

Run Multiple Microsoft® Applications Simultaneously With Hitachi Unified Compute Platform Director and Microsoft Hyper-V®

Reference Architecture Guide

By Rick Andersen, Steven Burns, Leo Nguyen, and Nathan Tran

April 28, 2015



Microsoft Partner
Gold Server Platform

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Table of Contents

Solution Overview.....	2
Key Solution Components.....	5
Hardware Components.....	5
Software Components.....	7
Solution Design.....	10
Microsoft Exchange Server 2013 Configuration.....	10
Microsoft SharePoint Server 2013 Configuration.....	10
Microsoft SQL Server 2014 Configuration.....	10
Microsoft Lync Server 2013 Configuration.....	11
Infrastructure Resources for Management Cluster.....	11
Infrastructure for Compute Resources.....	12
Storage Resources.....	19
Application Infrastructure for Microsoft Exchange Server 2013.....	30
Application Infrastructure for Microsoft SharePoint Server 2013	34
Application Infrastructure for Microsoft Lync Server 2013	41
Application Infrastructure for Microsoft SQL Server 2014	42
Engineering Validation.....	47
Microsoft Exchange 2013	47
Microsoft SharePoint Server 2013	50
Microsoft Lync Server 2013	52
Microsoft SQL Server 2014	53

Run Multiple Microsoft® Applications Simultaneously With Hitachi Unified Compute Platform Director and Microsoft Hyper-V®

Reference Architecture Guide

This reference architecture using hardware from Hitachi Data Systems describes how to run Microsoft® Exchange Server 2013, Microsoft SQL Server® 2014, Microsoft SharePoint® Server 2013, and Microsoft Lync® Server 2013 on a shared infrastructure. This solution uses Microsoft Windows Server® virtual machines running on Hitachi Unified Compute Platform with Microsoft Hyper-V® to host these applications. The benefits of this solution with Hitachi Unified Compute Platform Director and Hyper-V are the following:

- Faster deployment
- Increased scalability
- Increased reliability
- Reduced risk

The key components of this solution are the following:

- **Storage** — Hitachi Virtual Storage Platform G1000
- **Compute** — Hitachi Compute Blade 500
- **Network** — Brocade Network Switches
- **Virtualization** — Hitachi Unified Compute Platform and Hyper-V

This solution integrates Hitachi storage and servers with Microsoft Hyper-V to support multiple Microsoft server applications. These applications run on virtual machines in an environment that shares compute and storage resources. Hitachi Unified Compute Platform Director integrated with Microsoft System Center 2012 R2 manages the environment.

This tightly integrated server, storage, network, and management infrastructure is built using Microsoft validated reference architectures with preconfigured VLANs, server clusters, and storage pools.

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution Overview

This reference architecture is built on hardware and software virtualization platforms from Hitachi and Microsoft.

- Hitachi Unified Compute Platform Director
- Hitachi Compute Blade 500 supplies the compute resources
- Hitachi Virtual Storage Platform G1000 supplies the storage resources
- Brocade supplies Fibre Channel and converged networking resources

Hitachi Unified Compute Director takes the basic infrastructure components of compute, network and storage and unifies them into a single platform solution for application deployment in a private cloud.

Hitachi Unified Compute Platform (UCP) provides the following capabilities for this reference architecture:

- Two distinct Hyper-V failover clusters managed by Unified Compute Platform Director.
 - A multiple node tenant cluster that hosts the production application virtual machines.
 - A two node management cluster that contains the following:
 - Software to deploy virtual machines to the tenant cluster.
 - Unified Compute Platform Director software to manage the Hyper-V private cloud infrastructure.
-

This reference architecture provides the following capabilities:

- Ability to manage the physical and virtual infrastructure with a single pane of glass using Hitachi Unified Compute Platform Director. Accelerates the deployment of servers and storage into the infrastructure and achieves higher utilization of clusters and storage.
 - Rapid deployment at a customer site as Hitachi Unified Compute platform is a converged solution built and configured at Hitachi Data Systems distribution centers prior to shipment. Fast time to production implementation.
 - Gain flexibility by native integration of Unified Compute Platform Director with System Center Virtual Machine Manager 2012 R2. This provides a unified view of the entire converged platform, storage, compute and networking.
 - Delivers predictable performance.
 - Automation and inspection of the environment with Unified Compute Platform Director providing a consistent interface. This allows for easy problem determination and remediation.
 - Highly automated and rapid provisioning of Hyper-V servers, host clusters and storage. Service templates automate and speed up deployment removing the risk of human errors.
-

Figure 1 provides a high level design of the reference architecture in this technical paper. The two UCP Director management nodes are Hitachi Compute Rack 210H servers.

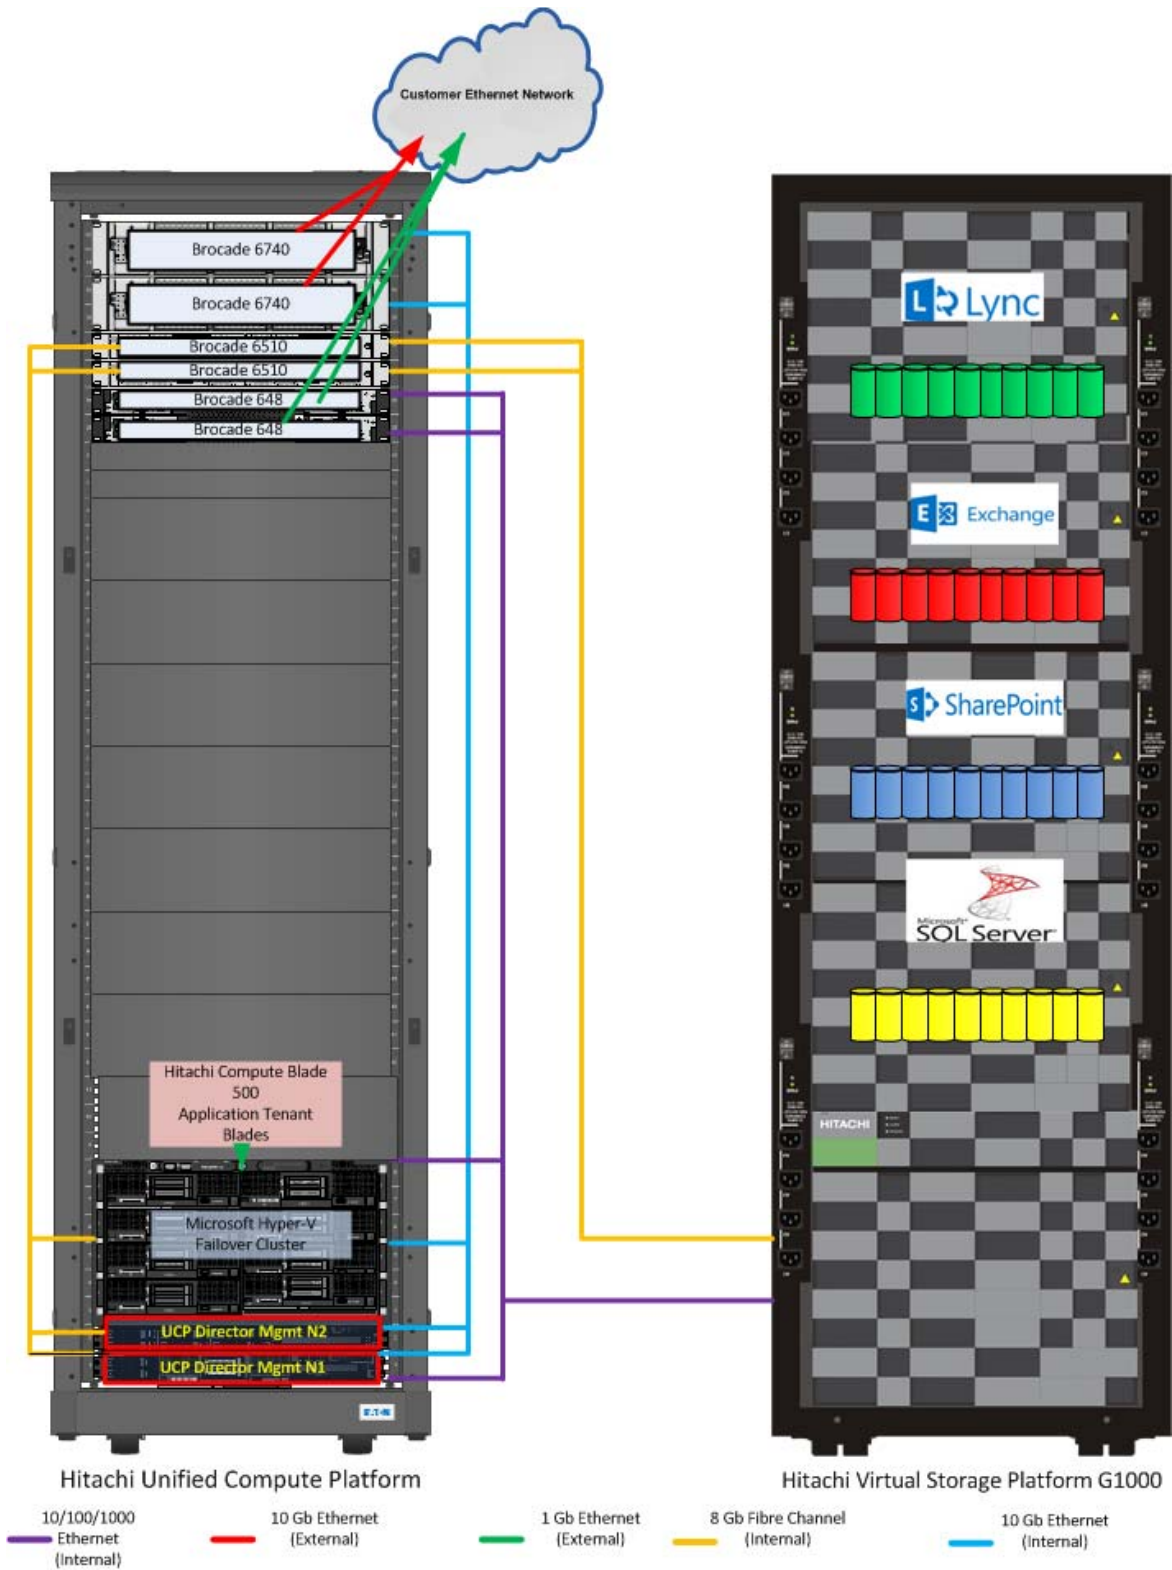


Figure 1

Key Solution Components

These are the key hardware and software components used to deploy this solution.

Hardware Components

Table 1 lists the hardware components used to deploy this solution.

Table 1. Hardware Components

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Virtual Storage Platform G1000	<ul style="list-style-type: none"> ■ Dual controller ■ 16 × 8 Gb/sec Fibre Channel ports ■ 512 GB cache memory ■ 392 × 600 GB 10k RPM SAS disks, SFF ■ 200 × 4 TB 7.2k RPM SAS disks LFF 	80-01-25-00	1
Hitachi Compute Blade 500 chassis	<ul style="list-style-type: none"> ■ 8-blade chassis ■ 2 Brocade 5460 Fibre Channel switch modules ■ 2 Brocade VDX6746 10 GbE DCB switch modules ■ 2 management modules ■ 6 cooling fan modules ■ 4 power supply modules 	A0206-A-9296 7.2.1 4.1.2a	1
CB520H B2 server blade (with CNA)	<ul style="list-style-type: none"> ■ Half-size blade ■ 2 × 12-core Intel Xeon E5-2697 processors, 2.70 GHz ■ 384 GB memory per blade ■ 1 Hitachi 8 Gb/sec 2-port Fibre Channel mezzanine card ■ 1 Emulex 10 GbE Onboard CNA 	04-30	8

Table 1. Hardware Components (Continued)

<i>Hardware</i>	<i>Description</i>	<i>Version</i>	<i>Quantity</i>
Hitachi Compute Rack 210H	<ul style="list-style-type: none"> ■ 1U rack server ■ 2 × 8-core Intel Xeon E5-2670 processors, 2.60 GHz ■ 192 GB memory per blade ■ 8 GB Dual Port HBA PCI 2.0 ■ 2 × 300 GB 10k RPM SAS disks ■ 10 Gb dual channel NIC 	01-09	2
Brocade 6740 Ethernet switch	<ul style="list-style-type: none"> ■ 40-port with 10 GbE speed 	4.1.2a	2
Brocade 648 Ethernet switch	<ul style="list-style-type: none"> ■ 48-Port with 10/100/100Mbps 		2
Brocade 6510 Fibre Channel switch	<ul style="list-style-type: none"> ■ 24-port with 8-16 Gb/sec speed 	7.2.1a	2

Hitachi Virtual Storage Platform G1000

[Hitachi Virtual Storage Platform G1000](#) provides an always-available, agile, and automated foundation that you need for a continuous infrastructure cloud. This delivers enterprise-ready software-defined storage, advanced global storage virtualization, and powerful storage.

Supporting always-on operations, Virtual Storage Platform G1000 includes self-service, non-disruptive migration and active-active storage clustering for zero recovery time objectives. Automate your operations with self-optimizing, policy-driven management.

The Hitachi Virtual Storage Platform G1000 has full support for Microsoft Offloaded Data Transfer, which provides rapid virtual machine provisioning by reducing the load on network and server resources.

Hitachi Compute Blade 500

[Hitachi Compute Blade 500](#) combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform.

Software Components

Table 2 lists information about the software components needed to implement this solution.

Table 2. Software Components

<i>Software</i>	<i>Version</i>
Hitachi Unified Compute Platform Director	3.5 Patch 2
Hitachi Command Suite	8.0
Hitachi Tuning Manager	8.0
Hitachi Compute Systems Manager	8.8
Hitachi Infrastructure Adapter for Microsoft® System Center Operations Manager	1.3.0
Hitachi Storage Adapter for Microsoft Windows® PowerShell®	5.0
Hitachi Storage Adapter for Microsoft® System Center Orchestrator	1.1.0
Hitachi Dynamic Provisioning	Microcode Dependent
Microsoft Systems Center Virtual Machine Manager	2012 R2
Microsoft Windows Server®	2012 R2 Datacenter 2012 R2 Standard for Applications
Microsoft Exchange Server	2013 SP1
Microsoft Lync® Server	2013
Microsoft SharePoint® Server	2013 SP1
Microsoft SQL Server®*	2014

Note: Multiple instances of Microsoft SQL Server were implemented in support of Microsoft Lync 2013 and SharePoint 2013.

Hitachi Unified Compute Platform Director Integration With Microsoft

[Hitachi Unified Compute Platform Director](#) offers the following for Microsoft cloud environments:

- Full parity across the RESTful API
 - Feature parity across the graphical user interface and command line interface
 - Provisions storage, server, and networking - all from single pane of glass
 - Bare-metal server provisioning
 - Management of virtual and physical servers
-

Unified Compute Platform Director software on the Unified Compute Platform integrates directly into Microsoft System Center Virtual Machine Manager (SCVMM) 2012 R2. It provides unified end-to-end infrastructure orchestration within a single interface.

Unified Compute Platform Director leverages your existing storage in one of two ways:

- Connect to your existing [Hitachi Virtual Storage Platform](#) or [Hitachi Unified Storage VM](#)
- Virtualize other storage arrays that you have from other vendors using Virtual Storage Platform or Unified Storage VM

Unified Compute Platform Director provides the following benefits:

- Centralization and automation of compute, storage, and networking components
- Significant reduction of time to value and operational costs across data centers
- Faster deployment of converged infrastructure with more efficient resource allocation
- Provides a foundation for the journey to the software defined datacenter using full support of the RESTful API

Hitachi Command Suite

[Hitachi Command Suite](#) manages virtualized storage and server infrastructures. With usability, workflow, performance, scalability, and private cloud enablement, Hitachi Command Suite lets you build sustainable infrastructures with leading storage technologies. It helps you flexibly align with changing business requirements and maximize return on IT investments.

Hitachi Compute Systems Manager

Hitachi Compute Systems Manager is the management software for Hitachi servers. Compute Systems Manager can be purchased with an optional Server Management Module, Network Management Module, or Server Deployment Module. Use Compute System Manager to introduce new servers into your data center environment.

Hitachi Storage Adapter for Microsoft Windows PowerShell

Hitachi Storage Adapter for Microsoft Windows PowerShell allows you to use PowerShell to manage Hitachi storage. This software provides a set of Hitachi storage cmdlets for discovering and managing Hitachi storage system information.

With the Hitachi storage cmdlets, you can create scripts with logic and other PowerShell cmdlets to automate complex storage tasks in the Hitachi Unified Compute Platform environment. Filter, sort, and group storage information by piping the output of Hitachi storage cmdlets to PowerShell cmdlets.

Hitachi Storage Adapter for Microsoft System Center Orchestrator

The Hitachi Storage Adapter for Microsoft System Center Orchestrator plugs into System Center Orchestrator and brings with it pre-defined tasks related to Hitachi storage systems (based on PowerShell cmdlets). Our provider also includes pre-defined activities that leverage a combination of the tasks. In addition to these pre-defined tasks and activities, the system administrator is able to create customized workflows for the Hitachi Unified Compute Platform environment using Hitachi-specific tasks.

Hitachi Infrastructure Adapter for Microsoft System Center Operations Manager

Hitachi Infrastructure Adapter for Microsoft System Center Operations Manager integrates with the SCOM server to report and monitor on Hitachi storage systems and Hitachi Compute systems and switches. The Hitachi Infrastructure Adapter for Microsoft SCOM software allows Hitachi storage and server administrators to monitor the UCP hardware components using a single, familiar, and easy to use interface. The System Center Operations Manager console provides numerous views that show the state, health, and other information of the various storage subsystem and server components in the Hitachi Unified Compute Platform environment. Administrators can also view alert information that is generated whenever any of the monitored subsystem components changes state. These management packs provide object views for the following:

- Storage subsystem
 - Storage controller
 - Port views and alerts
 - Array drive views and status
 - Dynamic provisioning pools view and status
 - RAID and LUN views and status
-

Solution Design

The combined architecture consists of multiple application infrastructures. Each application infrastructure is designed to serve a specific purpose. For the compute resource and storage resource application infrastructure, each has shared resources and dedicated resources.

Microsoft Exchange Server 2013 Configuration

The Microsoft Exchange Server 2013 configuration architecture consists of four mailbox servers designed to support 6,000 2 GB mailboxes using two database copies for mailbox resiliency.

Microsoft SharePoint Server 2013 Configuration

The Microsoft SharePoint 2013 architecture configuration consists of one application server, one Microsoft SQL Server 2014 database server, and eight web front end servers. The environment was designed to support 100,000 users, with a 1 percent concurrency rate, and a Microsoft SharePoint 2013 farm with 20 site collections using 20 different 200 GB content databases.

Microsoft SQL Server 2014 Configuration

The Microsoft SQL Server 2014 architecture configuration consisting of SQL server hosting and OLTP workload of varying sizes. This configuration models the activity and usage of an online stock brokerage firm with users and client access infrastructures that are distributed globally while accessing a back-end centralized database for customer and transaction records. Using this model, three separate configuration profiles were used to test and validate this solution.

The configuration profiles listed in Table 3 were used to evaluate the performance characteristic of this solution.

Table 3. Configuration Profiles

<i>Configuration</i>	<i>Database Size</i>	<i>Customer Record Count</i>	<i>vCPU</i>	<i>vRAM</i>
Small	1 TB	80,000	4 Cores	32 GB
Medium	2 TB	160,000	8 Cores	64 GB
Large	4 TB	320,000	16 Cores	128 GB

Microsoft Lync Server 2013 Configuration

The Microsoft Lync Server 2013 environment consists of three front end Microsoft Lync 2013 servers and one backend Microsoft SQL Server 2014 database server. The environment used to validate this reference architecture also included a dedicated mediation server. A certificate authority server was deployed in the environment to handle the certificate requirements of Lync 2013. The Lync 2013 environment was tested with 13,500 users. The number of users was based on Microsoft sizing recommendations for a Lync environment with three front end servers. This environment also allows for high availability because one front end server can be offline without exceeding the Microsoft recommended maximum number of 6,600 users per front end server.

Infrastructure Resources for Management Cluster

The management cluster for infrastructure resources is shared by all of the application environments. This infrastructure consists of a two node management cluster hosted on a pair of Hitachi Compute Rack 210H (CR 210) rack servers. The management for the infrastructure is provided by Hitachi Unified Compute Platform Director. The following software is used to manage the compute resources, storage resources, and Microsoft Hyper-V infrastructure:

- **Hitachi Unified Compute Platform Director** — Hitachi Unified Compute Platform Director integrates directly into Microsoft System Center Virtual Machine Manager 2012 R2. It provides unified end-to-end infrastructure orchestration within a single interface.
 - **Hitachi Command Suite** — This is the advanced data management tool that improves storage operations, provisioning, optimization, and resilience for Hitachi storage environments.
 - **Hitachi Compute Systems Manager** — This integrates into Hitachi Command Suite to provide general server management, power management, and system monitoring.
 - **Microsoft System Center Virtual Machine Manager 2012 R2** — This provides centralized management of virtualized hosts and virtual machines from a single console.
 - **Hitachi Infrastructure Adapter for Microsoft System Center Operations Manager** — This infrastructure adapter provides a management pack for object views of the storage subsystem, storage controller port views and alerts, array drive views and status, dynamic provisioning pools view, and RAID and LU views. This infrastructure adapter also provides a management pack for the discovery and monitoring of Hitachi Compute systems and switch families.
-

Infrastructure for Compute Resources

The infrastructure for compute resources consists of one Hitachi Compute Blade 500. This provides the basic foundation needed to start building out a Microsoft Hyper-V virtualized environment for Microsoft Exchange 2013, Microsoft SQL Server 2014, Microsoft SharePoint Server 2013, and Microsoft Lync Server 2013.

The infrastructure includes the following hardware:

- 1 Hitachi Compute Blade 500 chassis
- 8 server blades
 - 2 for hosting Exchange
 - 2 for hosting SharePoint
 - 1 for hosting SQL Server
 - 3 for hosting Lync
- 4 switch modules

Hitachi Compute Blade 500 Chassis Configuration

The chassis configuration uses the following component modules to power, manage, and monitor the chassis, server blades, and switches. These not only provide redundancy, but also provide ease of administration and maintenance.

The chassis includes the following hardware:

- 4 power supply modules
- 6 cooling fan modules
- 2 management modules

The following switch modules in the chassis support Ethernet and Fibre Channel connectivity:

- 2 Brocade 5460 Fibre Channel switch modules
- 2 Brocade DCB 10 GbE switch modules

Hitachi Compute Blade 500 Server Blade Configuration

The CB520H B2 server blade delivers performance, scalability, and configuration flexibility in this hardware configuration. The server blade hosts the Windows 2012 Hyper-V R2 hypervisor for the guest virtual machines.

Each server blade includes the following:

- Dual-socket Intel Xeon E5-2697 v2 processors
 - 12-core per socket
- 24 × 16 GB DIMM for 384 GB of RAM

Unified Compute Platform Director Network Configuration

The CB520H B2 server blade comes with a single on-board two-channel 10 GbE converged network adapter (CNA) card for network traffic. The CNA card is connected to the Brocade DCB 10 GbE switch modules through the chassis backplane. The Brocade DCB 10 GbE switch modules are then connected to the top-of-rack Brocade 6740 aggregate switches for uplink to the customers' network. The CNA card is configured into four logical NICs per channel, eight NICs per server blade, for performance and redundancy. Each channel has 10 Gb of bandwidth that can be divided across the four logical NICs.

Unified Compute Platform Director was used in conjunction with System Center Virtual Machine Manager 2012 R2 to configure the networks. Five networks were defined using SCVMM 2012 R2 as shown in Figure 2.

VM Networks and IP Pools (5)			
Name	Subnet	Gateway Connection	Available Addresses
Cluster-Netwk		No	
Cluster Pool	192.168.101.0/24		216
DAG-REP-Network		No	
Dag-Rep Pool	10.5.4.0/24		232
Live Migration-Network		No	
Live Migration Pool	192.168.102.0/24		216
Mgmt		No	
VM-Network		No	
VM Network Pool	10.5.3.0/24		220

Figure 2

For performance enhancement and security, the networks were isolated using different VLANs, as follows:

- **VLAN 100 Management Network** — This is for chassis management connections and primary management of the Microsoft Hyper-V hosts.
 - **VLAN 101 Cluster Network** — This is for cluster management connections.
 - **VLAN 102 Live Migration Network** — This is for the Live Migration network.
 - **VLAN 1053 for Client Network** — This is for communication between servers and clients.
 - Microsoft Exchange — Communication between Exchange servers and client access.
 - Microsoft SharePoint — Communication for the Network Load Balancers and clients to Web Front End servers.
 - Microsoft SQL Server — Communication between SQL Server and clients.
 - Microsoft Lync Server — Communication between Lync servers and clients.
 - **VLAN 1054 Network** — This is for replication communication between Exchange servers, and communication between SQL Server, SharePoint, and Web Servers.
 - Microsoft Exchange — Log shipping and database seeding.
 - Microsoft SharePoint — Communication between SQL Server, SharePoint application server and SharePoint Web servers for data.
-

Figure 3 shows the network configuration for the Brocade DCB 10 GbE switch modules 0/1 connected to the Brocade 6720 switches. You might need additional connections in your production environment.

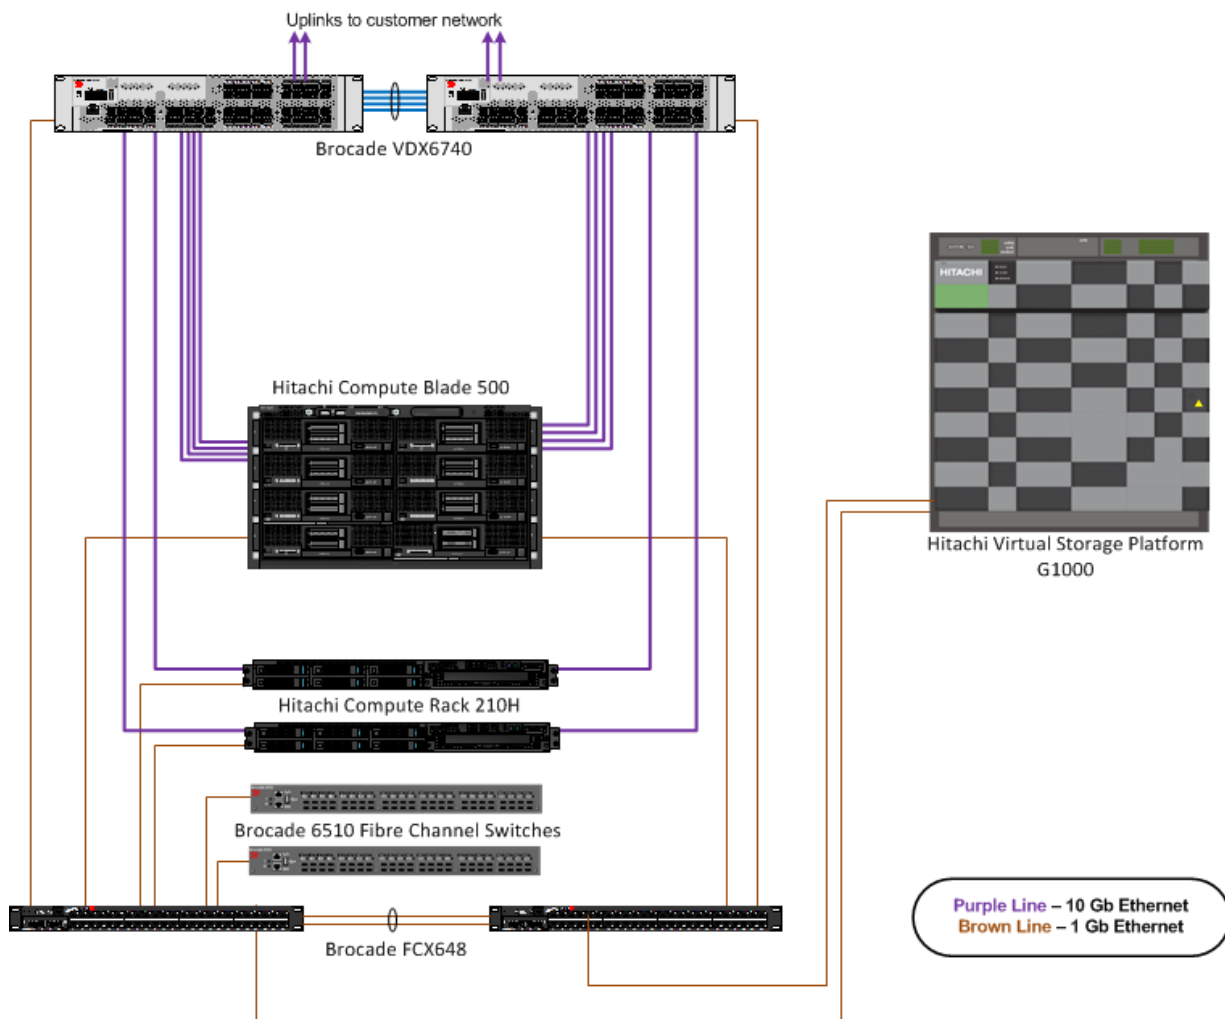


Figure 3

SAN Configuration

The Hitachi Unified Computer Director solution ships with a pre-configured SAN configuration. Table 4 shows the default configuration for the Unified Compute Platform Director management cluster and also logical devices that will be utilized by the UCP computer for command devices that will be attached to Hitachi Tuning Manager and Hitachi Device Manager.

Table 4. Unified Compute Platform Director Default Configuration

<i>Allocated to:</i>	<i>Parity Group</i>	<i>LDEV ID</i>
Hitachi Unified Compute Platform (UCP) Management	1-1	00:FD:00 & 00:FD:01
UCP Compute	1-2	00:FD:02 & 00:FD:03
UCP Compute	1-3	00:FD:04 & 00:FD:05
UCP Compute	1-4	00:FD:06 & 00:FD:07

There are sixteen front-end 8 Gb Fibre Channel ports. Four of these ports will be used for the UCP management servers and the remaining twelve will be used for the tenant node cluster as shown in Table 5.

Table 5. Front End Port Connectivity

<i>Allocation</i>	<i>Ports</i>
Hitachi Unified Compute Platform (UCP) Management Cluster	CL1-A,CL2A CL3-A,CL4A
UCP Compute Blade Cluster (Tenant)	CL5-A,CL6A CL7-A,CL8A CL1-B,CL2B CL3-B,CL4B CL5-B,CL6B CL7-B,CL8B

Hitachi Unified Compute Platform Director is delivered to the customer with a pre-configured high availability SAN architecture. The following best practices are provided by this design to ensure a secure, high-performance, and scalable Microsoft applications deployment:

- Hitachi Unified Compute Platform Director implements dual SAN fabrics, multiple HBAs, and host-based multipathing software. This allows for two or more paths from the application servers connecting to two independent SAN fabrics providing the redundancy required for critical applications.
 - Hitachi Unified Compute Platform Director zones the fabric for multiple, unique paths from HBAs to storage ports. Hitachi Unified Compute Platform Director implements two Fibre Channel switch fabrics to provide multiple independent paths to Hitachi Virtual Storage Platform G1000 to prevent configuration errors from disrupting the entire SAN infrastructure.
 - For SQL, the Queue Depth for all HBAs on each Hyper-V host was changed to 64 per best practice for Hyper-V hosts running SQL Server to match the allocation unit size of the disks of the virtual machines.
-

Figure 4 shows the SAN configuration for this architecture.

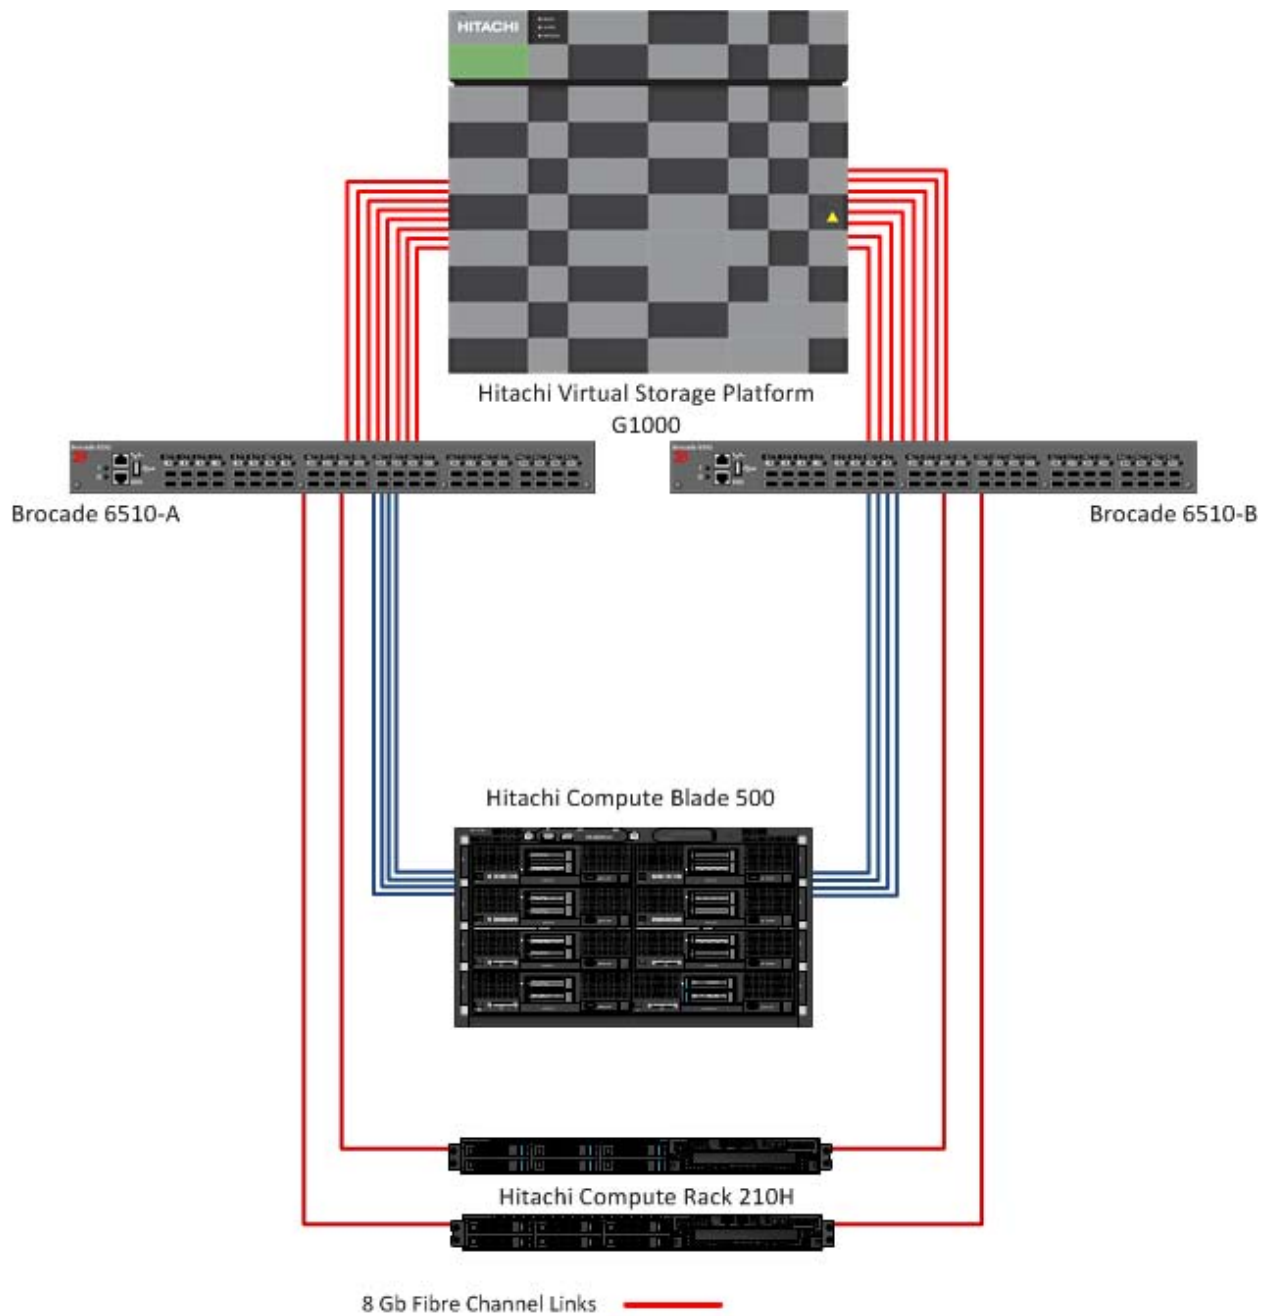


Figure 4

LUN Configuration for Unified Compute Platform Director Management Cluster and Tenant Cluster

For this solution storage LUNs were allocated as Cluster Shared Volumes and presented to the two Hyper-V hosts in the management cluster. The Cluster Shared Volumes in the management cluster contained the storage for all of the management virtual machines required by Unified Compute Platform Director. Also in Pool 127 is the quorum file for both the management cluster and tenant cluster. Table 6 shows the dynamic provisioning pools, LUN sizes, and storage port configuration.

Table 6. LUN Configuration

<i>Dynamic Provisioning Pool</i>	<i>LUN Allocation</i>	<i>LUN Size (GB)</i>	<i>LUN</i>	<i>Hyper-V Cluster</i>
Pool 127	Cluster Shared Volume 1	4000	00	Management
	Mgmt Cluster Quorum Disk	1	01	
	Tenant Cluster Quorum Disk	1	01	Tenant

Storage Resources

The shared storage resource component is one Hitachi Virtual Storage Platform G1000. This resource supports a maximum of 1,152 SFF or LFF drives.

Microsoft Exchange 2013 Storage Resources

The Microsoft Exchange 2013 infrastructure cell for storage resources consists of the following:

- 1 shared Hitachi Virtual Storage Platform G1000
- 60 × 4 TB 7.2k RPM drives
- 5 dynamic provisioning pools

Determining Storage I/O Requirements

The [Exchange 2013 Mailbox Server Role Requirements Calculator](#) was used to determine the storage I/O requirements.

From Table 7, the recommended IOPS per mailbox using the 150 messages sent or received per mailbox per day is 0.101. To ensure that Exchange can provide sufficient overhead for periods of extremely high workload, Microsoft recommends adding 20 percent overhead for production scenarios, for a total of 0.1212 IOPS.

Table 7. Estimated IOPS Requirements for Microsoft Exchange 2013

<i>Messages Sent or Received per Mailbox per Day</i>	<i>Estimated IOPS per Mailbox (Active or Passive)</i>
50	0.034
100	0.067
150	0.101
200	0.134
250	0.168
300	0.201

Table 8 shows the IOPS requirements per database, server, database availability group, and environment to support this reference architecture.

Table 8. Calculated Host I/O and Throughput Requirements

<i>Host I/O and Throughput Requirements</i>	<i>Per Database</i>	<i>Per Server</i>	<i>Per DAG</i>	<i>Per Environment</i>
Total database required IOPS	60	362	1447	1447
Total log required IOPS	13	77	309	309
Database read I/O percentage	60%	N/A	N/A	N/A
Background database maintenance throughput requirements	1 MB/sec	6 MB/sec	24 MB/sec	24 MB/sec

Determining Storage Capacity Requirements

The [Exchange 2013 Mailbox Server Role Requirements Calculator](#) was used to determine the storage capacity requirements. In addition to the requirement for mailbox capacity, you must consider the following:

- **White space**

The database always has free pages (white space) throughout it. During on-line maintenance, the deletion of items marked for removal frees those pages. Estimate the amount of white space in the database by knowing the number of megabytes of mail sent and received by the users with mailboxes in the database.

- **Safety Net**

Safety Net is a new feature in Exchange Server 2013 that prevents data loss by maintaining a queue of successfully delivered messages. It's essentially a delivery queue that exists on each mailbox server. Each time a message is delivered to a database, a copy of that message is held within the Safety Net. The message copy remains in the Safety Net until its expiration date, which is determined by the administrator.

- **Content Indexing**

Content indexing allows fast email items search. This contributes an additional overhead of about 20 percent to the total database size.

Table 9 shows capacity requirements per database, server, database availability group, and environment.

Table 9. Calculated Disk Capacity Requirements

<i>Disk Space Requirements</i>	<i>Per Database (GB)</i>	<i>Per Server (GB)</i>	<i>Per DAG (GB)</i>	<i>Per Environment (GB)</i>
Transport database space required	N/A	386	1545	1545
Database space required	1123	6735	26941	26941
Log space required	55	331	1324	1324
Database and log volume spare required	1712	348	1394	1394

Dynamic Provisioning Pool Configuration

Configure 60 4 TB disks as RAID 10 (2D+2D). Configure 15 parity groups. Each parity group contains four disks and is equal to 7294.38 GB. Create and format four LDEVs out of each parity group. Each LDEV is equal to 1823.59 GB. Provision five Dynamic Provisioning Pools and add the following number of parity group according to Table 10.

Table 10. Dynamic Provisioning Pool Configuration

<i>Dynamic Provisioning Pool</i>	<i>Number of Parity Groups</i>	<i>Number of Disks</i>	<i>Capacity (TB)</i>	<i>Purpose</i>
0	1	4	7.29	Windows Server operating system
1	6	24	43.76	Exchange active databases
2	6	24	43.76	Exchange passive databases
3	1	4	7.29	Exchange active logs
4	1	4	7.29	Exchange passive logs

Volume Configuration

Unified Compute Platform Director was used to configure storage and present new volumes to the Hyper-V failover cluster. This process automates and orchestrates all the administration tasks between the compute, storage, and Brocade fabric. Cluster Shared Volumes were defined as shown in Table 11.

Table 11 shows the detailed cluster shared volume configuration used for Microsoft Exchange 2013.

Table 11. Volume Configuration

<i>Dynamic Pool No.</i>	<i>Host</i>	<i>CSV Number</i>	<i>CSV Size (TB)</i>	<i>Purpose</i>
0	Blade 0	1	3	Microsoft® Windows Server® 2012 R2
	Blade 1			
1	Blade 0	2	25	Microsoft Exchange active databases
	Blade 1			
2	Blade 0	3	25	Exchange passive databases
	Blade 1			
3	Blade 0	4	3	Exchange active logs
	Blade 1			
4	Blade 0	5	3	Exchange passive logs
	Blade 1			

Microsoft SharePoint 2013 Storage Resources

The infrastructure for Microsoft SharePoint 2013 storage resources consists of the following:

- 1 shared Hitachi Virtual Storage Platform G1000
- 48 × 600 GB 10k RPM SAS drives

Determining I/O and Capacity Requirements

The [Capacity Planning for SharePoint Server 2013](#) and [Capacity management and sizing overview for SharePoint Server 2013](#) from Microsoft were used to determine the storage I/O and capacity requirements to support 100,000 users interacting with 20 site collections and 20 content databases of 200 GB each.

Dynamic Provisioning Pool Configuration

All of the dynamic provisioning pools and volumes for this environment are built using the following guidelines:

- 600 GB 10k RPM SAS drives
- RAID-10 (2D+2D) for the following:
 - Microsoft Windows Server operating system
 - Microsoft SQL Server database and log files
 - SharePoint Index volumes
 - SharePoint Search database
- RAID-10 offers the best in performance and reliability

Table 12 shows the number of RAID groups and drives needed to create the dynamic provisioning pools.

Table 12. Number of RAID Groups and Drives for Dynamic Provisioning Pools

<i>Dynamic Provisioning Pool</i>	<i>Number of RAID Groups</i>	<i>Number of Drives</i>	<i>Usable Capacity (TB)</i>	<i>Purpose</i>
20	2	8	2.1 TB	Virtual machine operating system and index
21	7	28	7.3 TB	SQL databases and Temp databases
22	2	8	2.1 TB	Database logs
23	1	4	1.0 TB	Search databases

Volume Configuration

Unified Compute Platform Director was used to configure storage and present new volumes to the Hyper-V failover cluster. Cluster Shared Volumes were defined as shown in Table 13.

Table 13 shows the detailed cluster shared volume configuration used for Microsoft SharePoint 2013.

Table 13. Cluster Shared Volume Configuration

Host	Virtual Machine Name	Pool Number	CSV Size	VHDX Size	Purpose
Blade 3	SP-DB-SPSQL	20	2.1 TB	300 GB	Microsoft® Windows Server® operating system
		21	7.3 TB	210 GB	SPContentDB1
				210 GB	SPContentDB2
				210 GB	SPContentDB3
				210 GB	SPContentDB4
				210 GB	SPContentDB5
				210 GB	SPContentDB6
				210 GB	SPContentDB7
				210 GB	SPContentDB8
				210 GB	SPContentDB9
				210 GB	SPContentDB10
				210 GB	SPContentDB11
				210 GB	SPContentDB12
				210 GB	SPContentDB13
				210 GB	SPContentDB14
				210 GB	SPContentDB15
				210 GB	SPContentDB16
				210 GB	SPContentDB17
				210 GB	SPContentDB18
				210 GB	SPContentDB19
				210 GB	SPContentDB20
				40 GB	TempDB1
				40 GB	TempDB2
40 GB	TempDB3				
40 GB	TempDB4				
40 GB	TempDB5				
40 GB	TempDB6				
40 GB	TempDB7				
40 GB	TempDB8				
40 GB	TempDB9				
40 GB	TempDB10				
40 GB	TempDB11				
40 GB	TempDB12				

Table 13. Cluster Shared Volume Configuration (Continued)

Host	Virtual Machine Name	Pool Number	CSV Size	VHDX Size	Purpose
		22	2.1 TB	50 GB	SPContentLog1
				50 GB	SPContentLog2
				50 GB	SPContentLog3
				50 GB	SPContentLog4
				50 GB	SPContentLog5
				50 GB	SPContentLog6
				50 GB	SPContentLog7
				50 GB	SPContentLog8
				50 GB	SPContentLog9
				50 GB	SPContentLog10
				50 GB	SPContentLog11
				50 GB	SPContentLog12
				50 GB	SPContentLog13
				50 GB	SPContentLog14
				50 GB	SPContentLog15
				50 GB	SPContentLog16
				50 GB	SPContentLog17
				50 GB	SPContentLog18
				50 GB	SPContentLog19
				50 GB	SPContentLog20
Blade 3	SP-DB-SPSC	20	2.1 TB	100 GB	Windows OS
		21	7.3 TB	40 GB	TempDB1
				40 GB	TempDB2
				40 GB	TempDB3
				40 GB	TempDB4
				40 GB	TempDB5
				40 GB	TempDB6
				40 GB	TempDB7
				40 GB	TempDB8
				40 GB	TempDB9
				40 GB	TempDB10
				40 GB	TempDB11
				40 GB	TempDB12
		22	2.1 TB	40 GB	SP Search DB Logs
				100 GB	TempDBLog
23	1.0 TB	100 GB	SP Search DBs		

Table 13. Cluster Shared Volume Configuration (Continued)

Host	Virtual Machine Name	Pool Number	CSV Size	VHDX Size	Purpose
Blade 4	SP-WS-01	20	2.1 TB	80 GB	Windows OS
	SP-WS-02			80 GB	Windows OS
	SP-WS-03			80 GB	Windows OS
	SP-WS-04			80 GB	Windows OS
	SP-WS-05			80 GB	Windows OS
	SP-WS-06			80 GB	Windows OS
	SP-WS-07			80 GB	Windows OS
	SP-WS-08			80 GB	Windows OS

Microsoft Lync 2013 Storage Resources

The infrastructure for Microsoft Lync 2013 storage resources consists of the following:

- 1 shared Hitachi Virtual Storage Platform G1000
- 88 × 600 GB 10k RPM SAS drives

Determining I/O and Capacity Requirements

The server hardware platform recommendations from [Microsoft TechNet](#) were used to size the storage for the Microsoft Lync Server 2013 environment. Microsoft recommends a minimum of eight 10,000 RPM disks for each front end and back end server.

NOTE: Not all types of Lync Server 2013 workloads were included in this reference architecture. Deploying additional workloads will impact sizing and performance requirements.

Dynamic Provisioning Pool Configuration

All of the dynamic provisioning pools and volumes for this environment are built using the following guidelines:

- 600 GB 10k RPM SAS drives
- RAID-10 (2D+2D) for the following:
 - Microsoft Windows Server operating system
 - Microsoft SQL Server database and log files for both the SQL Server 2012 back end database server and the instances of SQL 2012 on the front end servers
- RAID-10 offers the best in performance and reliability

Table 14 shows the number of RAID groups and drives needed to create the dynamic provisioning pools.

Table 14. Number of RAID Groups and Drives for Dynamic Provisioning Pools

<i>Dynamic Provisioning Pool</i>	<i>Number of RAID Groups</i>	<i>Number of Drives</i>	<i>Usable Capacity (TB)</i>	<i>Purpose</i>
10	2	8	2.1 TB	Microsoft® Windows® OS
11	6	24	6.3 TB	Microsoft Lync® Front End Data
12	8	32	8.4TB	Lync Back End Data
13	6	24	6.3 TB	Lync Archive Data

Volume Configuration

Unified Compute Platform Director was used to configure storage and present new volumes to the Hyper-V failover cluster. Cluster Shared Volumes for Microsoft Lync 2013 were defined as shown in Table 15.

Table 15. Microsoft Lync 2013 Cluster Shared Volume Configuration

<i>Host</i>	<i>Virtual Machine Name</i>	<i>Dynamic Pool Number</i>	<i>CSV Size</i>	<i>VHDX Size</i>	<i>Purpose</i>
Blade 5	LyncFE1	10	2.1 TB	200 GB	Microsoft® Windows® OS
		11	6.3 TB	1 TB	Microsoft Lync® Front End Data
	LyncCA	10	2.1 TB	200 GB	Windows OS
Blade 6	LyncFE2	10	2.1 TB	200 GB	Windows OS
		11	6.3 TB	1 TB	Lync Front End Data
	LyncMed	10	2.1 TB	200 GB	Windows OS
Blade 7	LyncFE3	10	2.1 TB	200 GB	Windows OS
		11	6.3 TB	1 TB	Lync Front End Data
	LyncBE	10	2.1 TB	200 GB	Windows OS
		12	8.4 TB	2 TB	Lync Back End Data
		13	6.3 TB	2 TB	Lync Archive Data

Microsoft SQL Server 2014 Storage Resources

The infrastructure cell for Microsoft SQL Server 2014 storage resources consists the following:

- 1 shared Hitachi Virtual Storage Platform G1000
- 272 × 600 GB 10k RPM SAS drives

Determining I/O and Capacity Requirements

The sizing of the storage for the Microsoft SQL Server environment was based on previous testing done in the Hitachi Data Systems labs. While the use of 600 GB drives in this solution results in excess disk capacity, the SQL environment requires the number of spindles to provide the necessary performance and disk throughput to support the workload used to validate this reference architecture.

The dynamic provisioning pools and LUNs for the SQL environment were built using the following guidelines:

- Place database and log files in separate dynamic provisioning pools for best performance.
- Place database and tempdb files in the same dynamic provisioning pool but on separate LUNs.
- Use a RAID-10 drive layout for the creation of dynamic provisioning pools for the best performance and availability.
- Account for approximately 20 percent additional overhead for the database, tempdb, and transaction logs when planning from a capacity perspective.

Dynamic Provisioning Pool Configuration

All of the dynamic provisioning pools and volumes for this SQL environment are built using the following guidelines:

- 600 GB 10k RPM SAS drives
 - RAID-10 (2D+2D) for the following:
 - Microsoft Windows Server operating system
 - Microsoft SQL Server database and log files
 - For best performance:
 - Place the Microsoft Windows Server operating system, SQL database, and SQL log files in separate dynamic provisioning pools
 - RAID-10 offers the best in performance and reliability
 - Reserve an additional two drives for spares
-

Table 16 shows the number of RAID groups and drives used to create the dynamic provisioning pools.

Table 16. Number of RAID Groups and Drives for Dynamic Provisioning Pools

<i>Dynamic Provisioning Pool</i>	<i>Number of Parity Groups</i>	<i>Number of Disks</i>	<i>Capacity</i>	<i>Purpose</i>
8	1	4	1.2 TB	Microsoft® Windows Server® operating system
9	36	144	43.2 TB	SQL LG Data Pool
10	18	72	21.6 TB	SQL MD Data Pool
11	9	36	10.8 TB	SQL SM Data Pool
12	2	8	2.4 TB	Shared SQL Log Pool
13	2	8	2.4 TB	Shared TempDb Pool

Volume Configuration

Unified Compute Platform Director is used to configure storage and present new volumes to the Hyper-V failover cluster. This process automates and orchestrates all the administration tasks between the compute, storage, and Brocade fabric. Cluster Shared Volumes used for SQL Server 2014 were defined as shown in Table 17.

Table 17. XSQL Server 2014 Cluster Shared Volume Configuration

<i>Dynamic Pool Number</i>	<i>CSV Number</i>	<i>CSV Size</i>	<i>Purpose</i>
8	1	1.2 TB	For Microsoft® Windows® OS Boot LUN
9	2	3 TB	For SQL SM Data file placement
10	3,4	3 TB	For SQL MD Data file placement
11	5,6,7,8	3 TB	For SQL LG Data file replacement
12	9	2.4 TB	For SQL Log file placement
13	10	2.4 TB	For SQL TempDB file placement

Application Infrastructure for Microsoft Exchange Server 2013

The application infrastructure for Microsoft Exchange consists of the compute, storage, and network components necessary to start building out the Microsoft Hyper-V infrastructure to run Microsoft Exchange workloads.

Virtual Machine Configuration

Configure the same vCPU, memory, and vNIC for each Exchange virtual machine as shown in Table 18.

Table 18. Virtual Machine Configuration

<i>Virtual Machine Name</i>	<i>vCPU</i>	<i>vRAM</i>	<i>vNIC</i>	<i>Purpose</i>
MBXCAS1	10	90	2	Combined mailbox and client access server
MBXCAS2				
MBXCAS3				
MBXCAS4				

Virtual Machine Processor and Memory Configuration

The [Exchange 2013 Mailbox Server Role Requirements Calculator](#) and [Processor Query Tool v1.1](#) were used to determine the computing CPU and memory requirements.

Microsoft recommends disabling hyper threading because it causes capacity planning and monitoring challenges. As a result, the expected gain in CPU overhead is not likely to be justified.

Using two 12-core Intel Xeon 2697 v2 processors, the Processor Query Tool recommended 957 as the spec rate value. Allocate ten vCPUs with 90 GB memory to each of the four Microsoft Exchange virtual machines.

Note — Microsoft recommends an additional 10% CPU overhead for Exchange virtual machines running in a Hyper-V environment.

Virtual Machine Hard Disk Configuration

Create the VHDX disks as fixed disk for performance. Assign them to the following virtual machines as shown in Table 19.

Table 19. VHDX Disk Configuration

Microsoft® Hyper-V® Host	Virtual Machine Name	CSV Number	vDisk VHDX	Purpose			
0	DC1	1	200 GB	Microsoft Windows Server® 2012 R2			
	DC2						
0	MBXCAS1	1	200 GB	Windows Server 2012 R2			
		2	1800 GB	Db1 (active)			
				Db2 (active)			
				Db3 (active)			
		3	1800 GB	Db10 (passive)			
				Db11 (passive)			
				Db12 (passive)			
		4	200 GB	Log1 (active)			
				Log2 (active)			
	Log3 (active)						
	5	200 GB	Log10 (passive)				
			Log11 (passive)				
			Log12 (passive)				
	MBXCAS2	1	200 GB	Windows Server 2012 R2			
					2	1800 GB	Db4 (active)
							Db5 (active)
		Db6 (active)					
		3	1800 GB	Db1 (passive)			
Db2 (passive)							
Db3 (passive)							
4		200 GB	Log4 (active)				
			Log5 (active)				
			Log6 (active)				
5		200 GB	Log1 (passive)				
			Log2 (passive)				
	Log3 (passive)						

Table 19. VHDX Disk Configuration (Continued)

Microsoft® Hyper-V® Host	Virtual Machine Name	CSV Number	vDisk VHDX	Purpose
1	MBXCAS3	1	200 GB	Windows Server 2012 R2
		2	1800 GB	Db7 (active)
				Db8 (active)
				Db9 (active)
				Db4 (passive)
				Db5 (passive)
				Db6 (passive)
		4	200 GB	Log7 (active)
				Log8 (active)
				Log9 (active)
				Log4 (passive)
		5		Log5 (passive)
	Log6 (passive)			
	MBXCAS4	1	200 GB	Windows Server 2012 R2
		2	1800 GB	Db10 (active)
				Db11 (active)
				Db12 (active)
				Db7 (passive)
				Db8 (passive)
Db9 (passive)				
4		200 GB	Log10 (active)	
			Log11 (active)	
			Log12 (active)	
			Log7 (passive)	
5			Log8 (passive)	
	Log9 (passive)			

Virtual Machine Network Configuration

Each Exchange virtual machine is configured to use two virtual network adapters. Adapter 1 is configured for MAPI connection access. Adapter 2 is configured for database replication. All of the blades share the same network configuration. This allows for failover of any of the Exchange virtual machines to any host in the failover cluster.

Microsoft Exchange Database Availability Group Configuration

Four Exchange virtual machines are configured to support 6,000 mailboxes using an active/active two database copies model. In the event of a server blade or virtual machine failure, Exchange continues to provide emails to end users without any disruption.

Configure active and passive databases for four Exchange virtual machines, as follows:

- **Exchange Virtual Machine: MBXCAS1**
 - DB1-3 as active databases
 - DB10-12 as passive database copies
- **Exchange Virtual Machine: MBXCAS2**
 - DB4-6 as active databases
 - DB1-3 as passive database copies
- **Exchange Virtual Machine: MBXCAS3**
 - DB7-9 as active databases
 - DB4-6 as passive database copies
- **Exchange Virtual Machine: MBXCAS4**
 - DB10-12 as active databases
 - DB7-9 as passive database copies

Each database in this DAG configuration supports 500 mailboxes. This design balances CPU and memory resources to provide better performance.

Figure 5 shows the database distribution for a four-member DAG configuration.

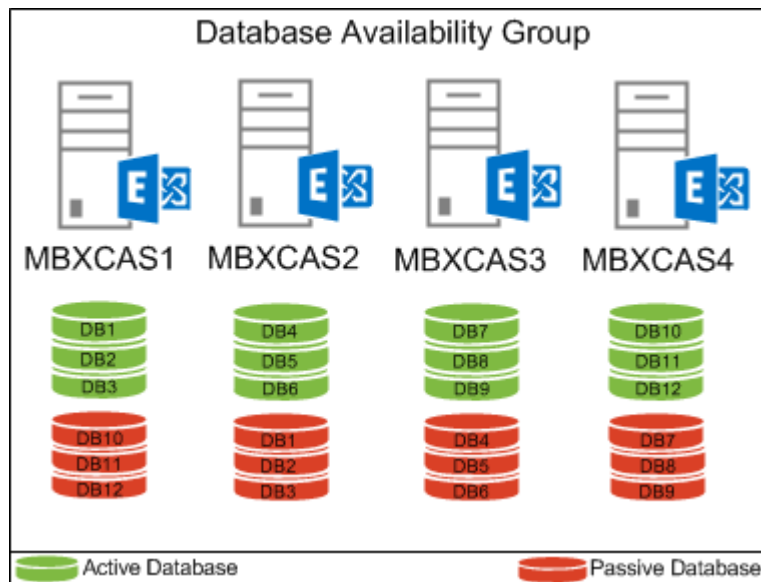


Figure 5

Application Infrastructure for Microsoft SharePoint Server 2013

The application environment for Microsoft SharePoint Server 2013 consists of the compute, storage, and network components necessary to start building out the Hyper-V infrastructure to run Microsoft SharePoint workloads.

Virtual Machine Configuration

For the environment used to validate this reference architecture, blades 3 and 4 of the Hitachi Compute Blade 500 are configured to run the Microsoft SharePoint virtual machines.

- **Blade 3** hosts the virtual machine running Microsoft SQL Server 2014, Data Center Edition. It is the main SQL server for the SharePoint farm.
- **Blade 3** also hosts a virtual machine running Microsoft SharePoint 2013 and Microsoft SQL Server 2012 R2, Data Center Edition. It is the application and search crawl server, with local (in SQL) search tables, and it hosts the central administration SharePoint site.
- **Blade 4** hosts eight virtual machines serving as Microsoft SharePoint 2013 web front-end servers.

All virtual machines for the SharePoint environment have Microsoft Windows 2012 R2, Standard Edition, as the installed operating system.

Table 20 shows the virtual machine configuration with vCPU, vRAM, and vNIC allocation.

Table 20. Virtual Machine Configuration

<i>Host</i>	<i>VM Name</i>	<i>vCPU</i>	<i>vRAM</i>	<i>vNIC</i>
Blade 3	SP-DB-SPSQL	12	150 GB	1
	SP-DB-SPSC	12	150 GB	1
Blade 4	SP-WS-01	4	40 GB	2
	SP-WS-02	4	40 GB	2
	SP-WS-03	4	40 GB	2
	SP-WS-04	4	40 GB	2
	SP-WS-05	4	40 GB	2
	SP-WS-06	4	40 GB	2
	SP-WS-07	4	40 GB	2
	SP-WS-08	4	40 GB	2

Virtual Machine Processor Configuration

The [Capacity Planning for SharePoint Server 2013](#) and [Capacity management and sizing overview for SharePoint Server 2013](#) from Microsoft were used to determine the computing requirements to support a 100,000 user Microsoft SharePoint 2013 server farm with 20 site collections using 20 200 GB content databases.

The server blade uses dual E5-2697 Intel Xeon processors (12-core).

Virtual Machine Memory Configuration

The [Capacity Planning for SharePoint Server 2013](#) and [Capacity management and sizing overview for SharePoint Server 2013](#) from Microsoft were used to determine the computing requirements to support a 100,000 user Microsoft SharePoint 2013 server farm with 20 site collections using 20 content databases of 200 GB.

The following are considerations for virtual memory configuration for Microsoft SharePoint Server 2013:

- Depending on the data, sites, and user security profiles in your environment, a WFE with 40 GB of memory allows for the following:
 - 10 GB for object cache.
 - 20 GB for SharePoint WFE services.
 - 10 GB to extend search components or any of the other 3 caches configurable for SharePoint Server WFEs.
 - Object cache is configured on the application server and all WFE servers to 10 GB.
 - If the search topology needs to be extended, each search component will consume memory.
 - Having a search index on each WFE increases search response time but increases processor utilization. Only extend search indexes to each WFE, as done in this reference architecture, if your index is above 10 million items.
 - If you have up to 10 million items, one partition suffices.
 - If between 10 million to 40 million items, use up to four partitions.
 - Refer to [Enterprise search architectures for SharePoint Server 2013](#) for further detailed explanation as well as additional planning guidance beyond the scope of this reference architecture.
-

- Distributed cache is reserved only to the application server in this configuration. Distributed cache, which can be extended to any server in the environment, depends on the following:
 - The number of user profiles
 - The way in which social media is used and designed within your environment

Refer to [Manage the Distributed Cache service in SharePoint Server 2013](#) for further explanation and consideration on configuring and extending this functionality of SharePoint within your environment.

- Set distributed cache to 10% of the server memory.
- If planning to have a dedicated distributed cache server for your environment, follow these best practices:
 - Determine the total physical memory on the server. This solution allocates 16 GB of memory to the server.
 - Reserve 2 GB of memory for other processes and services that are running on the cache host.

For example: Using the amount of memory allocated to the service in the reference architecture, you have 16 GB – 2 GB = 14 GB. Allocate 14 GB to the distributed cache service.

- Set half of the remaining memory (in megabytes) as the size of the distributed cache service.

For example: Using the amount of memory allocated to the distributed cache service in this reference architecture, you have 14 GB / 2 = 7 GB. Use 7000 MB as the cache size of the distributed cache service.

- Blob cache and page output cache are both outside the scope of this reference architecture. This is due to the amount of variations that your data, sites, and security profiles directly influence.

Both of these cache types can increase response times, depending on the data and sites within as well as the configuration of your environment. Allocate memory for both, as required.

See [Cache settings operations in SharePoint Server 2013](#) to determine if and what size to configure each cache.

Virtual Machine Hard Disk Configuration

Create LUNs in each dynamic provisioning pool as required and configure each LUN as a cluster shared volume in the failover cluster for failover and redundancy. See Volume Configuration on page 27 for the dynamic provisioning pool, cluster shared volume, and VHDX configuration.

- Configure the disks as fixed VHDX. When originally created within Hyper-V, format the disks as NTFS.
- On the SQL server use multiple SCSI controllers for the database, tempDB, and log drives.
- On the SharePoint application server spread the TempDB drives across two SCSI controllers.

Virtual Machine Network Configuration

All blades share the same network configuration. This allows for failover of any of the SharePoint virtual machines to any host in the failover cluster.

Microsoft SharePoint Server 2013 Configuration

This is how to configure Microsoft SharePoint Server 2013 in this environment.

Microsoft SQL Server Architecture

These are best practice recommendations from Microsoft for storage and SQL Server capacity planning and configuration for Microsoft SharePoint Server 2013. For further explanation on the following requirements, see [Storage and SQL Server capacity planning and configuration \(SharePoint Server 2013\)](#) from Microsoft.

■ **Operating System Disk Management**

All disks created within the operating system **must** have an allocation unit size of 64k.

■ **SQL Server Instance and Database Settings**

- You must do this step during the installation of Microsoft SQL Server. It cannot be changed afterwards, even with a complete rebuild of the master database.
-

During installation select the following:

■ **Collation** — Latin1_General_CI_AS_KS_WS

- For Microsoft SQL Server running SQL databases and content databases, set memory to 90% of the total memory of the server.
- For Microsoft SQL Server running Microsoft SharePoint search databases, set memory to no more than 40% of the total memory of the server.
- Under **Database Settings**, set MAXDOP to 1 for both SQL Server instances.
- Disable **auto create statistics** and **auto update statistics** for both SQL Server instances. Microsoft SharePoint handles this functionality.
- Enable start-up parameter **-T1118**. This instructs SQL Server to use a round-robin allocation strategy for the tempDB file, so tempDB files maintain the same size. This reduces resource allocation contention in the tempDB database to improve performance on complex queries.

■ **Microsoft SQL Server Database Files**

Create all content databases with the **initial size** set to 200 GB and **auto growth** set to 1000 MB.

Set **MAXSIZE** for each database file to a value that matches the capacity of the volume.

■ **Microsoft SQL Server Transaction Log File**

Create all content database logs as 20% of the size of the content database. Set the **initial size** set to 40 GB and set **auto growth** to 200 MB.

- **Microsoft SQL Server tempDB Files**

By default, tempDB only supports a single data file group and a single log file group with a default number of files set to 1. Microsoft recommends creating at least as many data files of equal size as the server has CPU cores. The number of concurrent threads is less than or equal to the number of CPU cores.

- **General Database Maintenance**

To keep your databases properly maintained, follow these recommended practices:

- **Monitor the database server** to make sure that it responds appropriately and **is not overloaded**. Key performance counters to monitor include the following:
 - **Network wait queue** — 0 or 1
 - **Average disk queue length (latency)** — Less than 20 msec
 - **Memory used** — Less than 70%
 - **Free disk space** — More than 25% for content growth
- Do not auto-shrink databases or set up any maintenance plans that *programmatically* shrink your databases.
- Shrink a database only when 50% or more of the content in it has been removed by user or administrator deletions.
- Shrinking databases is very resource intensive. It requires careful scheduling.
- Only shrink content databases.
- The configuration, central administration, and search databases do not usually experience enough deletions to contain sufficient free space.
- Avoid needing to shrink databases by **including growth allocations in your capacity planning**, including an overhead allocation of 10% to 20%.
- SharePoint has a two level recycle bin. Deleting items from a SharePoint site does not free up storage capacity until the item is deleted from both levels of the recycle bin, if the second level is enabled. The recycle bin is configured in the General settings for the web application.

For more information, see the Microsoft *TechNet* article "[Monitoring and maintaining SharePoint Server 2013.](#)"

Microsoft SharePoint Server 2013 Operating System Additional Configurations

Follow these best practices:

- For Microsoft Windows Server 2012 R2 to be able to open document repositories in Windows Explorer, enable Desktop Experience in Server Role.
 - To allow downloading and uploading of any document over 50 MB in Microsoft Windows Server 2012 R2, configure all servers within the SharePoint topology following the instructions on how to [enable download and upload of files over 50mb](#) from Microsoft.
 - When running NLB for your SharePoint farm, disable LoopBack to avoid a 401.1 error to the clients. Do the following on all WFE servers in the topology:
 - Add the URL for SharePoint to **Trusted Sites** in Microsoft Internet Explorer® for all WFE servers.
 - Run the following at a command prompt:

```
New-ItemProperty  
HKLM: \System\CurrentControlSet\Control\Lsa -  
Name "DisableLoopbackCheck" -value "1" -  
PropertyType dword
```
-

Application Infrastructure for Microsoft Lync Server 2013

The application environment for Microsoft Lync Server 2013 consists of the compute, storage, and network components necessary to start building out the Hyper-V infrastructure to run Microsoft Lync workloads.

Virtual Machine Configuration

For the environment used to validate this reference architecture blades 5, 6, and 7 of the Hitachi Compute Blade 500 are configured to run the Microsoft Lync virtual machines.

- **Blade 5** hosts one of the Lync 2013 front end servers and the certificate authority server required for a Lync environment.
- **Blade 6** hosts the second Lync 2013 front end server and a dedicated Lync 2013 Mediation server.
- **Blade 7** hosts the third Lync 2013 front end server and the Lync back end server running Microsoft SQL Server 2012.

All virtual machines for the Lync environment have Microsoft Windows 2012 R2, Standard Edition, as the installed operating system. Table 21 shows the virtual machine configuration with vCPU, vRAM, and vNIC allocation.

Table 21. Microsoft Lync Server 2013 Virtual Machine Configuration

<i>Host</i>	<i>VM Name</i>	<i>vCPU</i>	<i>vRAM</i>	<i>vNIC</i>
Blade 5	LyncFE1	12	32 GB	1
	LyncCA	4	8 GB	1
Blade 6	LyncFE2	12	32 GB	1
	LyncMed	8	16 GB	1
Blade 7	LyncFE3	12	32 GB	1
	LyncBE	12	32 GB	1

Virtual Machine Processor Configuration

The vCPU sizing for the Lync 2013 servers was based on [Microsoft recommendations](#).

The server blades use dual E5-2697 Intel Xeon processors (12-core).

Virtual Machine Memory Configuration

The virtual memory sizing for the Lync 2013 servers was based on [Microsoft recommendations](#).

Virtual Machine Hard Disk Configuration

Create LUNs in each dynamic provisioning pool as required and configure each LUN as a cluster shared volume in the failover cluster for failover and redundancy.

Configure the disks as fixed VHDX. When originally created within Hyper-V, format the disks as NTFS.

Virtual Machine Network Configuration

All three blades share the same network configuration. This allows for failover of any of the Lync virtual machines to any host in the failover cluster.

Microsoft Lync Server 2013 Configuration

There are a large number of variables that must be taken into account when designing and deploying a Lync environment. For information see the [Microsoft Lync Server 2013](#) site on TechNet. You can also download the document titled [Planning a Lync Server 2013 Deployment on Virtual Servers](#) for information on planning and testing a virtualized Lync 2013 environment. This document also includes information on the Key Health Indicators (KHIs) that need to be monitored. This paper covers one specific design scenario for a 13,000 user environment. The basic design is based on an N+1 design to ensure acceptable performance if one Lync 2013 front end server goes offline for any reason. The sizing is based on the Microsoft recommendation that you should have one front end server for every 6,660 users.

Application Infrastructure for Microsoft SQL Server 2014

The application infrastructure for Microsoft SQL Server 2014 consists of the compute, storage, and network components necessary to start building out the Hyper-V infrastructure to run Microsoft SQL Server workloads.

This configuration used a single CB520 B3 blade with two 12-core Intel Xeon 2697 v2 processors with 256 GB of RAM for the Hyper-V host in a clustered configuration. Other blades in the Hyper-V Failover Cluster are used to satisfy the clustering configuration, and allow for failover or live migration of SQL Server 2014 to other blades in the cluster.

Virtual Machine Configuration

Described in this section are the CPU, memory, network, and disk configuration for each SQL Server virtual machine.

Virtual CPU, Memory, and vNIC Configuration

Configure the vCPU, memory, and vNIC for each SQL Server virtual machine as shown in Table 22.

Table 22. Microsoft SQL Server 2014 Virtual Machine Configuration

<i>Host</i>	<i>VM Name</i>	<i>vCPU</i>	<i>vRAM</i>	<i>vNIC</i>	<i>Purpose</i>
Blade 2	SQL SM	4	32 GB	2 vNIC, load-balancing	Microsoft SQL Server for small configuration
	SQL MD	8	64 GB	2 vNIC, load-balancing	SQL Server for medium configuration
	SQL LG	16	128 GB	2 vNIC, load-balancing	SQL Server for large configuration

Virtual Machine Processor Configuration

The number of vCPUs allocated to the Microsoft SQL Server virtual machine and the size of the database to be supported is based on previous testing performed in the Hitachi Data Systems labs. The previous testing established the number of vCPUs based on the workload for a specific database size and user workload for that size database.

Virtual Machine Hard Disk Configuration

Create the VHDX disks as fixed disk for performance. Assign them to the following virtual machines as shown in Table 23.

Table 23. Virtual Machine Hard Disk Configuration

<i>Virtual Machine Name</i>	<i>CSV Number</i>	<i>vDisk VHDX</i>	<i>Purpose</i>
SQL SM	1	200 GB	Microsoft Windows Server® 2012 R2
	2	1800 GB	Data01
			Data02
			Data03
			Data04
	12	600 GB	Log01
	13	80 GB	Tempdb01
			Tempdb02
			Tempdb03
			Tempdb04

Table 23. Virtual Machine Hard Disk Configuration (Continued)

<i>Virtual Machine Name</i>	<i>CSV Number</i>	<i>vDisk VHDX</i>	<i>Purpose</i>
SQL MD	1	200 GB	Windows Server 2012 R2
	3	1800 GB	Data01
			Data02
			Data03
			Data04
	4	1800 GB	Data05
			Data06
			Data07
			Data08
	12	600 GB	Log01
	13	80 GB	Tempdb01
			Tempdb02
			Tempdb03
			Tempdb04
			Tempdb05
			Tempdb06
Tempdb07			
Tempdb08			

Table 23. Virtual Machine Hard Disk Configuration (Continued)

<i>Virtual Machine Name</i>	<i>CSV Number</i>	<i>vDisk VHDX</i>	<i>Purpose</i>
SQL LG	1	200 GB	Windows Server 2012 R2
	5	1800 GB	Data01
			Data02
			Data03
			Data04
	6	1800 GB	Data05
			Data06
			Data07
			Data08
	7	1800 GB	Data09
			Data10
			Data11
			Data12
	8	1800 GB	Data13
			Data14
			Data15
			Data16
12	600 GB	Log01	

Table 23. Virtual Machine Hard Disk Configuration (Continued)

<i>Virtual Machine Name</i>	<i>CSV Number</i>	<i>vDisk VHDX</i>	<i>Purpose</i>
SQL LG	13	80 GB	Tempdb01
			Tempdb02
			Tempdb03
			Tempdb04
			Tempdb05
			Tempdb06
			Tempdb07
			Tempdb08
			Tempdb09
			Tempdb10
			Tempdb11
			Tempdb12
			Tempdb13
			Tempdb14
			Tempdb15
			Tempdb16

Virtual Machine Network Configuration

All blades share the same network configuration. This allows for failover of any of the SQL virtual machines to any host in the failover cluster.

Engineering Validation

This section describes the test methodologies and test results used to validate this solution.

Microsoft Exchange 2013

This describes the validation of Microsoft Exchange 2013 in this environment.

Test Methodology

This is the test methodology used to validate storage disk subsystem and compute for this solution.

Jetstress Disk Subsystem Test

The purpose of using [Microsoft Exchange Server Jetstress 2013](#) is to validate the storage disk subsystem performance IOPS to support 6,000 mailboxes using two database copies on four Exchange virtual machines.

Microsoft recommends verifying the performance and stability of a disk subsystem prior to putting an Exchange server into production.

Table 24 shows the Jetstress parameters configuration. The testing goal was to verify that the achieved IOPS meet or exceed the target IOPS with latency less than 20 msec.

Table 24. Jetstress Parameters

<i>Jetstress Parameter</i>	<i>Value</i>
Number of Exchange virtual machines	4
Number of databases per host	6
Number of copies per database	2
Thread count	20
Target IOPS per Exchange virtual machine	362
Total IOPS per database availability group	1447
Duration	8 hours

Load Generator Peak Load Performance Tests

[Microsoft Exchange Load Generator 2013](#) is used as a simulation tool to introduce various types of client workloads into a non-production Exchange environment. This test sends multiple messaging requests from client virtual machines to the Exchange servers to see how Exchange responds to email loads.

Table 25 lists the Load Generator parameters used in configuring the tests.

Table 25. Load Generator Parameters

<i>Load Generator Parameter</i>	<i>Value</i>
Total number of mailboxes	6,000
Client type	Microsoft Outlook® 2007 Online
Action profile	Microsoft Outlook 150
Mailbox size	2 GB
Length of simulation day	10 hours
Total length of simulation	8 hours

Test Results

This section provides Jetstress and Loadgen test results.

Jetstress Disk Subsystem Throughput Test Results

The Jetstress disk subsystem throughput test results show that the storage is able to deliver the IOPS with latency less than 20 msec to support this reference architecture.

Table 26 and Table 27 show the Jetstress results.

Table 26. Database Sizing and Throughput Jetstress Results

<i>Database Sizing and Throughput</i>	<i>MBXCAS1</i>	<i>MBXCAS2</i>	<i>MBXCAS3</i>	<i>MBXCAS4</i>
Achieved Transactional I/O per Second	884.41	883.60	892.54	896.90
Capacity Percentage	100%	100%	100%	100%
Throughput Percentage	100%	100%	100%	100%

Table 27. Total I/O Performance Jetstress Results

<i>Virtual Machine Name</i>	<i>Exchange DB Instances</i>	<i>I/O DB Reads Average Latency (msec)</i>	<i>I/O DB Writes Average Latency (msec)</i>	<i>I/O DB Reads/sec</i>	<i>I/O DB Writes/sec</i>	<i>I/O DB Reads Average Bytes</i>	<i>I/O DB Writes Average Bytes</i>
MBXCAS1	Instance3004.1	14.210	1.367	102.077	45.300	33677.120	34175.262
	Instance3004.2	13.821	1.274	102.005	45.096	33700.977	34175.946
	Instance3004.3	13.797	1.168	102.079	45.198	33682.628	34172.848
	Instance3004.4	15.555	1.412	102.247	45.376	33705.868	34162.604
	Instance3004.5	15.545	1.297	102.149	45.298	33691.585	34163.101
	Instance3004.6	15.580	1.291	102.183	45.402	33693.220	34167.317

Table 27. Total I/O Performance Jetstress Results (Continued)

<i>Virtual Machine Name</i>	<i>Exchange DB Instances</i>	<i>I/O DB Reads Average Latency (msec)</i>	<i>I/O DB Writes Average Latency (msec)</i>	<i>I/O DB Reads/sec</i>	<i>I/O DB Writes/sec</i>	<i>I/O DB Reads Average Bytes</i>	<i>I/O DB Writes Average Bytes</i>
MBXCAS2	Instance1556.1	13.540	1.303	111.500	45.321	52706.982	34201.866
	Instance1556.2	13.912	1.265	111.302	45.157	52699.533	34213.259
	Instance1556.3	13.640	1.164	111.062	44.966	52771.431	34224.134
	Instance1556.4	15.502	1.409	111.219	45.152	52702.193	34216.711
	Instance1556.5	15.441	1.332	111.144	45.123	52766.594	34216.832
	Instance1556.6	15.569	1.329	111.230	45.109	52735.275	34213.120
MBXCAS3	Instance2808.1	13.488	1.276	111.946	45.581	52502.718	34190.937
	Instance2808.2	14.246	1.206	112.094	45.866	52490.144	34188.006
	Instance2808.3	15.504	1.368	112.109	45.677	52481.847	34190.518
	Instance2808.4	13.256	1.086	112.156	45.668	52487.410	34191.547
	Instance2808.5	15.218	1.259	112.240	45.768	52472.778	34183.701
	Instance2808.6	15.792	1.265	112.335	45.784	52420.815	34176.950
MBXCAS4	Instance1980.1	13.465	1.291	112.537	45.958	52470.890	34182.034
	Instance1980.2	13.357	1.219	112.800	45.958	52391.510	34169.825
	Instance1980.3	14.225	1.157	112.664	46.066	52428.565	34176.214
	Instance1980.4	15.452	1.377	112.514	45.817	52428.054	34182.272
	Instance1980.5	15.890	1.329	112.650	45.982	52416.761	34176.821
	Instance1980.6	15.171	1.278	112.709	45.926	52417.465	34174.972

Load Generator Peak Load Performance Test Results

The load generator test results show that the two Hyper-V hosts running four Exchange virtual machines were able to support this reference architecture successfully.

Table 28 shows the Load Generator test results.

Table 28. Memory Percent Committed and CPU Utilization Load Generator Results

<i>VM Name</i>	<i>Number of Mailboxes</i>	<i>CPU % Committed</i>	<i>Memory % Utilization</i>	<i>Test Result</i>
MBXCAS1	1,500	31.64%	44.86%	Passed
MBXCAS2	1,500	22.82%	39.30%	Passed
MBXCAS3	1,500	21.73%	43.46%	Passed
MBXCAS4	1,500	22.41%	43.55%	Passed

Microsoft SharePoint Server 2013

This section describes the validation of Microsoft SharePoint Server 2013 in this environment.

Test Methodology

Hitachi Data Systems used Microsoft Visual Studio® 2012 Ultimate edition with its Visual Studio test agents to test this reference architecture. Different load tests created a user load similar to what most companies would see on a production Microsoft SharePoint 2013 environment.

Table 29 describes the user load mix during testing.

Table 29. User Load Details

<i>Load Type</i>	<i>Percentage</i>
Browsing	75%
Upload	3%
Search	11%
Open	11%

For these tests, the Visual Studio host, the test agent, and test controller hosts were in a separate network dedicated for user data connectivity to the web servers in the SharePoint environment. Ten test agents evenly drove the load mix to all 20 site collections in the environment.

A total of 100,000 users were created and given contributor permission to the site collections in the environment. A 1% user concurrency rate was set. This means that at any given time, there were 1,000 users connected to the SharePoint environment using the load mix described in Table 29 above.

Data population and growth utilized real world files and documents of different types and sizes to ensure the closest similarity to real world environments. The files used to populate the content databases and test the environment had variable sizes ranging from a couple kilobytes to a couple megabytes. Testing maintained a 100 MB default limitation for upload size.

The content databases were attached to each of 20 site collections. Each was loaded with about 60 GB of files. There was a repository set aside of random files for the upload tests. During the collaboration test execution, the data files were randomly uploaded to all of the site collections by random test users.

The primary purpose of the test methodology validated whether the following was true for the environment, based on the calculation of a heavy user profile:

- Withstand a 1% user concurrency and execution of a mixed load type
 - Maintain acceptable response time levels
-

This means that overall the test needed to achieve and maintain an average of 34 requests per second.

Table 30 lists acceptable page response times that were used for the testing of the environment.

Table 30. Acceptable Page Response Times

<i>Load Type</i>	<i>Acceptable Page Response Time</i>
Browsing	Less than 1 second
Upload	Variable (dependent on file size)
Search	Less than 5 seconds (using the default SharePoint search page)
Open	Variable (dependent on file size)

The search tests used the SharePoint search page to execute searches against the different content databases as well as the overall SharePoint environment. Because documents ranged from 10 KB to 5 MB, there were various acceptable response times for the different load types.

A *modify workload* type uses the values for open and upload, as the overall act of modifying a file by individual users depends on the number of changes done by the user.

Although the primary purpose of the test was to verify that the response times for each load type remained at or below acceptable values, testing also monitored other environment resources during testing to verify that there were not bottlenecks.

Test Results

The test results show that the environment successfully handled the user load within the acceptable response time values.

Table 31 lists the response times for each of the load types.

Table 31. Load Type Response Times

<i>Load Type</i>	<i>Response Time</i>
Browsing	< 1 second
Upload	Variable < 2 second
Search	< 1 second
Open	Variable < 2 second

Microsoft Lync Server 2013

This describes the validation testing of Microsoft Lync Server 2013 in this environment.

Test Methodology

Hitachi Data Systems used the Lync Server Stress and Performance Tool to test the Lync 2013 environment by simulating workloads. There are many types of workloads with many possible variables that can be simulated. The testing for this reference architecture included the following workloads configured using the Lync Server 2013 Load Configuration Tool:

- Instant Messaging - Load Level set to High with conferencing and large conference enabled
- Audio Conferencing - Load Level set to High with adhoc and large conference enabled
- Application Sharing - Load Level set to High with adhoc and large conference enabled
- Data Collaboration - Load Level set to High with adhoc and large conference enabled
- Distribution List Expansion - Load Level set to High
- Address Book Web Query - Load Level set to High

Users were created in Microsoft Active Directory® with the Lync Server Stress Test - User Provisioning Tool using values based on Appendix A of the *Planning a Lync Server 2013 Deployment on Virtual Servers* document published by Microsoft.

Five client servers were used to run the batch files created using the Lync Server 2013 Load Configuration Tool.

The success of a test was verified by comparing data from the tests to the Lync Key Health Indicators. Information on the Key Health Indicators can be found [here](#). A poster documenting what Key Health Indicators are, how to collect them, and details on which indicators are relevant to which Lync server roles can be downloaded can be found [here](#).

Test Results

The test results show that the Lync 2013 environment successfully handled the user load with all Lync 2013 key health indicators within acceptable values.

Microsoft SQL Server 2014

This describes the validation testing of Microsoft SQL Server 2014 in this environment.

Test Methodology

To validate this reference architecture, an industry-standard OLTP workload that simulates a stock brokerage scenario was exercised on the SQL Server database. This section describes the results of the validation.

Table 32 lists the OLTP workload parameters used as a baseline for the tests.

Table 32. OLTP Workload Parameters

<i>Parameter</i>	<i>Description</i>
Test scenario	OLTP workload for a stock brokerage firm
Virtual machine CPU utilization	≥50%
Minimum database LU capacity usage	80%
Minimum individual disk busy rate	50%
Number of active users	Variable, scaled with architecture size
Test type	Performance
Test duration	≥4 hours

Table 33 lists results for the OLTP workload tests for the 80,000-user configuration.

Table 33. OLTP Workload Test Results for 80,000-user Configuration

<i>Metric</i>	<i>Microsoft® Success Criterion</i>	<i>Result</i>
Achieved Database Disk IOPS	Varies	16,000 IOPS
Database Avg. Disk Seconds/Read	≤ 20 msec	1.5 msec
Database Avg. Disk Seconds/Write	≤ 20 msec	0 msec
Transaction Log Avg. Disk Seconds/Read	≤ 5 msec	N/A
Transaction log Avg. Disk Seconds/Write	≤ 5 msec	1 msec

Table 34 lists test results for the OLTP workload tests for the 160,000-user configuration.

Table 34. OLTP Workload Test Results for 160,000-user Configuration

Metric	Microsoft® Success Criterion	Result
Achieved Database Disk IOPS	Varies	24,400 IOPS
Database Avg. Disk Seconds/Read	≤ 20 msec	1.5 msec
Database Avg. Disk Seconds/Write	≤ 20 msec	1 msec
Transaction Log Avg. Disk Seconds/Read	≤ 5 msec	N/A
Transaction Log Avg. Disk Seconds/Write	≤ 5 msec	1 msec

Table 35 lists test results for the OLTP workload tests for the 320,000-user configuration.

Table 35. OLTP Workload Test Results for 320,000-user Configuration

Metric	Microsoft Success Criterion	Result
Achieved Database Disk IOPS	Varies	37,856 IOPS
Database Avg. Disk Seconds/Read	≤ 20 msec	2 msec
Database Avg. Disk Seconds/Write	≤ 20 msec	1 msec
Transaction Log Avg. Disk Seconds/Read	≤ 5 msec	N/A
Transaction Log Avg. Disk Seconds/Write	≤ 5 msec	1 msec

Testing shows that all three configuration profiles for this solution meet or exceed all design goals for these architectures.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems [Global Services](#) website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Education](#) website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the [Hitachi Data Systems](#) website.



Corporate Headquarters

2845 Lafayette Street, Santa Clara, California 95050-2627 USA

www.HDS.com

Regional Contact Information

Americas: +1 408 970 1000 or info@HDS.com

Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@HDS.com

Asia-Pacific: +852 3189 7900 or hds.marketing.apac@HDS.com

© Hitachi Data Systems Corporation 2015. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. "Innovate with Information" is a trademark or registered trademark of Hitachi Data Systems Corporation. Active Directory, Outlook, Hyper-V, Lync, SharePoint, Visual Studio, Windows, Windows Server and SQL Server are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.