerEdge™ blade servers, Internet Small Computer Simple Interface (iSCSI) / 10GbE switching and Dell EqualLogic™ iSCSI storage to build private cloud solutions.

Virtualization software from VMware, including vSphere 5.1 with vCenter™ Server, provides fine-grained virtual resource management tools that enable advanced management of application resources to provide reliable performance guarantees and increased availability without tying VMs to individual servers. While the underlying infrastructure is becoming more feature-rich, accurately sizing and validating SharePoint Server 2010 on such an infrastructure requires a new approach. Some of the considerations include accurately defining resource requirements and maintaining application best practices at the same time, while ensuring application performance is within the thresholds that are prescribed by Microsoft. As a consequence, a methodology that correctly designs and deploys SharePoint Server 2010 on new converged infrastructure is essential.

One such methodology is described in this guide. This methodology helps reduce dependencies on the underlying infrastructure from a sizing and deployment perspective, while simplifying the management of SharePoint Server 2010 ecosystem. It does this by abstracting out the SharePoint solution in the form of a virtualized resource pool configuration consisting of memory, CPUs, and virtual network adapters that are provided by Active System 800v. The resource configuration also consists of an appropriate number of VMs pertaining to various server roles that are part of the SharePoint server deployment. The performance of SharePoint Server 2010 is guaranteed as long as the underlying platform is able to meet the virtualized resource pool configuration requirements.

The subsequent sections of this guide offer specific details around the virtualization best practices for SharePoint Server 2010 application, how Active System 800v design principles help in mapping these best practices, and presents a reference implementation for up to 5000 concurrent users in SharePoint collaboration profile to demonstrate the benefits of Active System 800v when implementing SharePoint 2010 application farms, and finally provides details on how the reference implementation in this paper performed under load testing.

Active System 800

Active System 800, a member of Dell Active Infrastructure family, is a converged infrastructure system that has been designed and validated by Dell™ Engineering. It is available fully racked, cabled, and delivered to your site, to speed up deployment. Dell Services will deploy and configure the solution tailored for business needs, so that the solution is ready to be integrated into your datacenter. Active System 800 is offered in configurations with either VMware® vSphere® (Active System 800v) or Microsoft® Windows Server® 2012 with Hyper-V® role enabled (Active System 800m) hypervisors. This guide is based on the reference architecture for the VMware vSphere based Active System 800v solution.

As a converged infrastructure solution, Active System 800 employs the Data Center Bridging (DCB) suite of standards and technology available on the Dell Force10™ network switching infrastructure to prioritize IPC, Data, or iSCSI traffic in a converged fabric. Active System 800v converged infrastructure is designed to enable Ethernet transport reliability for application workloads. The underlying design principles make Active System 800v a suitable converged infrastructure solution for virtualizing IO intensive workloads.

Figure 1 Active System 800 Configurations

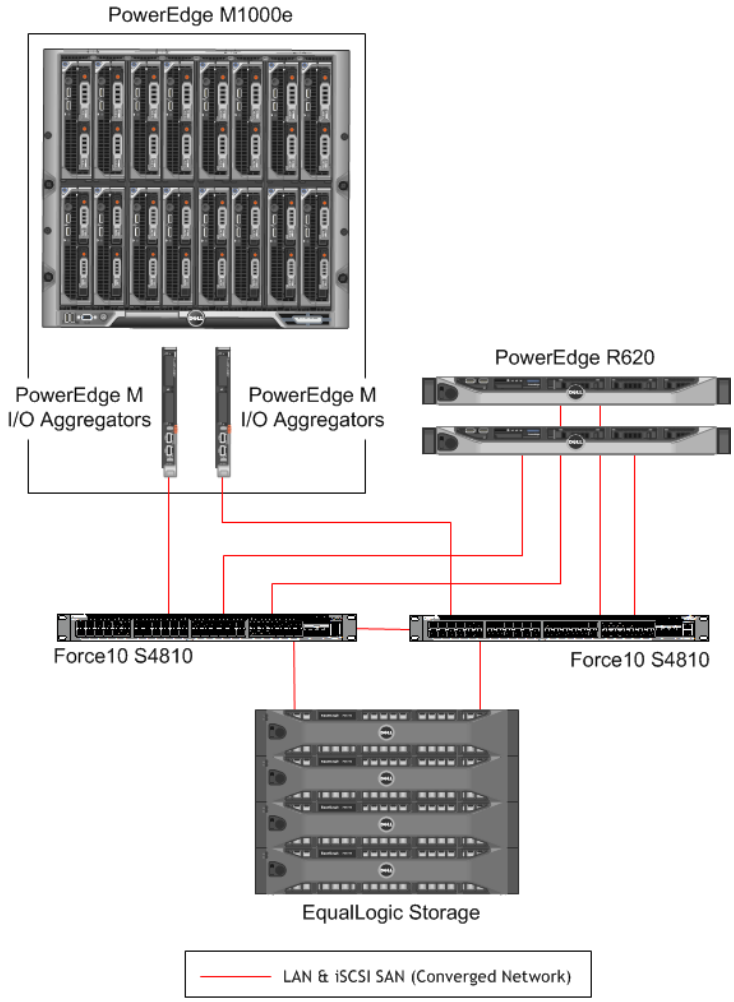

Active System 800v includes Dell PowerEdge™ M1000e blade chassis with Dell PowerEdge™ M I/O Aggregator, Dell PowerEdge™ M620 blades, Dell EqualLogic™ Storage, Dell Force10™ network switches, and VMware vSphere 5.1. The solution also includes Dell PowerEdge™ R620 servers as management servers. VMware vCenter Server, EqualLogic Virtual Storage Manager for VMware, and OpenManage™ Essentials, are included with the solution.

Table 1 highlights the different Active System offerings.

Table 1 Active System 800v Configurations

	8 Blades Configuration	16 Blades Configuration	24 Blades Configuration	32 Blades Configuration
Compute capacity (Total number of cores)	128	256	384	512
Memory Capacity (in TB)	1	2	3	4
Storage Capacity (Raw TB)	43.2	86.4	129.6	172.8
Number of drives	(2x24) 10K 900GB 2.5" SAS drives	(4x24) 10K 900GB 2.5" SAS drives	(6x24) 10K 900GB 2.5" SAS drives	(8x24) 10K 900GB 2.5" SAS drives

As can be seen in the Table 1, Active System provides increasingly larger amounts of compute, memory, and storage capacity in moving from the 8 blade to the 32 blade configuration. This compute capacity can be used as a pool of compute and storage for any application deployment or shared infrastructure for deploying multiple application workloads.

For more information, please refer to the Active System 800v documentation.

SharePoint Server 2010 on Active System 800v

SharePoint is 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, publishing intranets, business intelligence 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. Some of the SharePoint services, or sets of services, comprise predefined roles, and must be configured within the solution. These predefined roles include server roles, such as Web frontend (WFE) and application (APP) service roles. Other services and components are optional, but they provide additional features and functionality that are often desirable. These optional components — deployed usually as a part of APP server role — may include service applications, such as managed metadata and Excel® services. Some constraints and best practices help determine which components should be located on each server in the farm, and appropriate distribution of the components help enable the SharePoint farm to easily accommodate later growth.

NOTE: In SharePoint Server 2010, components generally provide functionality for a given service application. As a result, this paper may use the terms "role" and "component" interchangeably. In this

context, SharePoint roles refer to one or more components that provide a farm service, and should not be confused with Windows Server roles, which generally include one or more Windows services to provide operating system functionality.

The size and capacity of a SharePoint Server 2010 implementation vary based on several factors, such as the number of concurrent users, service applications in the farm, the expected uptime service-level agreement (SLA), and others. These factors dictate how many servers are required in the SharePoint farm and how the overall farm architecture looks. Based on the these factors, SharePoint Server 2010 farm implementations can be classified as small[1] farm , medium[2] farm and large[3] farm deployments.

As discussed above, Dell Active System 800v provides the converged infrastructure for deploying any application in virtual containers. This infrastructure solution is an ideal choice for virtualizing the SharePoint Server 2010 farm deployment. The pre-engineered and pre-integrated infrastructure components and the underlying design principles for achieving high-availability and optimal performance help reduce the planning and deployment cycles required for the farm deployment.

This paper provides details on how the SharePoint Server 2010 farm was designed, what SharePoint farm roles were implemented, how the underlying converged infrastructure design principles map into the best practices for a SharePoint farm deployment, and, finally, discusses how the farm performed under varying load conditions.

For more information on Dell SharePoint reference architecture guidance, performance sizing and best practices, refer to the white papers available at http://www.dell.com/sharepoint.

Dell Services offering

Dell Services provides customers a portfolio of professional services including deployment, consulting and support. Deployment services for Active System 800 helps customers deploy and configure the solution ready to be deployed into their datacenter. Dell Consulting for SharePoint service is designed to assist customers in the evaluation, assessment, design, and implementation of an upgrade or complete migration to the newest version of SharePoint technology platform. Dell Consulting for SharePoint can help customers simplify their SharePoint environment, minimize deployment times and maximize the value of business information sharing through secure pre-architected and standardized deployment, administration and governance models. Dell ProSupport™ for SharePoint provides customers with a consolidated, single point of contact for their software support needs and ensures solution availability. ProSupport allows customers unlimited access to Dell's 24x7 chat, email and phone support services to help with configuration issues, how-to-assistance, critical patches, updates and disaster recovery.

---

[1] Designing and implementing a small SharePoint 2010 farm: http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~dell-small-sharepoint-farm.pdf.aspx

[2] Designing and implementing a medium SharePoint 2010 farm: http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~dell-medium-sharepoint-farm.pdf.aspx

[3] Designing and implementing a large SharePoint 2010 farm: http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~dell-large-sharepoint-farm.pdf.aspx

For detailed services information, visit dell.com and contact Dell Services representatives.

Designing SharePoint 2010 on Active System 800v

A SharePoint Server 2010 farm deployment includes three-tier architecture with different farm components deployed at each tier. The traditional approach of deploying a three-tiered application uses resources such as physical servers. This approach requires expensive physical hardware and complex management methodology, and a typical application deployment may not be able to fully leverage the increased processing and memory capacity of a current generation physical server. Virtualization technologies such as VMware vSphere let IT administrators consolidate multiple application workloads onto physical servers and achieve application infrastructure high availability, complementing application high availability features and flexibility. Virtualization results in additional benefits, such as reduction in infrastructure costs required to deploy the application farms and enable dynamic on-demand scalability of the application.

The design principles of Active System 800v infrastructure help virtualize a multi-tiered application such as SharePoint Server 2010 without compromising application best practices.

Active System 800v for SharePoint Server 2010

Active System 800v provides an optimal virtualization infrastructure for deploying SharePoint Server 2010 farm architecture. As described in Table 1, Active System 800v is offered in different configurations, each providing a pool of resources for implementing SharePoint Server 2010 farm configurations. The underlying design principles for Active System 800v storage and network architecture map into the application best practices and recommendations for deploying SharePoint Server 2010.

The networking best practices, such as traffic isolation using virtual LANs (VLAN) and traffic prioritization for lossless iSCSI data transfers, are the core attributes of Active System 800v architecture. The EqualLogic iSCSI storage employed in Active System 800v infrastructure provides optimal performance levels for an application such as SharePoint Server 2010.

A SharePoint Server 2010 farm requires deployment of SharePoint farm roles as different VMs and requires that these farm roles be separated to provide farm role redundancy and high availability. Dell Active System 800v employs VMware vSphere 5.1 for virtualizing the host resources and provides an optimal virtual environment for achieving the application role isolation best practices and recommendations imposed by a multi-tiered application such as SharePoint Server 2010. Using VMware Distributed Resource Scheduler (DRS) rules, resource pools, and resource reservations, the minimum resource requirements for each farm role can be guaranteed. Using the DRS rules, farm role separation onto different hypervisor hosts can be achieved.

The following sections describe best practices in virtualizing a SharePoint Server 2010 farm deployment, explain how the underlying Active System 800v infrastructure design principles help adhere to the virtualization best practices, and provide details on SharePoint 2010 application-specific best practices.

Virtualizing SharePoint Server 2010

Virtualizing a SharePoint 2010 application[4] deployment requires certain best practices be followed to achieve optimal application performance, to ensure high availability of the farm components, and to realize the additional benefits mentioned above. The three-tier architecture governs the best practices to be followed and requirements while deploying SharePoint Server 2010 in a virtualized environment. As mentioned earlier, Active System 800v provides the infrastructure that is most efficient for virtualizing application workloads.

This section describes the recommendations and best practices for virtualizing a SharePoint 2010 farm and how Active System 800v helps in achieving these virtual infrastructure and application best practices.

Designing SharePoint Infrastructure High Availability

SharePoint farm role availability must always be considered when planning for a SharePoint farm deployment. Different tiers of a SharePoint farm deployment offer different options for high availability. When virtualizing the SharePoint farm roles, the underlying hypervisor high availability options, such as VMware HA[5], can also be considered when planning for application high availability. A combination of application high availability features and the hypervisor high availability features can be used to enhance the overall availability of the SharePoint Server 2010 farm.

VMware vSphere offers features such as Distributed Resource Scheduling (DRS) to configure high availability of the virtual machine infrastructure. By using DRS rules, the VM placement and failover can be planned to achieve virtual machine affinity and anti-affinity. For example, when virtualizing SharePoint farm roles, the VMs hosting database roles in a highly available configuration such as a SQL failover cluster or SQL database mirror should never be on the same physical host. This can be achieved by using DRS rules such as “Separate Virtual Machines” to create VM anti-affinity. This ensures that these VMs hosting similar farm roles never come online on the same physical host, complementing the application level high availability.

Estimating CPU and Memory Requirements

SharePoint Server 2010 can be deployed in different farm sizes. Several factors impact the number of servers and the minimum amount of CPU and memory resources required at each tier of the SharePoint farm.

[4] Plan virtual architectures: http://technet.microsoft.com/en-us/library/ff607811(office.14).aspx
[5] VMware vSphere HA: http://www.vmware.com/products/datacenter-virtualization/vsphere/high-availability.html

Table 2 SharePoint Server 2010 Farm Role Minimum Hardware Requirements

	Web Frontend	Application	Database
Number of Processors	4 processor cores	4 processor cores	4 processor cores
Memory	8GB	8GB	16GB
Hard disk space	80GB for system drive	80GB for system drive	80GB for system drive

Table 2 provides only minimum required resources[6] for each farm role in a SharePoint farm deployment. The actual amount of resources required at each tier is highly subjective to a combination of the following factors:

Workload Usage Profile: Refers to the distribution of user requests in the SharePoint farm. This defines a mix of typical SharePoint activities performed by the end users. For example, in a SharePoint 2010 farm, there could be a subset of users reading wiki pages, uploading documents, and performing searches, etc. The workload usage profile in this context refers to the percentage of users performing each of the mentioned activities at any given point in time. The SharePoint activities in a SharePoint farm have an impact on the CPU and memory resources required at the SharePoint database role. There are several usage models based on what features are deployed and used in a SharePoint farm. This includes usage models for Business Intelligence (BI), Enterprise Content Management (ECM), Search, publishing, collaboration, and so on.

This guide focuses on the recommendations and reference implementation for the collaboration workload only. Refer to Table 6 for details on the test mix used in this reference architecture.

Concurrent users: Refers to number of users actively making requests at the same time. The average concurrency in a SharePoint farm is typically a small percentage of all SharePoint users in an organization. This metric has a direct impact on the amount of processing power and memory required at each tier of the farm. This measure is influenced also by the workload profile or usage model of a SharePoint deployment. The number of concurrent users at the peak load time and a daily average are ideal measures of the amount of resources required.

Requests per second (RPS): Refers to the average number of requests made by each active user per second. Once again, the average number of requests depends on the workload usage profile. Typically, when RPS is used to determine the amount of resources, the size and type of the requests are not considered. When combined with the content database size, the type of usage profile, including the average size of files stored, determines the size of each SharePoint request. The size of a SharePoint request has a direct impact on the network requirements at

[6] Hardware and Software requirements for SharePoint farm deployment: http://technet.microsoft.com/en-us/library/cc262485.aspx#section2

each tier of the SharePoint farm and the amount of memory required at the database tier of the farm. The memory requirement for the SharePoint database role also depends on the size of the content database. For example, it is recommended to allocate 32GB or more memory for the SQL database role in a SharePoint 2010 farm hosting up to 2TB of content database.

Overall, the factors defined above-Workload usage profile, concurrent users, and RPS, have an impact on the hardware resource requirements in a SharePoint farm. For example, the more concurrent users, the heavier usage profile, and the greater RPS, the more hardware resources needed for the SharePoint farm deployment.

In addition to the above factors, the memory and processing requirements at the application tier of the farm depend on the service applications deployed in the SharePoint farm. In medium to large SharePoint deployments, the service applications, such as the search service application, should be deployed on dedicated physical or virtual servers. Active System 800v provides the infrastructure to support the above requirements, such as dense processing and memory requirements to host an application such as SharePoint Server 2010.

It is recommended to ensure that the physical host used for virtualizing a SharePoint farm deployment has adequate physical resources to accommodate the farm roles without over-subscribing the underlying physical resources. For example, if the ESXi host has 16 physical processing cores, the total vCPUs assigned to the VMs running on the host should not exceed 16. Similarly, memory over-commit is not recommended on hosts where SharePoint farm roles are deployed.

VMware vSphere supports a Virtual Symmetric Multi-Processing[7] (SMP) feature that enables co-scheduling of virtual processors assigned to the same virtual machine. This ensures that the guest operating system can perform concurrent execution of application activities by using multiple processors. As mentioned earlier, over-subscription of underlying physical CPU cores may hamper the co-scheduling of virtual processors and eventually result in performance degradation.

Also, in a virtual infrastructure hosting multiple workloads, implementing CPU and memory reservation for each virtual machine, guarantees resource allocations. This also helps ensure that the VMs move to the right hosts with the appropriate amount of resources during a failover or vMotion.

It is also recommended to enable Non-Uniform Memory Access (NUMA) to avoid remote memory access by a virtual machine. VMware ESX is designed to take advantage of the underlying NUMA architecture, and allocates resources in such a way that the virtual CPU and virtual memory assigned to the virtual machine are always local. However, the benefits of this approach are highly dependent on the size of NUMA node. For example, this approach will be beneficial on a system with 2-socket 8-core processors (size of the NUMA node is 8) and 64GB memory only if the virtual machine is allocated with eight or less vCPUs and memory less than or equal to 32GB. This ensures that the virtual machine resource allocation need not span multiple NUMA nodes and results in better performance.

As mentioned earlier and shown in Table 1, Active System 800v provides several configurations with increasing number of processor cores and storage capacity. The above requirements and best practices for optimal application and infrastructure performance are an integral part of Active System 800v architecture.

[7] VMware Virtual SMP: http://www.vmware.com/pdf/vsmp_datasheet.pdf

The resource pools feature in VMware vSphere 5.1 provides reservations and shares. Shares allow priorities for certain VMs over others, should there be over-subscription of resources. It is recommended to avoid usage of shares in a production environment. Instead of using shares, the recommendation is to avoid any over-subscription of CPU and memory available from the host, thereby avoiding contention and eliminating the need for prioritized access. A reservation allows a certain amount of CPU and memory to a VM. This enables stable performance of the application workload by ensuring the farm role virtual machine always get a minimum amount of resources required to support the desired workload.

Finally, another option is to use memory locking, which guarantees the Microsoft-specified memory requirement at all times to the VM. The drawback of this approach is that it does not allow dynamic amounts of memory to be allocated to VMs to release memory for other uses during periods of low utilization.

Once again, the infrastructure components of Active System 800v, such as VMware vCenter, make it easy to manage the configuration options described above.

Designing SharePoint Networks

For optimal performance, the SharePoint farm roles should be placed to ensure the communication latency between the SharePoint farm roles is minimal (less than or equal to one millisecond). It is recommended to use dedicated and separate network connections for:

- SharePoint farm public access (end-user facing).
- Farm role communication, such as Web frontend and application to database server communication and SQL Server witness communication.
- Farm private communication, such as SQL cluster communication and/or SQL database mirroring traffic.
- Farm management traffic such as Active Directory® communication, backup and other management related traffic.

Active System 800v network design principles offer the traffic isolation as described above by providing virtual LANs (VLAN). In a converged infrastructure, where both LAN and storage traffic use the same physical link, VLANs provide an efficient means to isolate the traffic between different farm roles. In a virtual SharePoint farm, the traffic isolation is achieved using different vSwitches for each traffic class and using multiple vNICs for each virtual machine hosting the SharePoint farm roles. Apart from network isolation achieved using the VLAN implementation in a converged infrastructure, the appropriate bandwidth priorities must be configured to ensure the storage traffic gets the priority over normal LAN traffic. In a converged infrastructure used for virtualizing SharePoint farm roles or any other application, technologies such as Data Center Bridging[8] must be employed to ensure priorities among different traffic classes.

[8] Data Center Bridging overview: http://en.wikipedia.org/wiki/Data_center_bridging

Designing SharePoint Storage

In general, all content in a SharePoint farm gets stored in a SQL Server database and, therefore, the storage design for a SharePoint database backend has a great impact on the farm performance. When virtualizing a SharePoint farm, the virtualization layer adds overhead to the storage performance. This can be mitigated by considering factors such as the type of disk storage used for virtual machine OS deployment and the disk storage used for the SharePoint content backend. For example, using thick-provisioned or Raw Device Mapping (RDM) disks for virtual machine storage provides better performance than the thin-provisioned virtual disks.

Also, it is recommended to use multiple virtual SCSI controllers to connect the underlying virtual storage, such as the raw devices to the VMs. VMware supports up to four SCSI controllers per virtual machine and up to 16 disks connected to each SCSI controller. This model of distributing the disks across multiple SCSI controllers eliminates contention for the SCSI bus at the virtual machine level and, therefore, improves overall storage performance.

For the overall virtual infrastructure storage, it is recommended to use Storage Area Network (SAN) volumes to provide higher performance and scalability levels. When using SAN volumes, multiple paths to the storage must be deployed to ensure there is path redundancy. These redundant paths should be deployed to provide both single link failure and failure at the switching infrastructure. Using host multipathing ensures that these redundant paths are load-balanced and the traffic fails over to the active path in case of any link failures.

The Active System 800 storage design principles implement the best practices defined by the EqualLogic iSCSI storage. For general guidance on EqualLogic iSCSI deployment recommendations for VMware vSphere, see http://www.equallogic.com/WorkArea/DownloadAsset.aspx?id=10799

SharePoint Server 2010 Farm Design

Apart from the infrastructure design principles and best practices for deploying SharePoint, there are SharePoint Server 2010 application requirements and/or best practices that must be considered while deploying SharePoint farms of any size. As mentioned earlier, the multi-tiered application architecture of SharePoint Server 2010 dictates what best practices need to be followed at each tier. This section of the paper describes the best practices and requirements for a farm deployment to achieve optimal performance and high-availability at all tiers.

Web Frontend (WFE) load balancing

A multi-server SharePoint farm deployment generally consists of multiple WFE servers. These WFE servers need to be load balanced using either hardware or software load balancer. This ensures equal load distribution among the WFE server in a SharePoint farm and provides high-availability for the SharePoint Web frontend role. Hardware load balancers such as F5 Networks BIG-IP 3900 LTM or software load balancers such as Microsoft Windows Network Load Balancing feature can be used for SharePoint WFE load balancing. The choice of a load balancer depends on the expected load on the WFE in terms of number of concurrent users and the number of WFEs supporting the concurrent user load. Also, other factors such as the SharePoint 2010 workload usage profile also have a great impact on the WFE resource utilization, and therefore, affect the number of WFEs required in the SharePoint farm.

For example, supporting a large SharePoint Server 2010 farm with 5,000 concurrent users and above in a collaboration usage profile may require more than four WFEs. For such a farm deployment, it is

recommended to deploy hardware load balancers because the native Windows NLB is known to have issues supporting more than four WFE servers.

SharePoint 2010 Database Design

SharePoint relies on the SQL Server database backend used for storing the SharePoint content databases. Therefore, the SQL Server plays a bigger role and affects the overall farm performance and availability. To a large extent, the SQL Server performance depends on the backend storage characteristics and the database layout used for deploying SharePoint databases. It is recommended to use RAID10 for the storage backend hosting SharePoint databases. The SharePoint databases, such as content database, search database, etc. should always be placed on dedicated volumes.

SharePoint farm availability depends on the SQL server availability and, therefore, it is important to design the SharePoint database backend to be highly available. There are multiple ways to achieve this including SQL Failover clustering[9] and SQL Database mirroring. The choice of method used for SQL database depends on the level of redundancy required at the database tier. When using a SQL failover cluster, a virtual instance name or cluster name is used to access the SQL server, and shared storage is deployed for the SharePoint clustered instance. Therefore, the SQL failover cluster provides a server level redundancy for multiple databases.

In the case of database mirroring, dedicated storage is deployed for hosting the principal and mirror databases. These SQL instances are hosted on separate SQL server hosts. SQL mirroring provides database-level redundancy. With the presence of a SQL witness server, SQL mirroring provides faster failover time than a SQL failover cluster. The cost implications of using SQL failover clustering versus SQL database mirroring should also be considered while deploying SQL high availability.

NOTE: All SharePoint databases support SQL mirror. However, it is not recommended to mirror SharePoint databases such as user profile synchronization database and WSS usage databases.

For more information on best practices for deploying SQL failover clusters, see http://msdn.microsoft.com/en-us/library/ms189910.aspx#BestPractices.

For more information on SQL mirroring best practices and recommendations, see http://technet.microsoft.com/en-us/library/cc917681.aspx.

Designing the Service Application High Availability

Provisioning SharePoint might require planning for dedicated application servers. This design decision depends on the size of the farm and expected daily or peak concurrent user load. In a farm with relatively high degree of concurrency, it is recommended to deploy service applications such as search service on dedicated servers, to ensure the expected levels of performance. In SharePoint 2010, several service applications are made HA aware and, therefore, can be deployed on multiple systems. This ensures that the services provided by these service applications are available even in the case of a single system failure.

[9] SQL 2008 R2 Fail-over Clustering: http://msdn.microsoft.com/en-us/library/ms189134(v=sql.105).aspx

The above design principles for SharePoint application deployment are independent of the underlying infrastructure platform. Active System 800v design principles provide the necessary infrastructure elements to support the above application best practices, while providing the most efficient virtualized converged infrastructure for deploying a multi-tiered application, such as SharePoint Server 2010.

The following section demonstrates using a reference implementation for SharePoint 2010 on Active System 800vthe benefits Active System 800v brings to a SharePoint Server 2010 farm deployment.

Reference Implementation for 5000 users

As discussed above, several factors impact the design decisions of a SharePoint 2010 farm deployment. A combination of all these factors must be considered while designing physical topology for the SharePoint farm. The reference configuration in this experimental implementation was sized for 5,000 users in a SharePoint collaboration usage profile with a content database up to two terabytes (TB) in size.

NOTE: This reference implementation is not a SharePoint appliance deployment and as a shared virtualization infrastructure, there could be more than one application deployment on a single instance of Active System 800v platform.

Figure 2 SharePoint 2010 Farm Virtual Machines

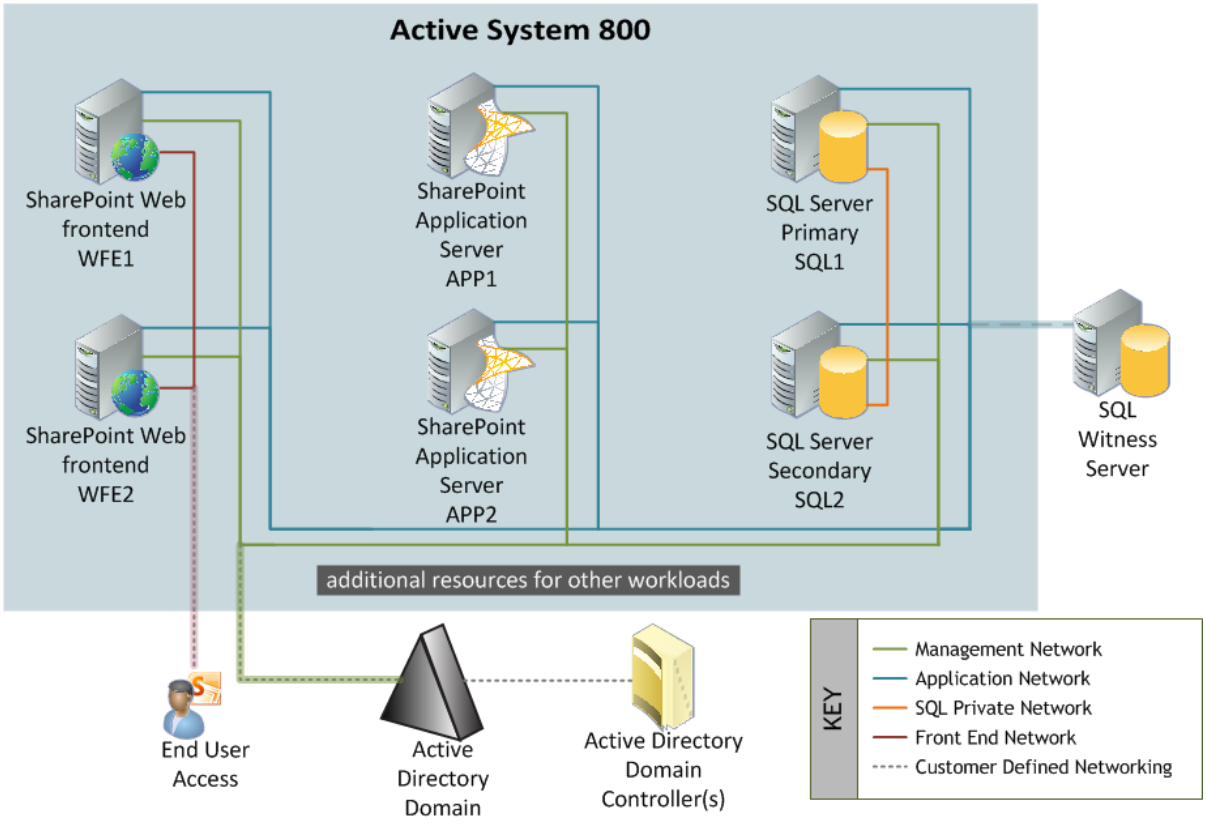

As discussed in the Designing SharePoint 2010 on Active System 800v section, the underlying Active System 800v design principles provide an efficient infrastructure to deploy a SharePoint Server 2010 farm.

The following sections describe how the SharePoint farm within this reference implementation was configured; the network and storage architecture implemented within this reference configuration; and how SharePoint-specific best practices were implemented.

The SharePoint farm reference implementation in this paper used Active System 800v configuration with eight ESX hosts – as described in Table 1. However, the entire compute capacity available with eight ESX hosts configuration was not used for deploying this SharePoint reference implementation. The virtual resources required for the SharePoint reference implementation were sized for 5,000 concurrent users and these requirements govern the number of physical ESX hosts used for SharePoint virtual infrastructure deployment. In this reference implementation, three EqualLogic PS6110x arrays were used – two arrays for hosting the SharePoint 2010 content databases and one for the virtual machine store. This indicates the minimum required storage capacity for a 5,000 concurrent SharePoint implementation described in this paper.

Virtual Infrastructure

As shown in Figure 2, six VMs were used to virtualize the SharePoint 2010 farm on Active System 800v infrastructure solution. The configuration in Table 3 of the SharePoint farm roles was derived after a series of test runs to understand the impact of SharePoint collaboration usage profile on a reference implementation as shown in Figure 2.

Table 3 SharePoint 2010 Farm Role Configuration in this Reference Implementation

	Web frontend	Application	Database
Number of VMs	2	2	2
Number of Processors	6 vProcs	4 vProcs	6 vProcs
Memory	16GB	8GB	64GB
Network	3x VMXNET3[10] vNIC	2x VMXNET3 vNIC	3x VMXNET3 vNIC
Storage Controller	LSI Logic SAS	LSI Logic SAS	LSI Logic SAS
Virtual Disk for OS	250GB Thick Eagerly Provisioned VMDK	250GB Thick Eagerly Provisioned VMDK	250GB Thick Eagerly Provisioned VMDK

The above virtual machine configuration describes the minimal resource configuration required to achieve the desired collaboration workload for up to 5,000 concurrent SharePoint users.

[10] VMXNET3 virtual adapters provide enhanced traffic management capabilities including multi-queue support, IPv6 offloads, and 10GbE support.

NOTE: The reference implementation described in this section is only a sample configuration implemented on Active System 800v and does not represent any resource limitations of the underlying infrastructure. It is possible to implement a much larger SharePoint farm on Active System 800v given the vast amount of compute as shown in Table 1.

The following sections describe how the SharePoint farm reference implementation leverages the underlying network and storage architecture defined by Active System 800v and the benefits this brings to the deployment of the SharePoint farm.

Network Infrastructure

As discussed earlier, SharePoint Server 2010 farm deployment best practices dictate that the networks associated with different farm role communication are isolated from each other to enable traffic isolation and security. Active System 800v infrastructure network architecture provides the network traffic isolation using VLANs. As a pre-integrated solution infrastructure, Active System 800v comes with the a few VLANs implemented for hypervisor management, storage, and virtual machine traffic. In addition to these VLANs, the SharePoint 2010 farm reference implementation described in Figure 2 implemented additional VLANs for separating the intra-farm role communication.

- Hypervisor Management VLAN
- VMware vMotion VLAN
- VM Network
    - *SharePoint farm Public access*
    - *SharePoint WFE-APP to SQL communication*
    - *SharePoint SQL to SQL communication*
- iSCSI data access VLAN

Figure 3 SharePoint 2010 Farm Network Architecture

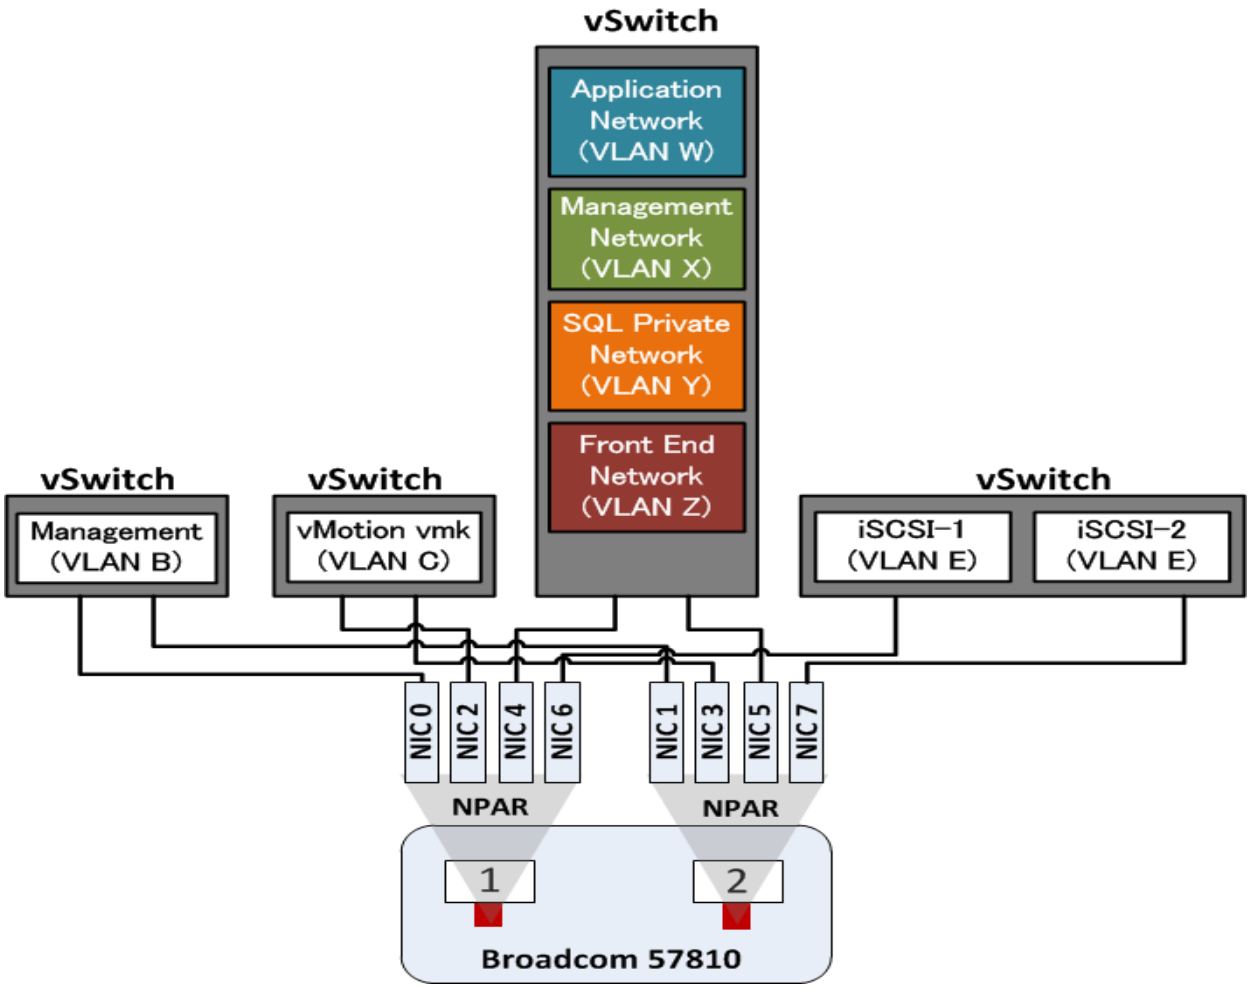

For more information on creation of custom VLANs for the SharePoint application, see "*Design and Implementation Guide for an Active System 800 with VMware vSphere*".

The traffic isolation between farm roles was enforced using the HOSTS file in the guest OS. For example, the SQL database mirror partners and the mirror witness server were allowed to communicate using the SharePoint SQL to SQL communication VLAN only by adding HOSTS file entries to reflect the IP addresses of these hosts in the respective subnet.

Active System 800v settings for Data Center Bridging were implemented without any modifications.

For more information on DCB Priority-Based flow control (PFC) and Enhanced Transmission Selection (ETS) settings, please refer to the Active System 800v documentation.

Storage Infrastructure

Active System 800v storage architecture provides the necessary VM space to store the SharePoint farm roles as VMs. The reference implementation in this paper leverages the underlying VM store implementation by Active System 800v. However, for the SharePoint 2010 databases and search query index storage, this reference implementation employed additional Dell EqualLogic PS6110x arrays.

A typical SharePoint Server 2010 deployment consists of databases such as:

- Configuration database
- Admin Content database
- Content database(s)
- Search Service Application databases such as Crawl Store, Property Store, Search Admin database
- WSS Usage database
- Any other databases for the configured service applications

The names of these databases can be customized and, therefore, this design and implementation guide refers only to the type of the databases instead of calling out the actual names of the databases.

Figure 4 SharePoint Server 2010 Databases

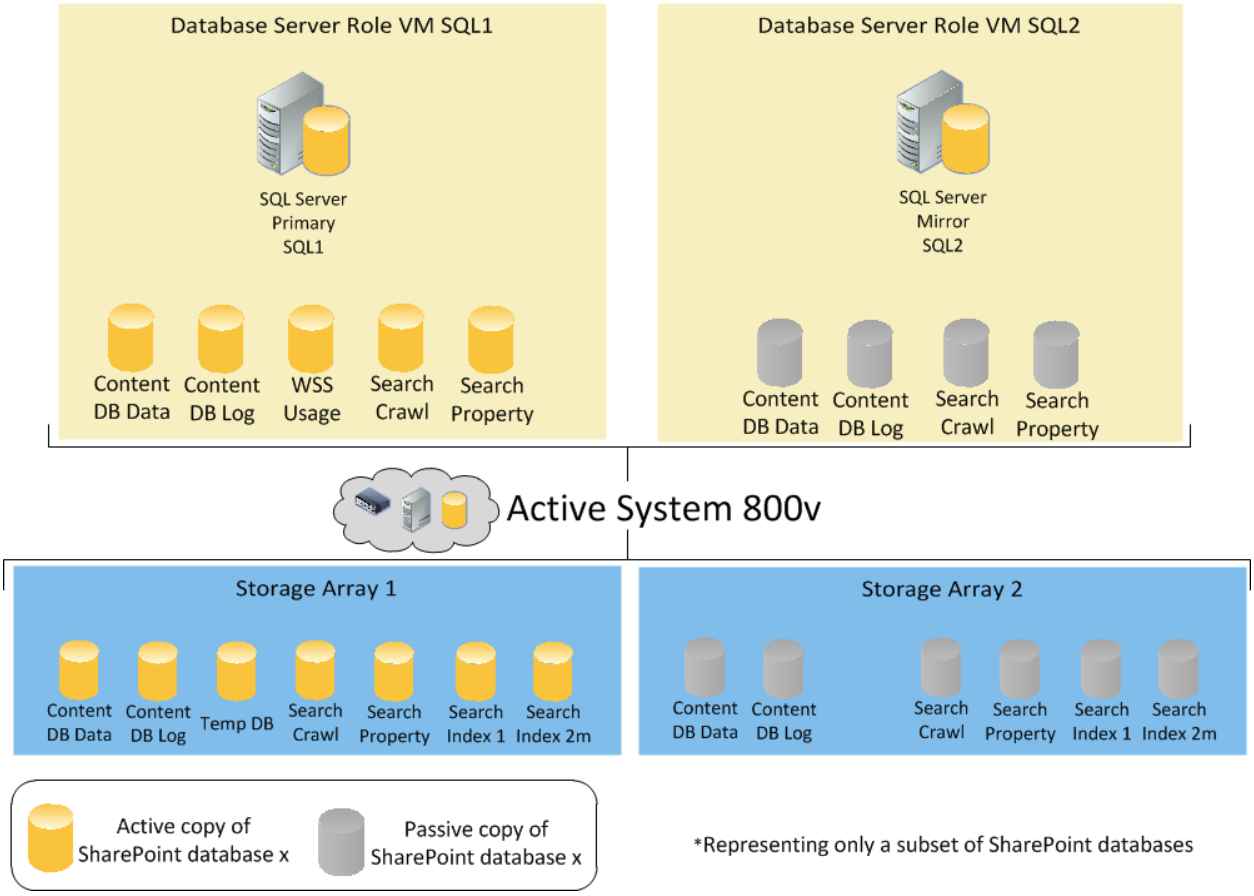

As a part of this reference implementation, only the search service application was deployed. Figure 4 illustrates only a subset of the SharePoint databases deployed in this reference implementation and how these databases are stored on the EqualLogic iSCSI storage backend.

Also, as part of the storage configuration, EqualLogic multipathing extension module (MEM) for VMware vSphere was implemented to support multiple iSCSI sessions, automatic load balancing of the iSCSI sessions, and improve the storage performance.

For more information on EqualLogic MEM, see
http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=3064

The storage arrays – 1 and 2 – as shown in Figure 4, were used to store the principal and mirrored copies of SharePoint databases. The active and mirrored copies of the search query index were stored on separate arrays to ensure high availability of search service application and query content.

As a best practice and a performance optimization measure, the iSCSI volumes from the storage arrays were mapped to the SharePoint farm role VMs as raw device mappings. For this purpose, multiple SCSI controllers were used to distribute the RDM devices.

For more information on mapping raw devices to VMs, please refer to the Active System 800v documentation for more details.

Figure 5 SharePoint 2010 Farm Role SCSI Controller Configuration

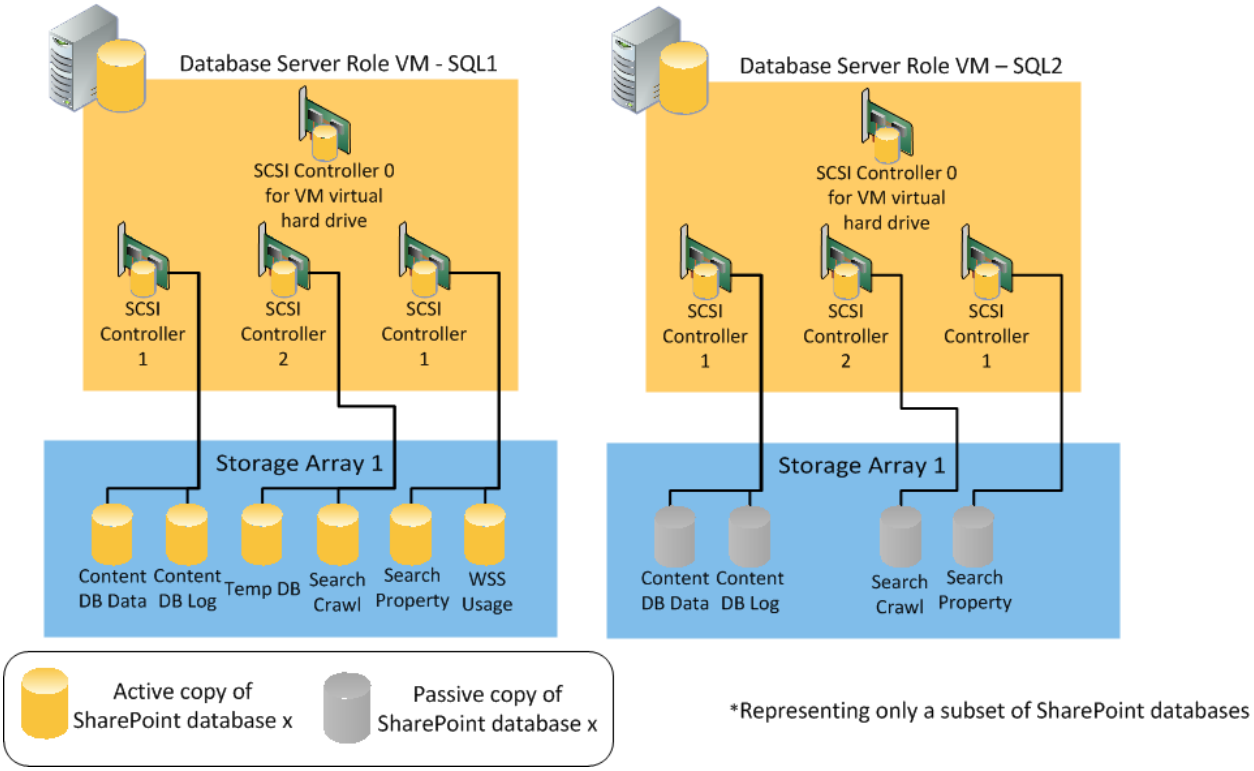

Figure 5 shows how some of the iSCSI volumes on the ESX hosts were attached to the principal SQL virtual machine role as raw device mappings. This approach ensures optimal performance at the storage tier with the virtual machine. Similarly, the same implementation was followed for the mirrored SQL instance and the application VMs hosting the search index and query volumes.

High Availability Configuration

The reference implementation in this paper leveraged the high availability features of the underlying virtualization platform to complement the application high availability features. One such implementation is the Distributed Resource Scheduler (DRS) rules. In this reference implementation, the HA or failover configuration was managed using DRS VM anti-affinity rules. These anti-affinity rules govern the virtual machine separation in the event of a failover. Within DRS affinity rules, the "Separate Virtual Machines" rule was deployed to achieve control over VM separation.

For example, in this reference implementation, the VMware DRS rules are used to achieve VM affinity or anti-affinity.

Figure 6 DRS Rules in the Reference Implementation

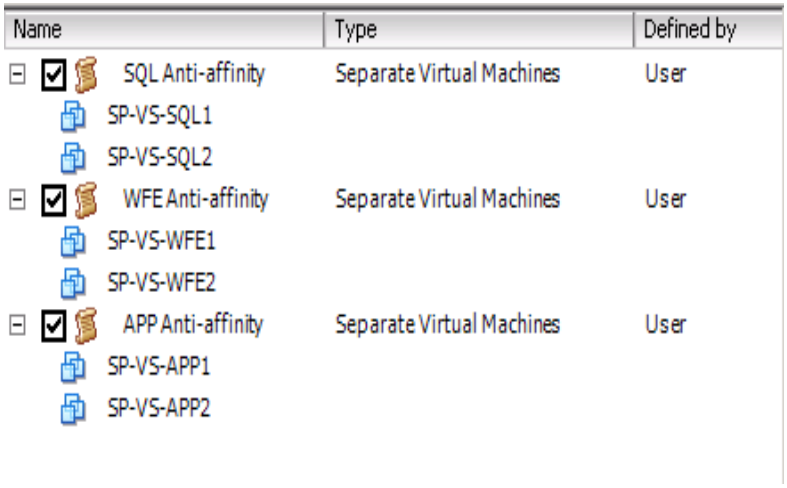

The above DRS rules in Figure 6 were implemented to separate the VMs hosting similar SharePoint farm roles.

SharePoint Server 2010 Farm

The SharePoint 2010 farm configuration, as shown in Figure 2 included two Web Frontend (WFE), two Application Servers (APP), and two database servers in a database mirror configuration. The application servers in the farm were used to host the search query and crawler roles. Table 4 lists the operating system and software editions used in the farm configuration.

Table 4 SharePoint 2010 Farm Software Configuration

	WFE and application servers	Database servers
Operating system	Windows Server 2008 R2 SP1 Enterprise Edition	
SharePoint server	SharePoint Server 2010 Standard Edition SP1	NA
Database server	NA	SQL Server 2008 R2 x64 Standard Edition SP1

NOTE: Step-by-step instructions for installing and configuring a SharePoint farm and any service applications used in this performance study are outside the scope of this paper. For more information and resources, refer to the References section of this paper.

Within this reference implementation designed for 5000 concurrent users, one SharePoint 2010 web application with several site collections was used. The content database size of the web application was approximately 2 Terabytes (TB).

Configuration of Web Frontend (WFE) Servers

This SharePoint 2010 farm design included two WFE servers. The software matrix for these WFE servers is shown in Table 4.

Within the scope of this paper, the Network Load Balancing (NLB) feature of Windows Server 2008 R2 was deployed. Using NLB, stateless applications such as SharePoint WFE servers are made scalable by adding servers when the load increases. In this reference architecture, the application server VMs were not a part of the NLB cluster. But, these application server VMs can be added when additional capacity at the WFE tier is desired. Therefore, NLB was configured to use multicast mode to avoid IP address conflicts[11] in the farm. To access the SharePoint NLB cluster from different IP subnets, an ARP entry for the NLB cluster name must be added on the LAN switch fabric.

For more information on implementing the NLB feature and configuring the NLB cluster, see http://technet.microsoft.com/en-us/library/cc732149.aspx.

Search service application configuration

SharePoint 2010 changed the search architecture and introduced high availability at the application tier (for instance, the search crawler). The new search service application architecture in SharePoint 2010 includes greater redundancy. This new design provides flexibility and lets the query and crawler roles be scaled-out separately on an as-needed basis. Search crawlers are now stateless; they do not store a copy of the index. However, the index does still propagate and is stored locally on the query servers. Two VMs hosting both crawler and query roles were used in the SharePoint 2010 farm.

[11] Using teamed adapters with network load balancing may cause network problems: http://support.microsoft.com/kb/278431

Figure 7 Search Service Application Architecture

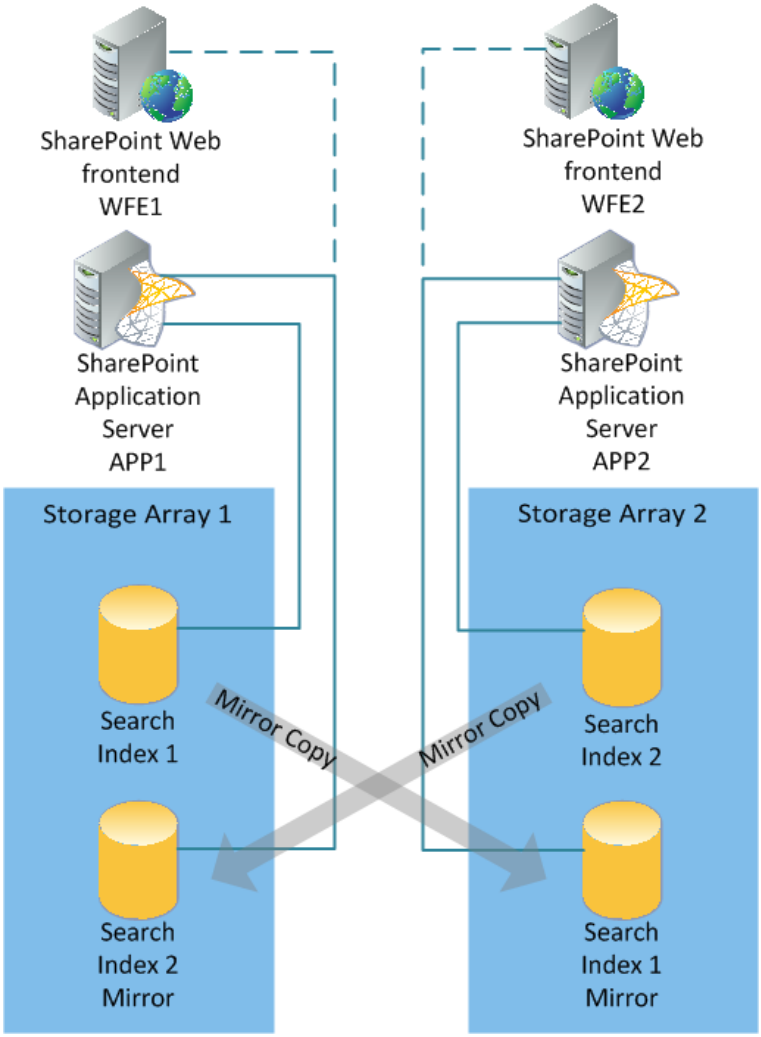

The farm configuration used the search service application configuration shown in Figure 7. The search crawler and query roles were hosted only on the application servers for improved search performance. Both the crawlers were associated with the same crawl database, and indexed the same content source.

The search architecture shown in Figure 7 is a logical representation of how the SharePoint Search service was configured. This configuration provides complete redundancy for both search crawler and query roles.

In Figure 7, the index partitions are represented as index 1 and index 2. The two primary index partitions were mirrored for redundancy; mirror, which is a copy of the index partition, was placed on each application server. On both of the application servers, VMware raw-device mapping was used to connect the volumes created on EqualLogic iSCSI storage arrays. Four volumes, two per application server were used to store each index partition.

For more information on creating and configuring Search Service Application in SharePoint 2010, see http://technet.microsoft.com/en-us/library/gg502597.aspx.

HTTP Request Throttling

SharePoint 2010 offers resource throttling features that are configured to help increase server performance and protect server resources during peak usage times. SharePoint 2010 has a default timer job that checks server resources compared to configured throttle levels. By default, Server CPU, Memory, Request in Queue, and Request Wait Time are monitored. After three unsuccessful checks, the server enters a throttling period and remains in this state until a successful check is completed. Requests that were generated prior to the server's entering throttling mode are completed. Any new HTTP GET and Search Robot request generates a 503[12] error message and is logged in the event viewer.

The throttle settings can be modified to increase the overall load supported by the farm servers. However, this requires a complete study to find accurate throttle setting recommendations for any given user load or requests per second. The default HTTP throttle monitor settings prevent an extensive load testing to find out the *real* capacity of the farm servers. As a result, HTTP request throttling was *turned off* during the load testing of SharePoint.

For more information on understanding and configuring resource throttling in SharePoint 2010, see http://technet.microsoft.com/en-us/magazine/gg491393.aspx

Configuration of the Database Server

As shown Figure 2, the SharePoint 2010 farm deployed two database servers in a mirroring configuration as VMs. A SharePoint farm's performance depends on the performance of the database server and the database storage backend. In this reference architecture, the storage resources on the database server were limited, so the SharePoint 2010 databases in this farm configuration were stored on the EqualLogic iSCSI volumes configured as VMware raw-device mappings inside the SQL server VMs.

In this reference implementation, along with the VMware HA feature, SQL database mirroring was used to enable application high availability. SQL mirroring requires a witness server for the completely automated failover. The SQL witness server deployment was external to the SharePoint farm virtual environment in this reference implementation.

As shown in Figure 2, the SQL witness server was connected to both the farm management VLAN and the WFE-APP to SQL application VLAN. This implementation ensures that the witness can still monitor the SQL mirror failures even in the event of a SQL server private communication network failure.

NOTE: Step-by-step instructions for configuring a SQL database mirror are outside the scope of this paper. For more information, see http://technet.microsoft.com/en-us/library/dd207314.aspx

The backend storage was implemented to store different SharePoint databases; separate iSCSI volumes were used to store each of these databases. Table 5 shows different SharePoint databases deployed in this reference implementation.

[12] Throttling starts alert- Events 8032 8062 - http://technet.microsoft.com/en-us/library/ee513044.aspx

Table 5 SharePoint 2010 Databases in the Reference Implementation

Database	Number of LUNs	LUN Size
1 x SharePoint Content Databases	1x for SQL Primary 1x for SQL Mirror	2TB
1 x SharePoint Content Logs	1x for SQL Primary 1x for SQL Mirror	200GB
Temp DB	1x for SQL Primary 1x for SQL Mirror	50GB
Search DB (Crawl, Property, and Admin)	1x for SQL Primary 1x for SQL Mirror	200GB
WSS Usage DB	1x for SQL Primary 1x for SQL Mirror	100GB
Other SharePoint Databases (Config and AdminContent)	1x for SQL Primary 1x for SQL Mirror	100GB

Memory configuration

By default, SQL Server service uses all available memory[13] as SQL Server dynamically grows and shrinks the size of the buffer pool, depending on the physical memory reported by the operating system. However, this behavior can be adjusted to limit the amount of memory used by SQL Server. Within the scope of this paper, SQL Server memory was limited to 80 percent of the physical memory available. For example, out of 64GB of memory on the database hosts, 52GB was allocated to SQL Server.

Reference Implementation Verification

Microsoft SharePoint 2010 is a versatile platform that is used in a large variety of ways. Some SharePoint workloads work almost out of the box, others require or allow significant customization, and still others are the result of completely custom-developed applications. This flexibility results in a multitude of ways of using SharePoint, which makes it almost impossible to accurately size servers and storage for a SharePoint farm. In addition, there is no standard benchmark for sizing SharePoint workloads at this time. It is very important to provide proper guidance to customers when it comes to recommending infrastructure elements of a SharePoint implementation. This led to the development of the Dell SharePoint Load Generation framework used to perform load testing of a SharePoint farm.

This internally-developed load generation framework[14] was used to understand the performance characteristics of the SharePoint farm. This framework includes load testing of SharePoint out-of-the box usage profiles, such as SharePoint collaboration, publishing, and search.

This section provides details on how the infrastructure built for the instance of SharePoint Server 2010 farm was validated to understand the steady and degraded state performance, how the infrastructure sustained the component failures, and how the DRS rules helped achieve high availability of the farm.

---

[13] SQL Server memory options - http://msdn.microsoft.com/en-us/library/ms178067.aspx
[14] How Dell does SharePoint load testing? http://en.community.dell.com/techcenter/b/techcenter/archive/2010/08/16/how-dell-does-microsoft-sharepoint-load-testing.aspx

SharePoint Server 2010 Infrastructure Verification

As a part of the SharePoint Server 2010 virtual infrastructure verification, several tests were performed to understand the validity of DRS rules and application / hypervisor level high availability configuration. The following sections show the different infrastructure verification tests performed and how the SharePoint 2010 farm roles implemented on Active System 800v get impacted during the failover process of the virtual machine roles.

DRS rules verification

The DRS rules, as described in the High Availability Configuration section of this paper, were used to ensure that the SharePoint farm roles hosting similar services never come online on the same ESX host. The virtual resources for each SharePoint farm role in this reference implementation were sized to support full farm capacity even during a failure or failover of some of the farm roles. The implementation of DRS rules assures that the VMs hosting similar SharePoint farm roles do not move to the same ESX host in the infrastructure.

To validate this DRS rules implementation, an ESX host failure was induced to understand how the VMs hosting SharePoint farm roles failover to other ESX hosts in the cluster. In this verification test, an ESX system hosting the WFE1 and APP2 VMs roles was powered down. This triggered a movement of these VMs (WFE1 and APP2) to other available ESX hosts in the virtual infrastructure. As a part of this verification, both VM failover and placement due to ESX host failure and the impact of VM failover during a load test on SharePoint connectivity was verified.

Figure 8 WFE Web Requests and Connection Failover During VM Movement

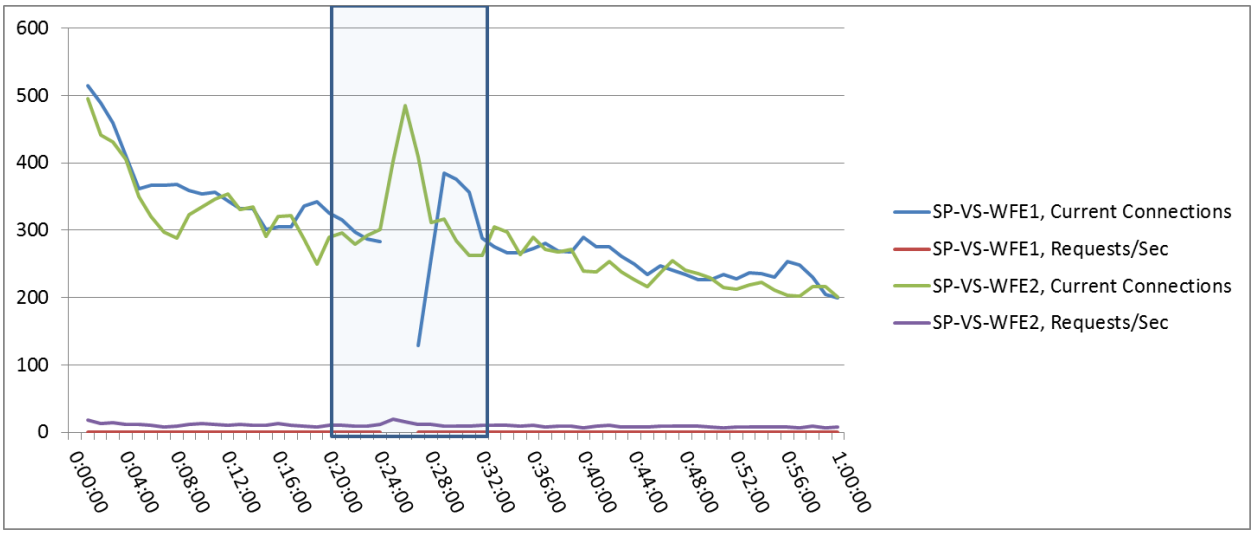

As defined in the DRS rules, the VMs hosting similar farm roles were moved to separate ESX hosts to complement the application HA features.

Also, as shown Figure 8, during the failover of WFE1 and APP2 VMs, WFE2 (spike in the green line) takes the additional SharePoint connections load while virtual machine with WFE1 farm role is being migrated to another ESX host. As soon as the WFE1 VM comes online, the NLB cluster starts distributing the requests evenly between all available frontend nodes in the SharePoint farm. This ensures that there is very minimal disruption in farm connectivity during a host failure.

SharePoint Database Server Failover

As a part of the SQL Database mirror failover verification, a one hour load test was performed and SQL principal server failure was simulated while the load test was in progress. This test showed that the SQL mirror instance was able to take over the principal role within a few seconds of principal failure and started to serve the incoming SharePoint user requests.

Figure 9 SharePoint Database Requests Transition While Mirror Failover

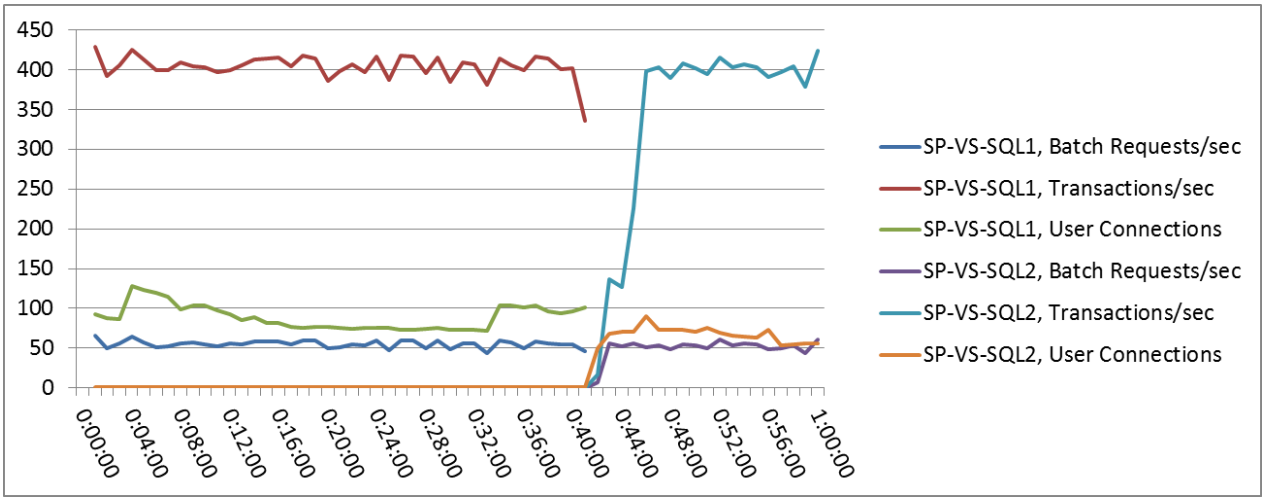

As shown in Figure 9, the simulated failure of SQL principal server triggered the database failover to the mirror instance. The SQL mirror instance becoming active can be clearly seen in Figure 9. During the SQL failover to a mirror database instance, the users may experience brief connectivity issues[15]. However, during the above verification test, it was observed that the mirror instance becomes active in less than 10 seconds, which is well within the thresholds imposed by SharePoint database connectivity best practices. This can be seen in the above chart, and once the farm connections to the mirrored database instance were restored, the load test continued without any errors.

SharePoint Server 2010 Performance Verification

As mentioned earlier, within the scope of this paper, a collaboration usage profile was used to perform load testing via the internally developed SharePoint load generation tool. Table 6 provides an overview of the test profiles used in this load test.

[15] User experience during a SQL database failover: http://technet.microsoft.com/en-us/library/dd207314(v=office.14).aspx#Section4

Table 6 Collaboration Test Mix

Action	Number of actions/hr/user
Read Site Home Page	6
Read Survey	2
Read Lists	2
Read Document Library	1
Read Wiki Page	1
Read Picture Library	1
Create Wiki Page	1
Upload Document	1
Search Site	3
Respond to Survey	1
Edit Wiki Page	1
Total tests/hour/connected user	20

The intent of this performance study was to understand the capacity of a SharePoint farm as shown in Figure 2 with the configuration described in Table 3. Several load test iterations were conducted with incremental user loads. For example, an initial user load of 500 virtual users was used and then incremented by 500 users until the farm resources reached an optimal level of usage.

The data set used to build the content database included several different types of files, such as Microsoft Office documents, Adobe PDF documents, and several image formats. Table 7 shows a distribution of file content sizes in each Web application used in this performance study.

Table 7 Data Set Used in This Study

Average File size	Number of files
1KB to 10KB	224122
10KB to 100KB	47235
100KB to 1MB	138262
1MB to 16MB	31517
16MB to 128MB	617
Greater than 128MB	12

The aggregated SharePoint content database size was around 2TB. During the load test duration, this content DB grew by almost 20 percent. A full content crawl was performed once at the beginning of the load tests. There were no subsequent crawls after load test or during the load test duration.

The following sections of this paper described the performance data and how several components within the farm performed at incremental user loads. As a part of this performance study, several performance metrics were collected and analyzed. This section describes how various components of the SharePoint farm performed under incremental user load. As mentioned earlier, this study included a SharePoint collaboration workload. All results shown here are relative to the workload used and may differ from any other implementation outside of the test mix shown in Table 6.

The load testing on SharePoint 2010 farm was performed in two different states of the farm:

1. **Steady state** describes the performance of the farm at the maximum available capacity. In this reference implementation, this implies the configuration of the farm as shown in Table 3.
2. **Degraded state** refers to the farm capacity reduced to a minimum number of validated farm roles. For example, within the degraded state validation, this reference implementation used only half the available farm roles, i.e., one WFE, one APP, and one DB server.

Table 8 Validated Farm Capacity

	SharePoint Farm Performance in steady state	SharePoint Farm Performance in degraded state
Maximum concurrent[16] user load validated	5000	2500
Requests per second[17] at Max concurrent user load	101/sec	47.3
Avg. Response time at the maximum validated user load	0.6 seconds	0.15

The above metrics indicate that the farm configuration used for this performance study could support faster, sub one-second farm response times even at the maximum validated concurrent user load. The degraded state verification was performed to understand the total capacity of the farm with reduced number of farm roles. Therefore, only a summary of the degraded state performance is described here and not a detailed analysis.

Table 9 Farm Performance Metrics in a Degraded State with Minimal Farm Roles

	SharePoint Farm Performance		
	WFE1	**APP1**	**SQL1**
Average Processor Utilization (%)	41.7	8.68	23.9
Avg. Disk Transfers / Sec (_Total)	511		
Avg. Disk Seconds / Transfer (_Total)	0.0017		
Avg. Disk Bytes / Transfer (_Total)	61688 (61KB)		

The metrics for the maximum validated concurrent user load in a degraded state – in Table 8 and Table 9 - show that the reference implementation in this paper was capable of supporting up to 2,500 concurrent users with the farm capacity reduced to half the actual implementation in terms of number of farm roles available.

[16] Concurrency refers to number of simultaneous requests to the farm servers

[17] This number indicates the avg. requests per second generate during the load test duration and this is a Visual Studio reported metric

The following charts show average farm response time for various user load iterations with full farm capacity or in steady state as described earlier.

Figure 10 Average Farm Response Time (in Seconds)

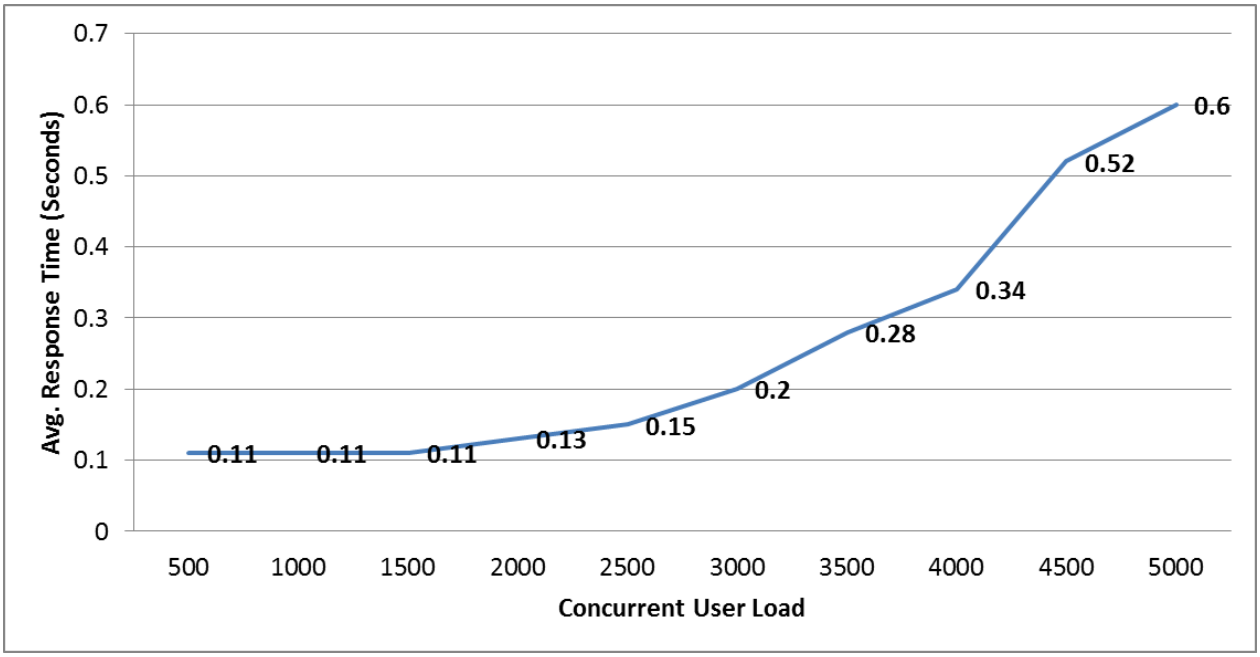

As seen in Figure 10, it is clear that the average farm response time was consistently less than a second and stayed almost flat at the lower concurrent user loads.

The average processor utilization – as shown in Figure 11 - on the Web frontend servers in the farm configuration was minimal even at the maximum validated user load. The following charts show the processor utilization metrics for several user load iterations in the SharePoint farm.

Figure 11 Average Processor Utilization Across Farm Servers

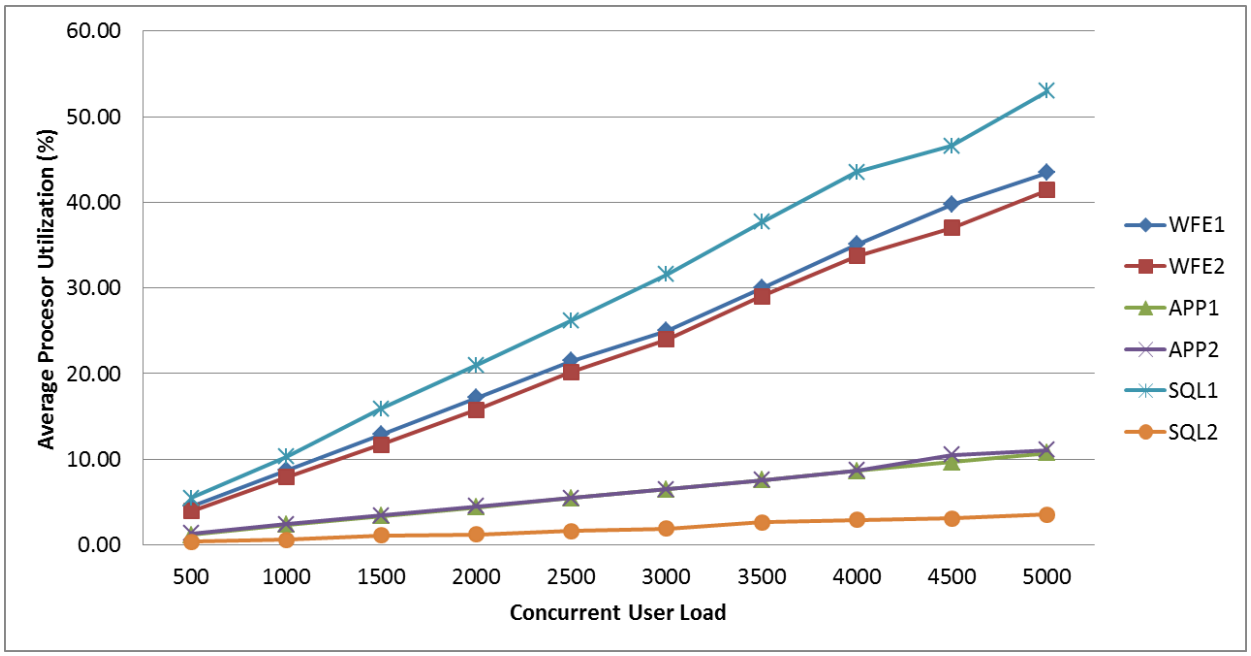

As shown in Figure 11, the processor utilization on the farm servers (APP1 & APP2) hosting application services, such as search service application, was lower than 10% even at the maximum validated user load. This may indicate that one application virtual machine is sufficient to support the concurrent user load shown in this paper. However, the second application virtual machine is required to support seamless high availability of the SharePoint farm services. Also, the processor utilization on the SQL mirror is attributed to the fact that the transaction logs get shipped to the mirrored instance and updated to the database. However, this utilization, as shown above, is very minimal.

In a heavily loaded scenario, even though the average processor usage is below 45-50 percent on the VMs hosting Web frontend farm roles, the SharePoint farm configuration in this reference implementation for 5,000 concurrent users may not support a larger number of users than what is shown in Table 8. This is mainly because of the ASP.NET and IIS request queue length limitations. The out-of-the-box IIS and ASP.NET queue length settings can be tweaked to go beyond the concurrent user load shown in this performance study paper. However, this is outside the scope of this paper and may require an in-depth study in itself.

The overall network usage on Web frontend and Application servers was very minimal and in the range of 100 Mbps. Considering the available throughput, the network usage on WFEs and APP servers was not a significant load at all. The Database server network utilization was also found to be very minimal.

The storage arrays used in this configuration for SharePoint content databases was found capable of supporting a 5,000 concurrent SharePoint collaboration workload. The total SharePoint content size was approximately 2TB. Also, as shown in Figure 4 ,all the other SharePoint databases such as Search, Usage Data, and SQL TempDB were also stored on the Dell EqualLogic storage arrays.

Table 10 shows the I/O read-write statistics and overall IOPS achieved at the maximum validated concurrent user load.

Table 10 Storage Performance Metrics

IO Performance metric	SharePoint Large Farm
Avg. Disk Transfers/Second (_Total)	~1223
Avg. Disk Writes/second (_Total)	~718
Avg. Disk Reads/Second (_Total)	~505
Avg. Disk seconds/Transfer (_Total)	0.001 or 1ms
Avg. Disk Queue Length (_Total)	1.90
Avg. Disk Bytes/Transfer (_Total)	110863 (108KB)

The above storage metrics indicate that the workload had an I/O mix of ~42 percent reads and ~58 percent writes with an approximate I/O size of 108KB.

In this reference implementation, a dedicated Dell PowerEdge PS6110x EqualLogic array was used for storing the principal copy of the SharePoint content database. The Dell EqualLogic PS6110x array with 24x 10K SAS drive backend is – theortically – capable of providing approaximately 2,160 IOPS in a RAID 10[18] configuration. In this reference implementation, the IOPS generated - approximately 1,223 IOPS - at the maximum verified user load is well below the theoretical maximum of the underlying storage. This showed that the Dell EqualLogic array was capable of handling SharePoint collaboration workload of 5000 concurrent users with enough room to accommodate the future growth of the SharePoint farm.

Conclusion

A SharePoint 2010 farm consists of multiple servers, each provisioned with different SharePoint components. There are several design and best practices to consider while building the infrastructure for a SharePoint farm. These design considerations include aspects such as high availability of all SharePoint farm roles, service application architecture considerations, and storage and networking best practices. Also, virtualizing an application like SharePoint requires additional planning and an approach to implement and map the infrastructure elements to suit application best practices.

The Dell Active System 800v virtualized infrastructure has been designed to reduce the time-to-value and enable rapid application deployment. This infrastructure platform is designed to help the application realize the benefits of a converged infrastructure and leveraging the underlying design principles for deploying a SharePoint farm derives several benefits. This includes rapid time-to-value for SharePoint application, high availability at all tiers of the SharePoint farm, infrastructure capacity that can accommodate any future needs of a SharePoint farm, and of course, a virtualized infrastructure that is pre-designed, pre-validated, and pre-built with several networking and storage best practices

[18] Assuming that each 10K SAS drive is capable of approximately 180 IOPS

Overall, the sample reference implementation of SharePoint Server 2010 farm for 5,000 concurrent users in a collaboration usage profile performed well on Active System 800v infrastructure with an average response less than one second, and other resource utilization levels, such as processor utilization, were well within the expected threshold.

References

- Dell solutions for SharePoint: Http://dell.com/sharepoint
- Dell Data Center Systems Management Solutions: http://content.dell.com/us/en/enterprise/data-center-system-management
- Microsoft SharePoint page: http://sharePoint.Microsoft.com
- Microsoft SharePoint TechNet: http://technet.microsoft.com/en-us/sharepoint/default.aspx
- SharePoint 2010 operations framework and checklists: http://technet.microsoft.com/enus/library/gg277248.aspx
- SharePoint 2010 health monitoring: http://technet.microsoft.com/en-us/library/ee681489.aspx
- Consulting Services: http://content.dell.com/us/en/enterprise/by-service-type-it-consulting-communication-collaboration.aspx
- ProSupport: http://content.dell.com/us/en/business/d/services/smb-prosupport-for-software-us