



# Reference Architecture - Microsoft SharePoint Server 2013 on Dell PowerEdge R630

A Dell reference architecture for 5000 Users

Dell Global Solutions Engineering  
June 2015

THIS REFERENCE ARCHITECTURE 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.

**Copyright © 2015 Dell Inc. All rights reserved.** This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



# Table of contents

Executive summary .....	4
1 Introduction .....	6
1.1 Scope .....	6
1.2 Audience .....	6
1.3 Solution components .....	6
1.3.1 Dell PowerEdge R630 .....	7
1.3.2 Dell EqualLogic PS6210 .....	8
1.3.3 Dell Networking S4810 .....	8
1.3.4 SharePoint Server 2013 .....	8
1.4 Design Considerations .....	9
1.4.1 High Availability .....	9
1.4.2 Application Performance and Best Practices .....	10
1.4.3 Hardware Abstraction .....	10
2 Solution architecture .....	11
2.1 Hyper-V cluster architecture .....	11
2.1.1 Storage architecture and configuration .....	13
2.1.2 Network architecture and configuration .....	14
2.2 Application architecture .....	15
2.2.1 Storage architecture .....	18
2.2.2 Network architecture .....	21
2.3 Virtualization rules .....	23
3 Verification .....	25
3.1 Methodology .....	25
3.2 Results .....	27
4 Solution specifications .....	29
5 Solution summary .....	31
A Additional resources .....	32



## Executive summary

Heterogeneous hardware, hardware sprawl, and a proliferation of various system management tools lead to inefficient infrastructure. The impact can include increased total cost of ownership (TCO), increased system maintenance downtime and slow application performance, which in turn affect business results. Many of these problems can be alleviated through hardware consolidation and virtualization. The processing power, memory, and 1U form factor of the Dell™ PowerEdge™ R630 server make it an efficient and powerful rack-based building block for virtualizing enterprise applications, such as Microsoft® SharePoint® Server 2013.

SharePoint Server 2013 is an enterprise platform that provides a rich set of features enabling end-user collaboration, web portals, electronic content management, e-Discovery, business intelligence, process flows and other customized applications. SharePoint Server 2013 is designed to run on dedicated physical hardware with best practices for performance and availability. (For more information, refer to Microsoft and partner documentation.)

With the widespread adoption of virtualization in the data center, Microsoft provides best practices for deploying 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. While this practice ensures appropriate performance and application availability during hardware issues, it also leads to creation of hardware resource silos dedicated to specific applications, resulting in management complexity for IT administrators.

Virtualizing SharePoint results in greater consolidation of application resources and reduces the complexity involved in deploying and managing application infrastructure. The overall benefits of deploying applications on a single virtualized infrastructure are:

- Simplification of deployment, configuration, maintenance and administration of hardware infrastructure and applications
- Single view of server, storage and networking for multiple applications
- Reduced hardware cost and administration time
- Optimized application uptime and enhanced business continuity by complementing application native availability with hypervisor-level and infrastructure-level high availability (HA)
- Greater IT efficiency by eliminating infrastructure silos designed for individual applications

The study in this paper describes the reference architecture for Microsoft SharePoint Server 2013 solution leveraging Microsoft Windows Server 2012 R2 Hyper-V® with PowerEdge R630 servers. This reference architecture is designed to leverage the benefits of virtualizing the underlying infrastructure. The infrastructure HA design ensures that there is no single point of failure of any component within the architecture. Certain percentages of PowerEdge R630 resources are reserved for the solution, while the



remaining resources can be leveraged for running additional applications – which is at the administrator’s discretion.

This paper is structured in the following sections:

- [Section 1](#): Describes the key components used in the solution reference architecture. These components are engineered to make the solution complete and effective. Understanding the components is essential for comprehending the reference architecture. This section also describes the design considerations for virtualizing Microsoft SharePoint and the benefits achieved.
- [Section 2](#): Describes the actual reference architecture leveraging Hyper-V on Dell PowerEdge R630 and provides information about how each component of the architecture is implemented according to the considerations described in Section 2.
- [Section 3](#): Provides a summary of verification performed and an analysis of the results gathered from performance verification.
- [Section 4](#): Provides solution specifications and additional information about components used in this architecture.
- [Section 5](#): Describes a summary of the overall solution and summarizes the benefits achieved by deploying SharePoint Server solution on PowerEdge R630.



# 1 Introduction

A typical SharePoint Server 2013 farm deployment includes a three-tier architecture with different solution components such as server systems, storage arrays, and networking equipment deployed at each tier. The traditional approach of deploying a three-tiered architecture uses resources such as physical servers. This approach requires expensive physical hardware and complex management methodology. Moreover, 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 Hyper-V let IT administrators consolidate multiple application workloads onto physical servers and achieve application infrastructure HA. Virtualization results in additional benefits such as reduction in infrastructure costs required to deploy the application farms and dynamic on-demand scalability of the application.

Apart from server virtualization, these application farms require shared storage that can provide the capacity and performance to support the application deployment and network equipment such as top-of-rack switches for providing an optimal network architecture for application connectivity.

This section describes the components deployed to support the SharePoint Server 2013 farm architecture designed for 5,000 concurrent users. The design considerations for architecting the farm for optimal performance, high-availability, and application best practices are also discussed.

## 1.1 Scope

This paper explores implementing Microsoft SharePoint Server 2013 on PowerEdge R630. This paper concentrates on infrastructure design principles, which include HA, application best practices, hardware abstraction and resource consolidation. A sample implementation for up to 5,000 users is given using these design principles.

## 1.2 Audience

This paper is intended for IT professionals and administrators interested in designing and deploying a Microsoft SharePoint solution for up to 5,000 users on PowerEdge R630 with Hyper-V. The paper provides an overview of the components in this solution architecture. However, the reader is expected to have sufficient understanding of Hyper-V, SharePoint 2013, PowerEdge R630, Dell EqualLogic™ PS6210, and Dell Networking 10GbE switches.

## 1.3 Solution components

The solution architecture described in this paper employs PowerEdge R630 servers for virtualizing SharePoint farm roles and EqualLogic PS6210x arrays for the shared storage required for the Hyper-V cluster and the application data. Dell Networking S4810-ON switches are used to connect physical servers and storage arrays to Local Area Network (LAN) and Storage Area Network (SAN).



The following sections provide a brief overview of components in the solution architecture described in this paper.

### 1.3.1 Dell PowerEdge R630

PowerEdge R630 is a 1U rack server from Dell that supports up to 36 processing cores and up to 768 GB of physical memory. Increased compute capacity makes PowerEdge R630 the most suitable platform for virtualizing enterprise applications such as SharePoint and other applications that perform compute-intensive tasks. A wide variety of supported network controllers provide the Local Area Network (LAN) and Storage Area Network (SAN) connectivity required for the virtualized infrastructure and applications. The enhanced systems management capabilities make it much easier to deploy, manage, and monitor the server infrastructure and get insights into the performance and utilization of the server in real-time.

Table 1 describes the components supported in a PowerEdge R630 server.

Table 1 PowerEdge R630 high-level technical specifications

	Supported in PowerEdge R630
CPU	Up to 2 Intel E5-2600v3 family processors
Memory	24 DIMMs; up to 768GB
Networking	QLogic or Intel 1GigE or 10GbE NDC Optional QLogic or Intel 1GigE or 10GbE PCIe Add-on NICs
Form Factor	1U rack
PCIe slots	3 slots: Two x16 slots with x16 bandwidth, half-height, half-length and One x16 slot with x8 bandwidth, half-height, half-length OR 2 slots: One x16 slot with x16 bandwidth, full-height, 3/4 length and One x16 slot with x16 bandwidth (or x8 with one processor only), half-height, half-length
Internal Storage	Up to 28 TB; 10 x 2.5-inch SAS or SATA drives (with no CD / DVD drive)  Up to 9.6 TB; 8 x 2.5-inch SAS or SATA drives (with CD / DVD drive)  Up to 22.5 TB; 24 x 1.8-inch SAS or SATA drives (with no CD / DVD drive)
Storage – RAID	PowerEdge RAID Controller H330 PowerEdge RAID Controller H730 PowerEdge RAID Controller H730P PowerEdge RAID Controller H830



	PowerEdge S110 on-board controller
Systems Management	iDRAC8 Express or iDRAC8 Enterprise

### 1.3.2 Dell EqualLogic PS6210

The EqualLogic PS6210X iSCSI storage arrays provide the high performance and dense storage required in this solution. Each array supports up to 24 x 10K SAS drives and up to 28.8 TB raw disk capacity. The dual controller configuration in the PS6210X arrays provides 16 GB non-volatile memory per controller and two 10GbE SFP+ for fiber or twin-ax copper or two 10GBASE-T with RJ-45 cabling per controller. Each controller in the storage array provides a 100BASE-T management network connectivity. In this reference architecture, two EqualLogic PS6210X arrays are used for Hyper-V VM store and SharePoint application data.

### 1.3.3 Dell Networking S4810-ON

The Dell Networking S-Series S4810 is a 10/40GbE top-of-rack (ToR) switch purpose-built for applications used in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact S4810 design provides 48 dual-speed 1/10GbE (SFP+) ports as well as four 40GbE QSFP+ uplinks, which free up valuable rack space and also simplify the migration to 40Gbps in the data center core. In the solution architecture, two S4810 switches are deployed as LAN switches and two S4810 switches are deployed as SAN switches in a non-converged network design.

### 1.3.4 SharePoint Server 2013

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. This can be represented in a three-tier architecture with different servers, as follows:

- **Front End Servers:** SharePoint Server is installed on these servers to provide end-user access through a hardware or software network load balancer.
- **Application Servers:** SharePoint Server is installed on these servers; additional SharePoint application services, such as search indexing, are configured.
- **Database Servers:** A highly available SQL Server is installed on these servers to provide storage for the SharePoint application. This design utilizes mirroring to provide high-availability; therefore, an additional SQL witness server is needed.

Unlike other applications such as Exchange Server and Lync Server, SharePoint Server does not have roles that can be installed separately during deployment. Instead, for Front End servers and Application servers, the SharePoint Server application is installed. The role of a SharePoint Server changes based on the





specific services for which it is configured. In a typical multi-tiered SharePoint farm, Front End servers are configured for client access, and Application servers are configured for providing SharePoint application services, such as Search. The SharePoint database is a SQL Server deployment that is connected to the Front End and Application servers. The Front End access has a network load-balanced path to the SharePoint users through either a hardware or software or virtual load balancer. The Application server has search services configured and runs the search crawl. [Section 1.4](#) describes the design considerations for virtualizing SharePoint Server 2013 applications.

## 1.4 Design Considerations

The details of design principles for each individual application and each layer of the infrastructure are listed below, followed by their impact to the SharePoint servers.

For virtualizing any application, the underlying physical infrastructure should be built with a number of key design principles, such as infrastructure HA, virtualization and networking best practices. The following additional principles must be considered for the deployment and configuration of SharePoint Servers. These additional design principles drive every aspect of the application design on a Microsoft Hyper-V cluster built using PowerEdge R630:

- High Availability
- Application best practices and performance
- Hardware abstraction

Deploying HA of the solution is critical as a principle to ensure minimal business downtime. HA must be considered at each layer of the solution, including the infrastructure, VM and application. Installing the application correctly and ensuring that the performance of the application remains above appropriate thresholds even during degraded scenarios is necessary for the business to function properly.

One of the goals of this virtualized infrastructure is to abstract hardware from the applications that run on that hardware. This hardware abstraction reduces administrative complexity, but it must be done while ensuring that the HA and application performance principles are not compromised. Finally, consolidating VMs, storage arrays, network load balancing, and databases -- while maintaining application best practices, reduces resource cost and administrative time.

### 1.4.1 High Availability

HA must be considered at every layer. The three layers of HA are:

- Infrastructure availability: In this reference architecture, if a hardware resource fails, other resources such as server, network path, storage drive, storage array, and so on are available to provide the same functionality.
- Hypervisor availability: During hardware failures, the application VM backup is moved to other physical hosts.
- Application-level availability: The pools of application servers are utilized to provide a highly available service.

If, for example, an infrastructure component such as a server fails, infrastructure availability provided by the Hyper-V cluster built on PowerEdge R630 servers ensures that there is an additional hardware resource, and the application-level availability keeps the application service up for end users. However, because fewer VMs are running the application, these services run in a degraded state. To reduce the time an application is in a degraded state, hypervisor availability automatically moves VMs to other available physical hosts. These two levels of availability maximize end-user access to the application service. In addition, the application data should sustain a storage failure. The EqualLogic storage arrays used in this reference architecture provide redundancy at various levels within the array. For example, each of these storage arrays is configured with redundant power supplies and redundant controllers connected to redundant top-of-rack (ToR) switches. This design ensures that component level failures within the array do not impact application availability. In this reference architecture, there are no redundant storage arrays to provide HA after an array-level failure. This design considers only a component-level failure within the storage arrays.

### 1.4.2 Application Performance and Best Practices

In addition to HA, application performance is essential for a good end-user experience. Following application best practices helps to prevent performance bottlenecks, and validating and verifying performance under load generation is key to developing performance recommendations. Several best practices and performance recommendations are listed here:

- Do not use dynamic memory for memory-intensive applications.
- Maintain the virtual CPU to physical core ratio at one-to-one.
- Allow no coexistence of critical application services on the same host.
- Use dedicated hard disk drive spindles for certain application-critical data, as specified by best practices. The use of fixed VHDX instead of pass-through disk mapping is recommended for optimal performance and ease of management for the applications that require external storage for data.

### 1.4.3 Hardware Abstraction

Abstracting specific hardware from application services reduces complexity and ensures that the application is load-balanced across the cluster. Applications can be decoupled from their underlying hardware, as long as best practices from the infrastructure and application layer are maintained. VM failover allows the application VMs to be moved around from server to server within the cluster for maintenance or in the event of server failure. However, certain application VMs cannot co-exist amongst themselves or with other application VMs for performance reasons. In such cases, rules can be created so that certain VMs do not co-exist on the same physical server with other VMs.

## 2 Solution architecture

This section describes the key configurations for the solution architecture. The solution architecture proposed in this paper for SharePoint 2013 incorporates the application best-practices and design principles described in [section 1.4](#).

The design goals for the application architecture are:

- Providing infrastructure HA for all the application VMs, while leveraging the application HA features
- Abstraction to decouple applications from the underlying hardware

The common design goals are implemented with the help of virtual infrastructure design principles and the hypervisor features. The reference implementation in this paper maps the design goals—both common and application-specific—to a set of specific hardware and software resources.

### 2.1 Hyper-V cluster architecture

The Hyper-V infrastructure design in this reference architecture employs the best practices for SharePoint virtualization to provide HA and optimal performance for the application Virtual Machines (VM). In this reference architecture, a single Hyper-V cluster is created using four PowerEdge R630 servers. All the SharePoint workload VMs are deployed in this Hyper-V cluster. Optionally, management software such as Microsoft System Center suite can be deployed in the same cluster to manage the VMs and applications deployed in the VMs. A high-level overview of this is shown in [Figure 1](#).

Table 2 Hyper-V cluster node configuration

Resource	Description
Compute Nodes	4 x PowerEdge R630 servers
Processors	2 x Intel® E5-2600v3 Family processors in each PowerEdge R630 server
Memory	128 GB in each PowerEdge R630



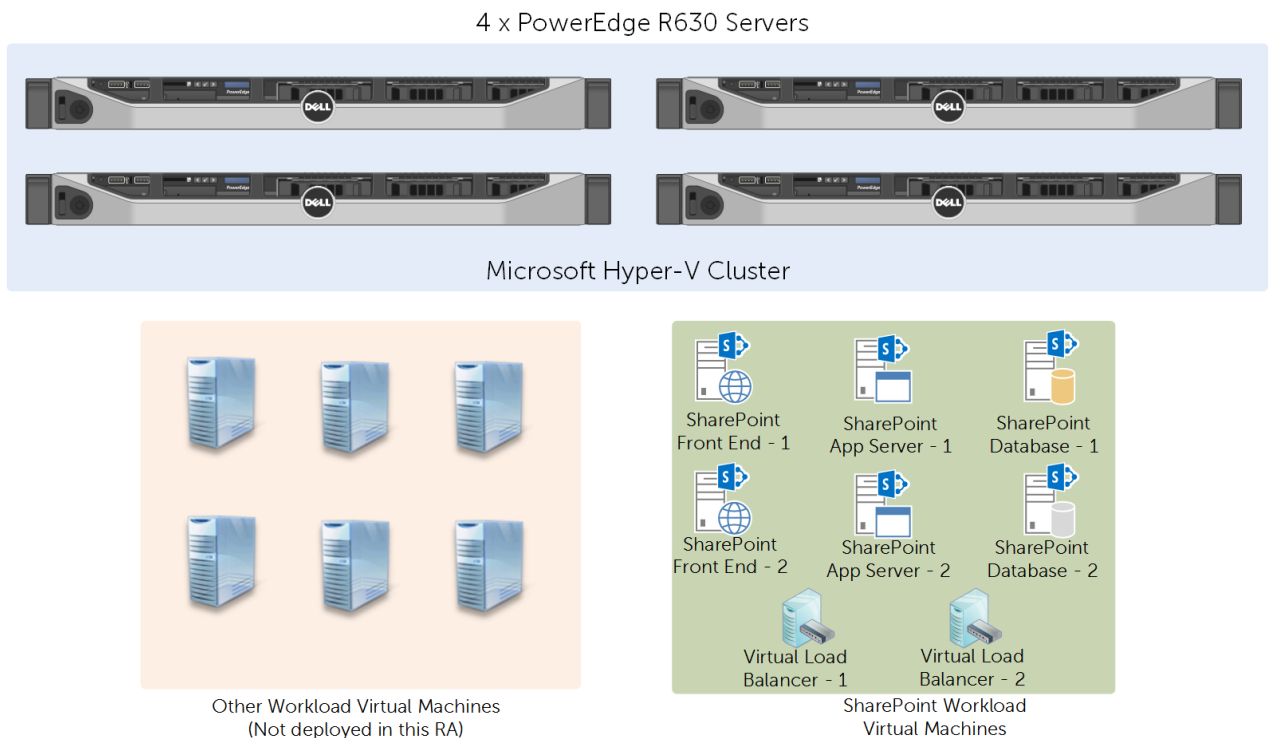


Figure 1 Hyper-V cluster on PowerEdge R630

Because this solution architecture deploys four PowerEdge R630 servers in a Hyper-V cluster, it provides more capacity than what is required by the SharePoint solution for a 5,000 concurrent-user deployment with an average farm response time of less than one second. The additional capacity can be used to deploy other workloads and/or management virtual machines in the same cluster. Also, for the SharePoint farm virtual machines, this solution architecture provides resiliency for up to two-node failure in the cluster without degradation in the farm capacity.

Windows Server 2012 R2 Hyper-V increases operational efficiency by enabling server virtualization. With Windows Server 2012 R2, Microsoft has introduced a number of improvements that allow customers to take advantage of newer server, storage and network hardware technologies.

The key benefits of this reference architecture to Microsoft SharePoint solution are:

- Efficient application management and best practices achieved through hardware abstraction at various levels of the infrastructure, such as storage, network and server.
- Complemented HA for the application servers
- Lower TCO vs. with traditional deployments.

As mentioned earlier, this reference architecture includes PowerEdge R630 servers hosting a virtualization solution based on Hyper-V with Dell Networking 10GbE switches as a network backbone and the EqualLogic PS6210X as a storage area network (SAN). [Table 2](#) lists the compute resources used in this reference architecture.

### 2.1.1 Storage architecture and configuration

This reference architecture employs two EqualLogic PS6210X--one used for the VM data store and SharePoint Search Index 1, and the other used for the SharePoint database and Search Index 2 storage. Each of the EqualLogic PS6210X includes 24 x 1.2 TB 10K SAS drives in a RAID 10 configuration. These storage arrays connect to the SAN using 10GbE SFP+ connections. As a part of the infrastructure design, the SAN and LAN fabrics are separated. The SAN fabric uses Dell Networking 10GbE switches for the top-of-rack (ToR) connectivity.

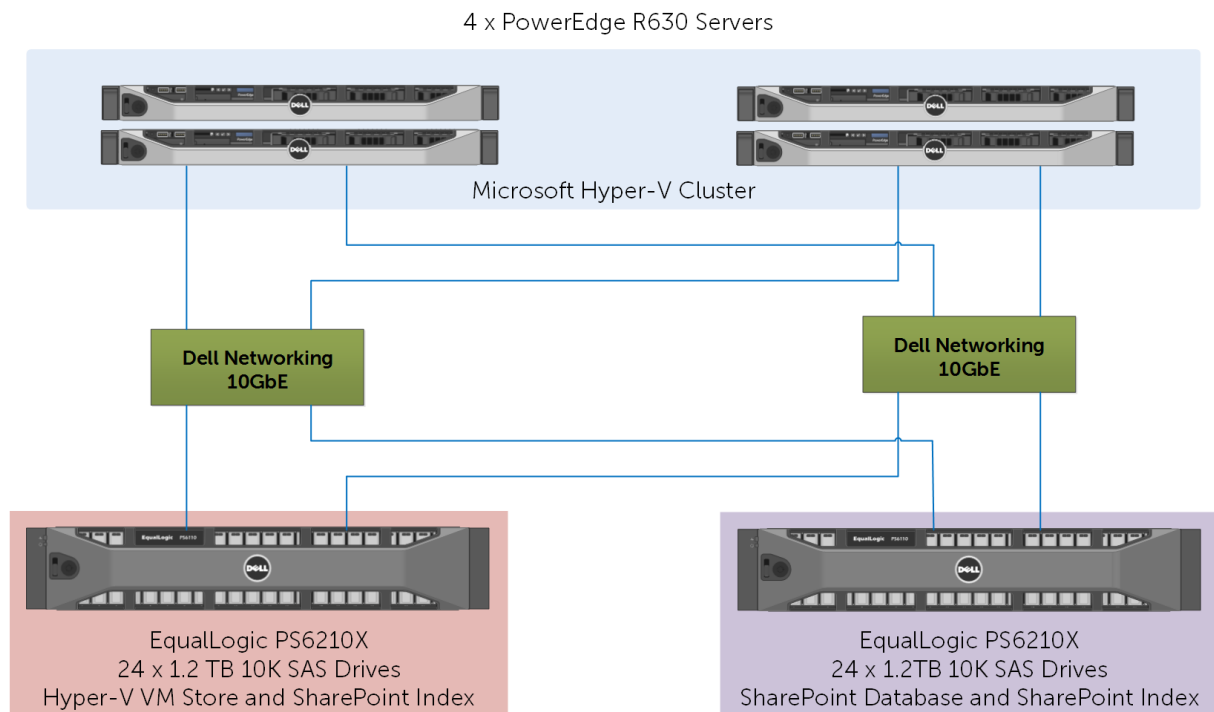


Figure 2 Storage architecture for Hyper-V cluster

As shown in [Figure 2](#), the two storage arrays are configured as two different storage pools to separate the disk spindles used for application storage from the VM store. A single large volume is created on the EqualLogic array used for the VM store and configured as a Clustered Shared Volume (CSV) for the Hyper-V cluster. A number of volumes – as described in [Section 2.2.1](#) – are created on the second EqualLogic array for SharePoint application data.

## 2.1.2 Network architecture and configuration

PowerEdge R630 has several network connectivity options. Apart from Network Daughter Card (NDC) connectivity, PowerEdge R630 also offers expansion through three PCIe slots. In this reference architecture, to build HA for the network connections, a QLogic 57800 dual port 10GbE SFP+ NDC and an additional QLogic 57810 dual port 10GbE PCIe add-on NIC are used. As shown in [Figure 3](#), the network connection redundancy is achieved by connecting Port 0 of NDC and add-on NICs to LAN fabric, Port 1 of NDC, and add-on NIC to SAN fabric.

The LAN and SAN fabrics are separated to provide isolation between the data and application traffic. The two TOR switches in each fabric are lagged together through an inter-switch-link (ISL), which provides 80-Gb bandwidth between the two switches. The solution provides four 10GbE uplinks from each switch to link into an existing core network infrastructure. As shown in [Figure 3](#), a native Windows network team is created with Port 0 of NDC and add-on NIC connections. A Hyper-V virtual switch is then deployed with the teamed NIC, and network Quality of Service (QoS) is enabled to make sure that the Live Migration, Cluster and VM management networks are allocated minimum bandwidth weight. This configuration enables traffic throttling when specific network traffic requires the bandwidth. For the application network connections in the VMs, virtual network adapters are configured with VLAN IDs to provide network traffic isolation at the LAN fabric.

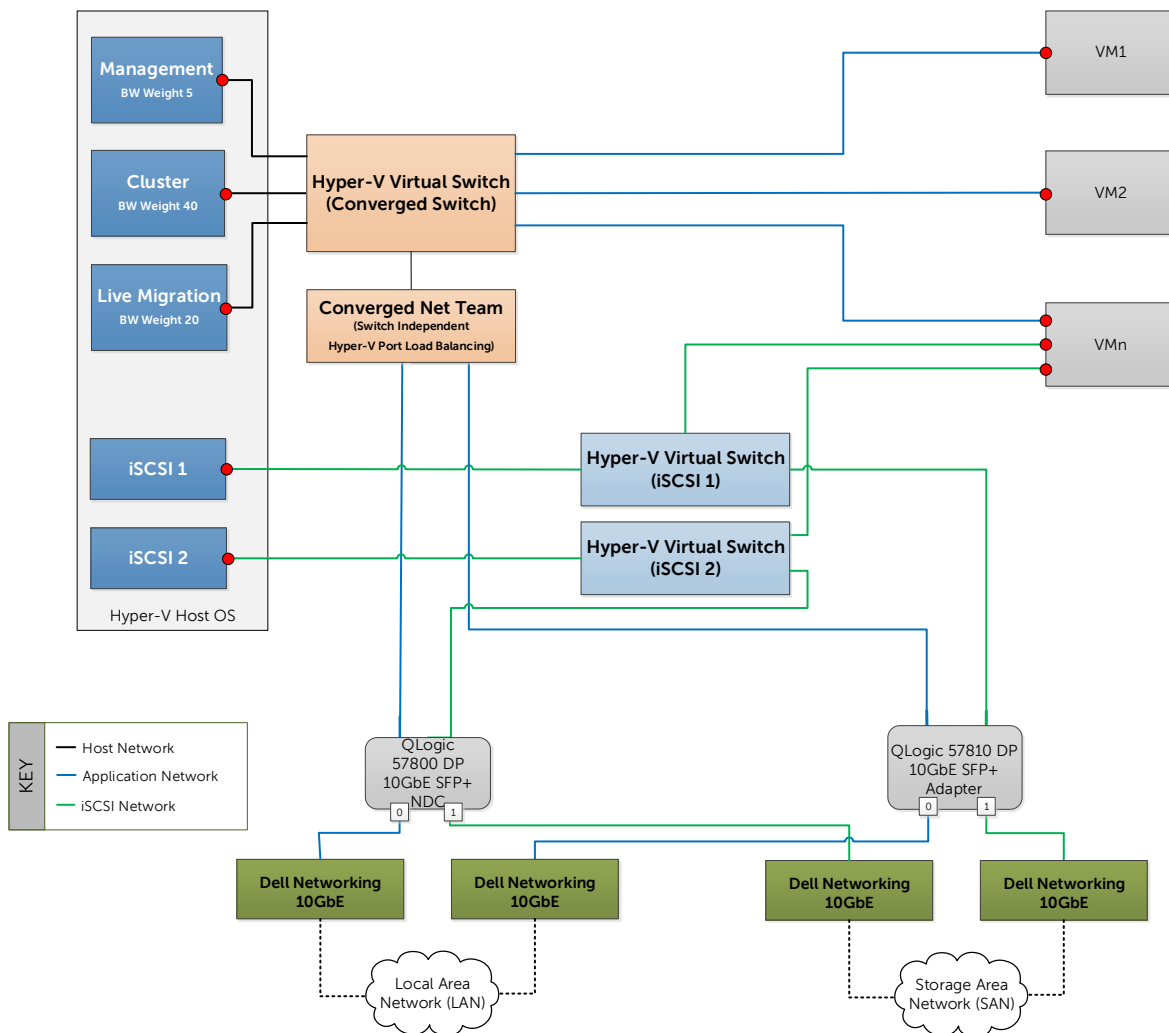


Figure 3 Network architecture for Hyper-V cluster

For the SAN connectivity in the host Operating System and the VMs that require in-guest iSCSI connections, two Hyper-V virtual switches are deployed, each with Port 1 of NDC and add-on NIC connections.

## 2.2 Application architecture

The application infrastructure implemented for SharePoint Server farm is virtualized and deployed on the Hyper-V infrastructure shown in [Figure 1](#).

Table 3 summarizes the VM resource requirements for deploying the SharePoint infrastructure.

Table 3 Farm virtual machine configuration

Application Role	Number of VMs	Total Virtual CPU	Total Memory (GB)	Networks
SharePoint Web Front End	2	8 (4 per VM)	32 (16GB per VM)	3 x Network Adapters per VM
SharePoint Application Server	2	8 (4 per VM)	16 (8GB per VM)	4 x Network Adapters per VM
SQL For SharePoint Databases	2	12 (6 per VM)	128 (64GB per VM)	6 x Network Adapters per VM
Virtual Load Balancer Appliances	2	4 (2 per VM)	8 (4GB per VM)	2 x Network Adapters per VM

[Figure 4](#) provides a high-level logical view of the application architecture for SharePoint deployment in this reference architecture. Each VM in this solution runs Windows Server 2012 R2 Data Center Edition and is configured with a fixed VHDX hard disk drive format as OS drive. The fixed VHDX format is recommended instead of pass-through disks for simplified management. Utilizing a virtual load balancer appliance reduces complexity of the network design.



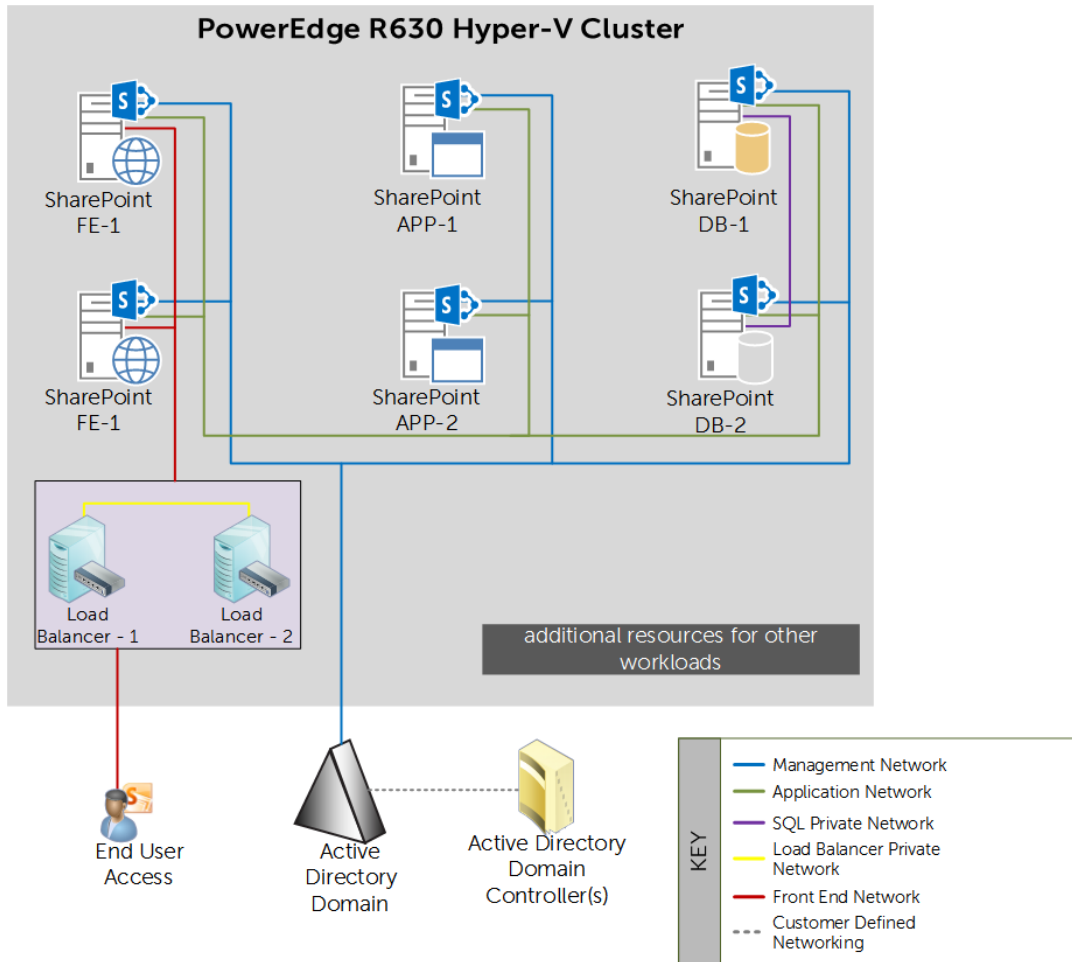


Figure 4 Application architecture for the SharePoint farm

This paper provides suggestions for the virtual load balancer configuration, but the reader must refer to the respective vendor's load balancer virtual appliance manual for appropriate configuration.

The SharePoint 2013 farm configuration, as shown in [Figure 4](#), includes two Web Front End (WFE), two Application Servers (APP) and two Database Servers in a failover cluster configuration. The Application Servers in the farm are 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 Software deployed in this farm configuration

	WFE and Application Servers	Database Servers
Operating System	Windows Server 2012 R2 Data Center Edition	
SharePoint Server	SharePoint Server 2013 Standard Edition SP1	NA
Database Server	NA	SQL Server 2012 SP1 Enterprise Edition with CU3

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 Additional Resources section of this paper.

A SQL failover cluster is used to provide SQL database services HA. Leveraging iSCSI connections in the SQL guests enables direct connectivity to the iSCSI shared storage. The SQL failover cluster ensures that a single SQL VM downtime does not impact overall SharePoint application services availability.

For SharePoint Front End load balancing and HA, two virtual load balancers are used. These virtual load balancers work as an HA pair. The HA configuration of the load balancers helps ensure that SharePoint services are highly available.

### 2.2.1 Storage architecture

Two storage components are employed in the application architecture described in this paper. The first component is the Cluster Shared Volume (CSV) for VM store, which is required to configure Hyper-V HA. With the VMs stored on shared iSCSI EqualLogic storage, each PowerEdge R630 host can run the SharePoint Application VMs. A single RAID 10 volume is created on an EqualLogic PS6210X array to host the CSV VM store.

The second component contains the application data storage. Similarly, a single RAID 10 volume is created on another EqualLogic PS6210X array for SharePoint content databases. For more details, see [Table 5](#).

Table 5 Overview of farm storage

Storage Array	Purpose
VM Storage Array	CSV Volume for VM Store and Search Index volumes 1 x Dell EqualLogic PS6210X; 24 1.2 TB 10K hard drives in RAID 1 configuration
SharePoint Storage Array	SharePoint Database storage and Search Index volumes 1 x Dell EqualLogic PS6210X; 24 1.2 TB 10K hard drives in RAID 1 configuration

The storage arrays used for VM and application storage are deployed as different storage pools to build storage isolation at the array level. As shown in [Figure 2](#), a single CSV volume is used for VM storage. For the application storage, as shown in Figure 5, multiple volumes are created on the EqualLogic array dedicated for application storage. Within each VM, with application storage requirements, software iSCSI initiator is used to connect to the iSCSI storage targets.

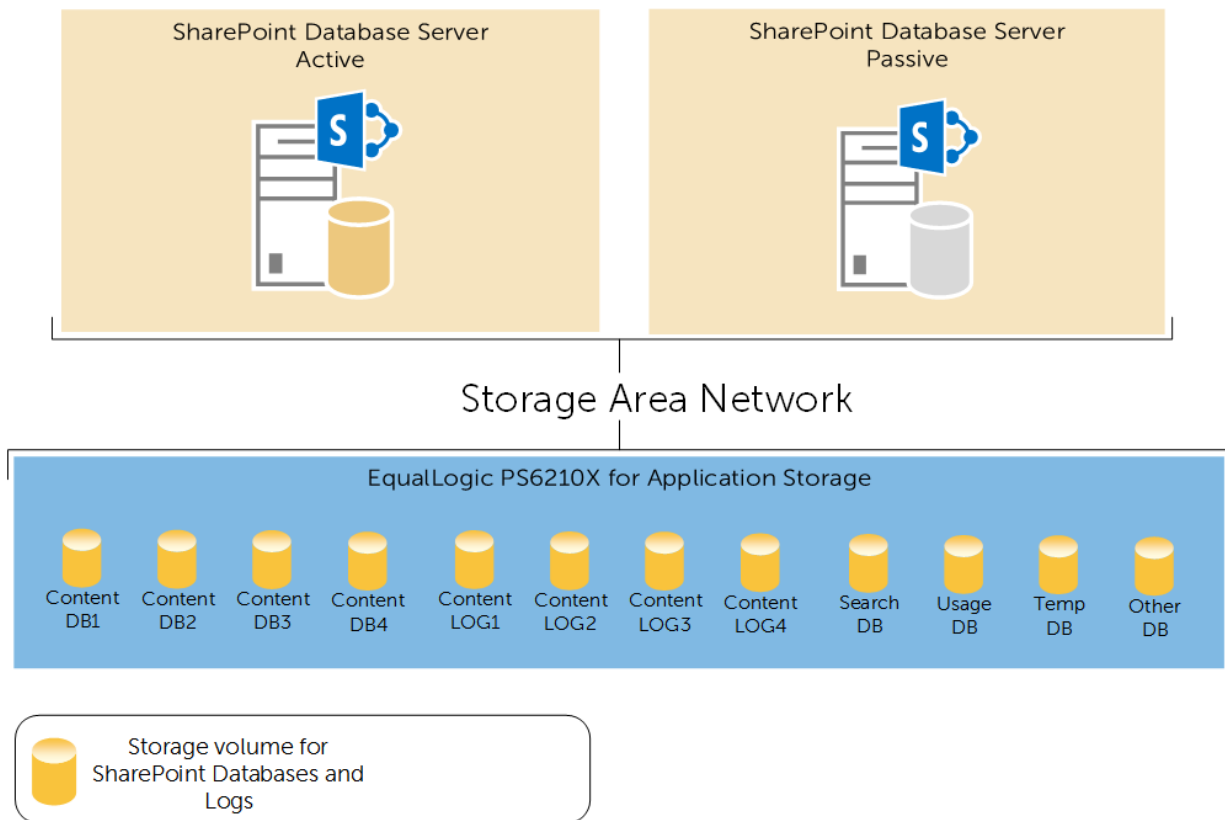


Figure 1 SharePoint database backend layout

Within this reference architecture, one SharePoint 2013 web application with four site collections is deployed. Each site collection is placed in a separate content database and therefore, four content databases are created.

Table 6 provides the details for each application volume created on the EqualLogic array.

Table 6 Overview of SharePoint databases

SharePoint Database	Number of volumes	Size of each volume
Content Database	4	650 GB
Content Database Logs	4	100 GB
Search Database	1	200 GB
Usage Database	1	200 GB
Temp DB	1	50 GB
Other SharePoint Databases (SharePoint Configuration and Admin Content)	1	100 GB

Similar to SharePoint 2010, the SharePoint 2013 search architecture includes HA at the application tier (for instance, the search crawler) and includes greater redundancy. This design provides flexibility and allows the query and crawler roles to scale-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. In this reference architecture, two virtual machines hosting both crawler and query roles are used in the SharePoint 2013 farm.



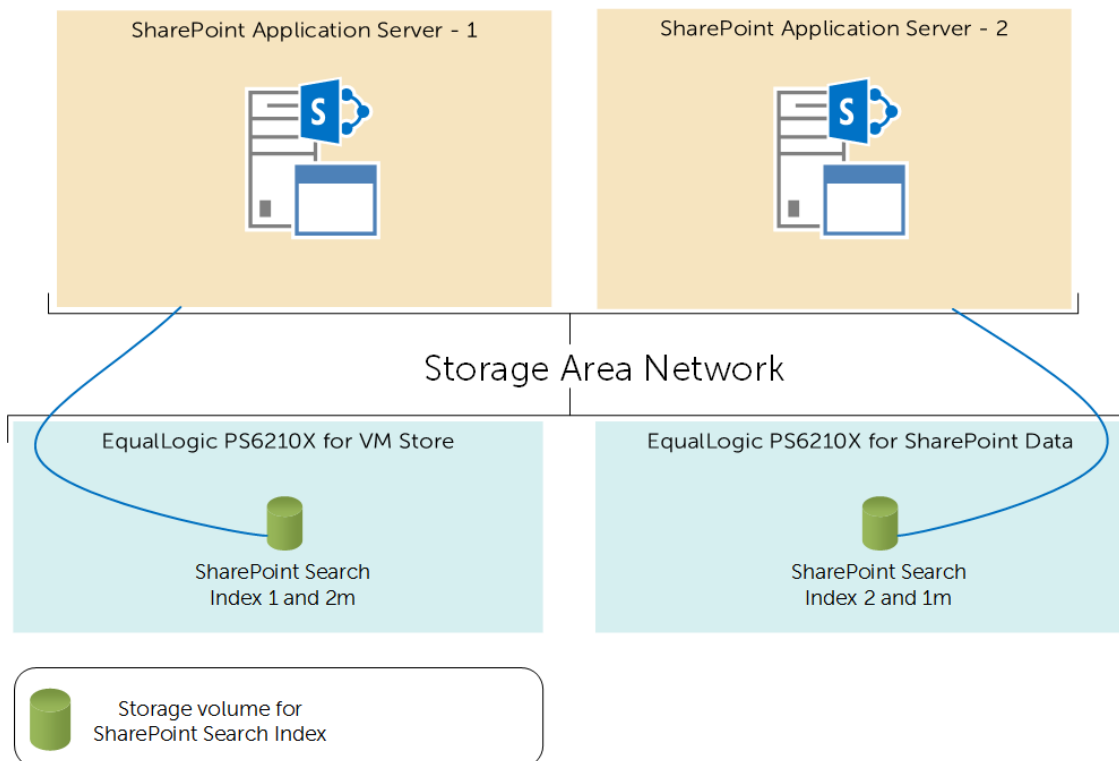


Figure 2 SharePoint search storage architecture

The farm configuration used for the search service application configuration is shown in [Figure 6](#). The search crawler and query roles are hosted only on the application servers for improved search performance. Both the crawlers are associated with the same crawl database, and index the same content source.

The search architecture shown in [Figure 6](#) 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 6](#), the primary index partitions are represented as index 1 and index 2. The two primary index partitions are replicated for redundancy; a replica, which is a copy of the index partition, is placed on each application server. On both of the application servers, software iSCSI initiators are used to connect to the iSCSI targets created for storing search index content. Two volumes—one per application server—are used to store each index partition. To provide highly available search index content, these volumes are created on separate iSCSI arrays—one on the VM store array and second on the SharePoint data store array.

## 2.2.2 Network architecture

According to application best practices and infrastructure design principles, each application network is deployed as a separate workload VLAN that is defined in the data center core network. The Dell PowerEdge R630 server provides two 10GbE ports labeled as QLogic 57800 NDC in [Figure 3](#), and two

network connections are provided by the add-on QLogic 57810 PCIe NIC. This configuration provides a total of four network connections within the host operating system for network connectivity. As mentioned in [Section 2.1.2](#), for the LAN connectivity of the host operating system, a native network team is created using the two NDC network connections, and a converged virtual network switch is implemented. All the workload VLANs are created as virtual network adapters on the converged virtual switch across these four network connections. The virtual switches created using the additional 10GbE network connections are used for iSCSI connectivity in the host operating system and the VMs that require in-guest iSCSI connectivity.

[Figure 7](#) illustrates this network architecture and the virtual switch implementation in each PowerEdge R630 server in this reference architecture.

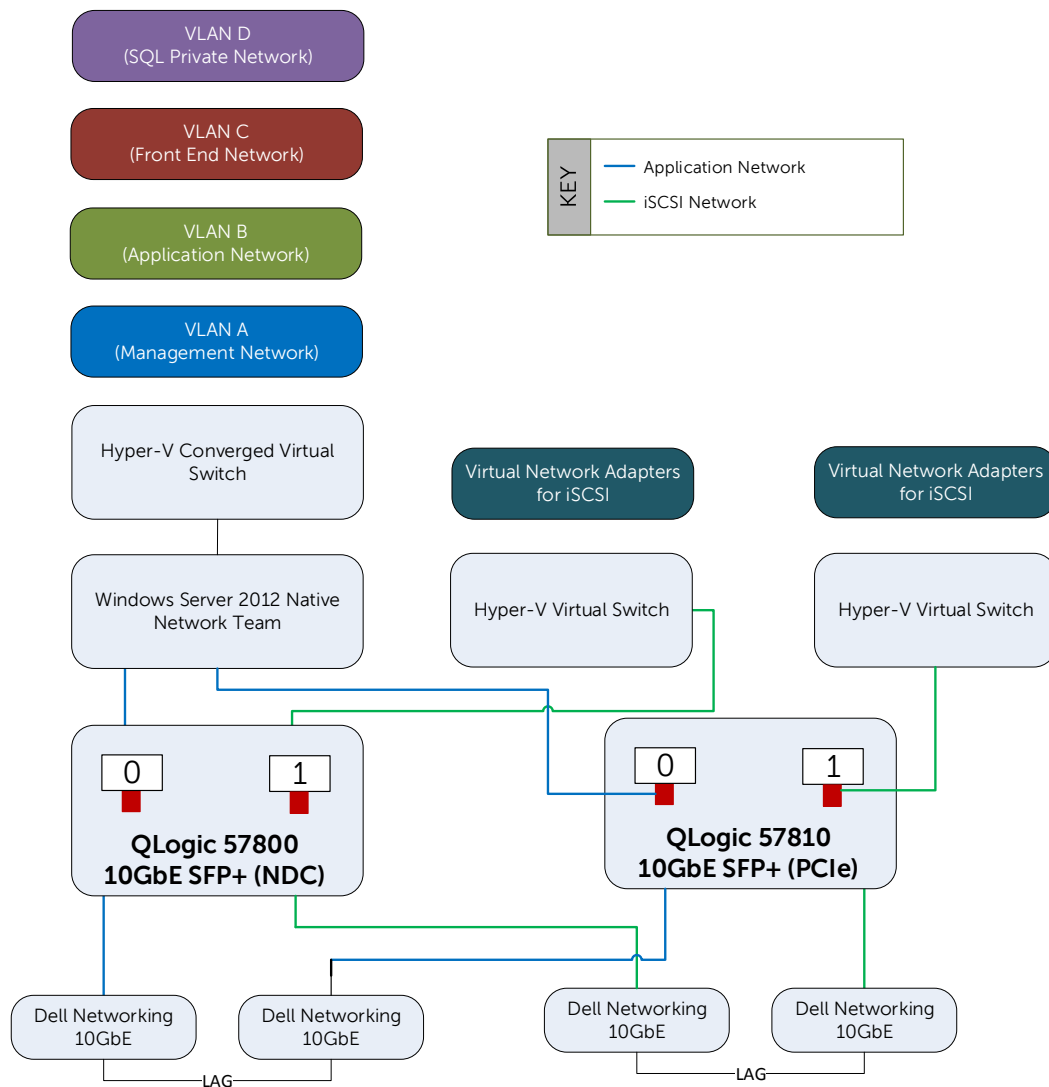


Figure 3 VM network architecture for farm servers

## 2.3 Virtualization rules

Hyper-V failover clustering enables application recovery in case of hardware failure. Using failover features built into the Hyper-V cluster implementation and System Center Virtual Machine Manager (SCVMM) Dynamic Optimization (DO), the applications can be restored to their healthy state almost instantaneously. These features can be configured independently of the guest operating system and the application running in the VM.

Whenever there is a change in the demand for resources, SCVMM DO enables automated resource load balancing, thus helping IT administrators offload the task of moving workloads within the cluster. One of the significant benefits of Live Migration is to enable administrators to proactively prepare the environment for maintenance tasks by live-migrating virtual machines to other available hosts. In the absence of DO, an IT administrator should know where to move specific VMs to maintain application best practices. By combining SCVMM DO with an availability set configuration, an administrator need not worry about specific VM placements in the cluster. Instead, the administrator can migrate the VMs and allow DO to ensure that application best practices are not compromised. As a part of this solution, four availability sets are defined.

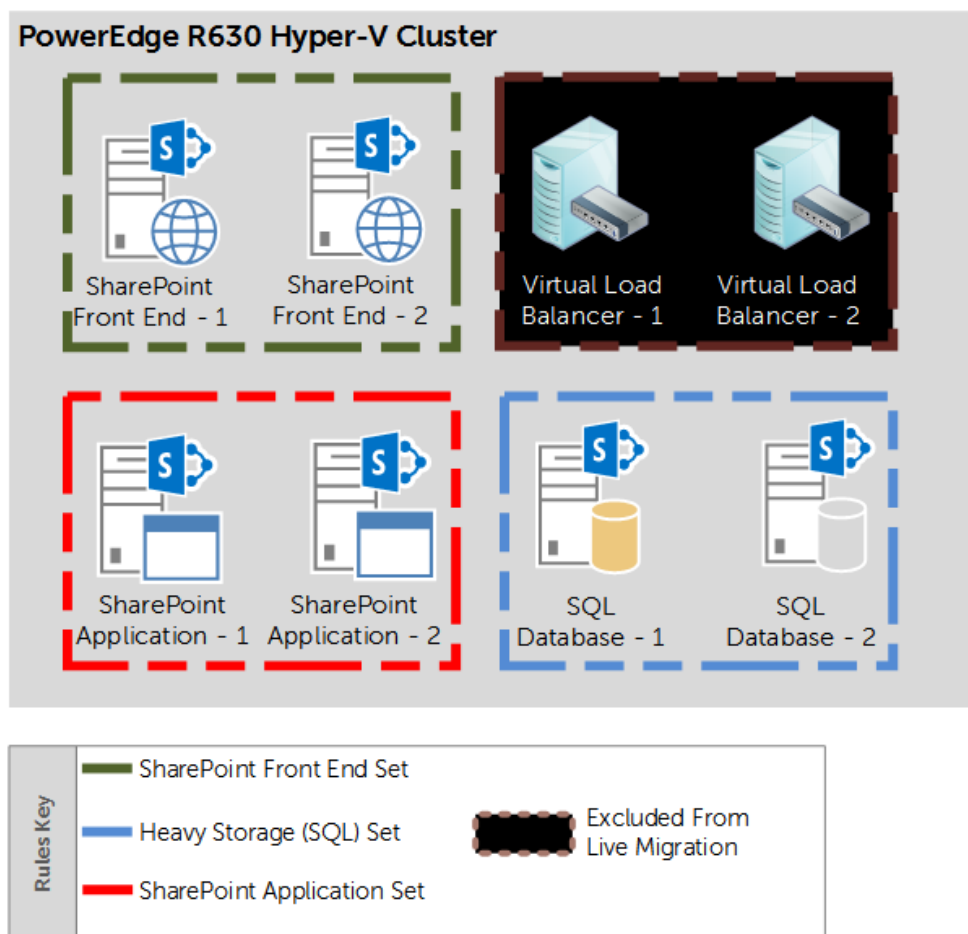


Figure 4 Virtualization rules for separating VM roles

As shown in [Figure 8](#), four availability sets are created within SCVMM. The "Heavy Storage (SQL) Set" ensures that the VMs running SQL applications are never placed on the same physical host. Based on this configuration, a minimum of two hosts are needed to run all SharePoint application VMs. Also, as shown in [Figure 8](#) and based on virtual load balancer appliances are excluded from Live Migration.

If SCVMM is not available to manage the hypervisor cluster, the availability sets shown in Figure 8 can be implemented as [Anti-Affinity Class names](#) within the cluster and can be configured using Windows PowerShell.



## 3 Verification

The reference architecture presented in this paper is built to implement the design considerations described in [Section 1.4](#). It is important to verify the solution architecture implementation to ensure that the underlying virtual infrastructure delivers optimal performance and is highly available. To this end, several tests are performed to verify that the application services are not impacted due to infrastructure component failure or planned failover activities of application VMs. It is observed that the application services do not get impacted when a component such as a top-of-rack network switch fails or during the planned Live Migration of application VMs from one host to another. A physical host failure triggers movement of VMs from the failed host to another available host. This causes interruption in application services for less than five minutes during which the SharePoint application farm works in a degraded capacity. As soon as the application VM restarts on another available host, the application farm returns to normal state and operates at full capacity. It is also observed that the SCVMM DO rules ensure that no two VMs running similar application services run on the same physical host.

The performance of the SharePoint application architecture in this solution is verified using an in-house SharePoint Load Generation tool. The average farm response time is observed to be less than a second and in line with the design goal for this reference architecture. The physical servers in this reference architecture are sized appropriately to support full application capacity—5,000 concurrent users—even in the failure of two PowerEdge R630 servers, without compromising the infrastructure design considerations described in [Section 1.4](#).

### 3.1 Methodology

As mentioned earlier, within the scope of this paper, a collaboration usage profile was used to perform load testing through the internally developed SharePoint load generation tool. 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](#).

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 Dataset used in the performance verification

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

Table 8 provides an overview of the test profiles used in this load test.

Table 8 Overview tests performed during the load test

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 aggregated SharePoint content database size was around 1 TB. During load test, 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.



## 3.2 Results

As mentioned earlier, the performance verification of the SharePoint farm was conducted considering the availability of only two hypervisor hosts in the infrastructure. This scenario represents a degraded state for the Hyper-V infrastructure. This section describes some of the performance metrics gathered during the performance verification and discusses the significance of the same.

Table 9 Overview of farm performance at peak user load

	Observations
<b>Maximum concurrent<sup>1</sup> user load validated</b>	5,000
<b>Requests per second<sup>2</sup> at Max concurrent user load</b>	192/sec
<b>Avg. Response time at the maximum validated user load</b>	0.098 seconds

Table 9 provides a high-level overview of the farm performance at the validated peak concurrent user load. The above metrics indicate that the farm configuration used for this performance study can support faster, below one-second farm response times even at the maximum validated concurrent user load.

Table 10 Overview of farm server performance at peak load

	SharePoint Large Farm in degraded state		
	Front End Servers	Application Servers	Database Servers
<b>Average Processor Utilization (%)</b>	61%	19%	48%
<b>Avg. Disk Transfers / Sec (_Total)</b>	1,089		
<b>Avg. Disk Seconds / Transfer (_Total)</b>	0.013		
<b>Avg. Disk Bytes / Transfer (_Total)</b>	55,659,870 (53KB)		

---

<sup>1</sup> Concurrency refers to number of simultaneous requests to the farm servers

<sup>2</sup> This number indicates the avg. requests per second generate during the load test duration and this is a Visual Studio reported metric

As shown in [Table 10](#), the average processor utilization on the farm servers is around the threshold limits (60%) and indicates that there is still room for future growth in number of concurrent users or to support the application needs. The average disk transfer per second indicates that the EqualLogic PS6210X is capable of delivering the needs of a SharePoint farm that has 5,000 concurrent users with a content database of 1 TB.



## 4 Solution specifications

Table 11 Overview of virtualization host configuration

Component	Details	
Virtualization Hosts	4 x Dell PowerEdge R630 Servers	
	Processor	2 x Intel Xeon E5-2660v2 Family
	Memory	128 GB; 8 x 16GB DDR3 DIMMs
	HDD	2 x 600GB 15K SAS in RAID 1 for OS volume
	Network	QLogic 57800 dual port SFP+ NDC QLogic 57810 dual port SFP+ Add-on NIC
	OS	Windows Server 2012 R2 Data Center Edition
Storage Arrays	2 x EqualLogic PS6210X for VM Store and SharePoint Data as separate Storage pools in a single EqualLogic Group 24 x 1.2 TB 10K SAS Drives per array in RAID 10 configuration	
Networking	2 x Dell Networking S4810-ON for Local Area Networking (LAN) 2 x Dell Networking S4810-ON for iSCSI Storage Area Networking (SAN)	

Table 12 Overview of virtual machine configuration

Component	Details	
SharePoint Front End Servers	2 x Hyper-V Virtual Machines	
	Processor	4 x Virtual Processors
	Memory	16 GB
	OS VHDX	1 x 160 GB stored on the VM Store array
	Network	3 x Network Adapters for Management, SQL Application and Public connectivity
	OS	Windows Server 2012 R2 Data Center Edition
	SharePoint	SharePoint Server 2013 Standard Edition
SharePoint Application Servers	2 x Hyper-V Virtual Machines	
	Processor	4 x Virtual Processors
	Memory	8 GB

	OS VHDX	1 x 160 GB stored on the VM Store array
	Data Volumes	1 x 200 GB iSCSI volume per VM
	Network	2 x Network Adapters for Management and SQL Application connectivity 2 x Network Adapters for iSCSI connectivity
	OS	Windows Server 2012 R2 Data Center Edition
	SharePoint	SharePoint Server 2013 SP1 Standard Edition
SharePoint Database Servers	2 x Hyper-V Virtual Machines	
	Processor	6 x Virtual Processors
	Memory	64 GB
	OS VHDX	1 x 160 GB stored on the VM Store array
	Data Volumes	4 x 650 GB iSCSI volume for SharePoint Content DB 4 x 100 GB iSCSI volume for SharePoint Content Log 1 x 50 GB iSCSI volume for Temp DB 1 x 200 GB iSCSI volume for Search DB 1 x 200 GB iSCSI volume for Usage DB 1 x 100 GB iSCSI volume for Other SharePoint DB
	Network	3 x Network Adapters for Management, SQL Application and SQL Cluster Private connectivity 2 x Network Adapters for iSCSI connectivity
	OS	Windows Server 2012 R2 Data Center Edition
	SQL	SQL Server 2012 SP1 Enterprise Edition CU3



## 5 Solution summary

The solution architecture described in this document addresses common problems associated with hardware sprawl which can be alleviated through virtualized SharePoint deployments that follow best practices for hardware abstraction and HA.

The PowerEdge R630 server is an efficient and powerful rack-based building block for deploying and virtualizing enterprise applications, such as Microsoft SharePoint Server 2013. This reference architecture includes four PowerEdge R630 servers configured and deployed to deliver HA and meet the performance objectives stated in [Section 2.2](#).

The recommended solution architecture described in [Section 3](#) takes into account the best practices laid out by Dell and Microsoft. The design points emphasize the use of the latest hardware features and capacity of PowerEdge R630 server to enable application consolidation, hardware abstraction and infrastructure HA. The infrastructure HA features ensure that the application services are continuously available and deliver optimal performance.

The storage architecture described in this reference architecture uses EqualLogic PS6210X arrays and provides both performance and capacity required by the application for optimal response times as described in the design goals.

The converged virtual switch used for VM network access enables high quality-of-service (QoS) for different classes of traffic and ensures that the demands for network bandwidth are met. Therefore, the bandwidth requirements for VM Live Migration and cluster traffic are given priority to enable seamless movement of VMs during optimization and/or planned downtime of the hypervisor hosts. Leveraging iSCSI connectivity in the SQL guests enables the implementation of SQL failover clustering to build HA for the SQL database and SharePoint application services.

Abstracting the application VMs completely from their physical hosts requires dynamic VM placement. VMs must be load-balanced across servers, depending on their workload type and their relationship to other VMs. To achieve server load balancing, the cluster is configured with Microsoft System Center Virtual Machine Manager (SCVMM) Dynamic Optimization (DO) to automatically live-migrate VMs to less-utilized hosts. SCVMM availability sets are used to ensure that the VMs with similar application roles or high resource requirements do not coexist on the same hypervisor host. If SCVMM is not available for management, the failover cluster anti-affinity class names can be used to achieve VM isolation.



## A Additional resources

[Dell Consulting Services for Microsoft SharePoint](#) is focused on meeting your needs with proven services and support.

**DellTechCenter.com** is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

- Dell PowerEdge R630

Referenced or recommended Microsoft publications:

- [Overview of farm virtualization and architectures for SharePoint 2013](#)
- [Best practice configurations for the SharePoint 2013 virtual machines and Hyper-V environment](#)
- [Configuring Dynamic Optimization in VMM](#)
- [Configure Availability Sets in VMM](#)
- [What's new in Windows Server 2012 Virtualization](#)

