TECHNICAL PAPER

Hitachi Unified Compute Platform for Microsoft® Private Cloud for Microsoft Exchange

Reference Architecture Guide

By Leo Nguyen

April 29, 2015

Microsoft Partner
Gold Server Platform

Hitachi Data Systems

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

Hitachi Unified Compute Platform for Microsoft® Private Cloud for Microsoft Exchange

Reference Architecture Guide

This Hitachi Unified Compute Platform for Microsoft® Private Cloud for Microsoft Exchange reference architecture uses the following:

- Hitachi Virtual Storage Platform G600.
- Hitachi Compute Blade 500.
- Microsoft Windows Server® 2012 R2.
- Microsoft Hyper-V®.
- Microsoft Exchange Server 2013.

This Private Cloud Fast Track solution provides you the following:

- **Faster deployment** with the ability to adapt to market pressures by leveraging the scalable design of the architecture.
- **Reduced risk** with a tested and validated solution with end-to-end interoperability of compute, storage and network components.
- **Predictability** with a consistent experience to the hosted workloads with standardized physical servers, network devices and storage systems.

This solution provides rapid provisioning and de-provisioning of virtual machines, with tight integration of the storage architecture. The automation in this solution includes the ability to provision a new LUN to support additional virtual machines and to add that LUN to a host cluster.

Storage architecture, critical for Hyper-V cloud solutions, is a topic that rapidly changes with new standards, protocols and implementations. Hitachi Virtual Storage Platform G600 scales as you host additional workloads in the cloud, providing high availability to these workloads.

Unified Compute Platform for Microsoft Private Cloud supports the shift from server operator to a service provider. There is a set of services to accompany the infrastructure, such as reporting, usage metering, and self-service provisioning.

Without these services, the cloud service layer is unavailable and your investment is little more than a traditional data center. For this reason, Unified Compute Platform for Microsoft Private Cloud provides high availability to the management systems.

This reference architecture is for IT administrators involved in data center planning and design, particularly a Hyper-V private cloud infrastructure. You need some familiarity with the following:

- Hitachi Virtual Storage Platform Gx00.
- Hitachi Storage Navigator.
- Microsoft Windows Server 2012 R2.
- Microsoft Hyper-V failover clustering.
- Microsoft Exchange Server 2013.

Solution Overview

This reference architecture is built on a hardware and software virtualization platform from Hitachi, Brocade and Microsoft.

The Hitachi Compute Blade 500 (CB 500) configuration in this solution has the following:

- Two-node Microsoft Hyper-V failover cluster.
- Infrastructure virtual machines.
- Management virtual machines.
- Exchange virtual machines:
  - 4,000 mailboxes that are 2 GB in size.
  - Two database copies for mailbox resiliency.

The Hitachi Virtual Storage Platform G600 (VSP G600) configuration in this solution has the following:

- One Hitachi midrange enterprise storage.
- 48 × 4 TB SAS disks.

The Brocade network configuration in this solution has the following:

- Two Brocade 6740 10 Gb/sec Ethernet switches.
- Two Brocade 6510 16 Gb/sec Fibre Channel switches.

Figure 1 shows the physical topology of this reference architecture.

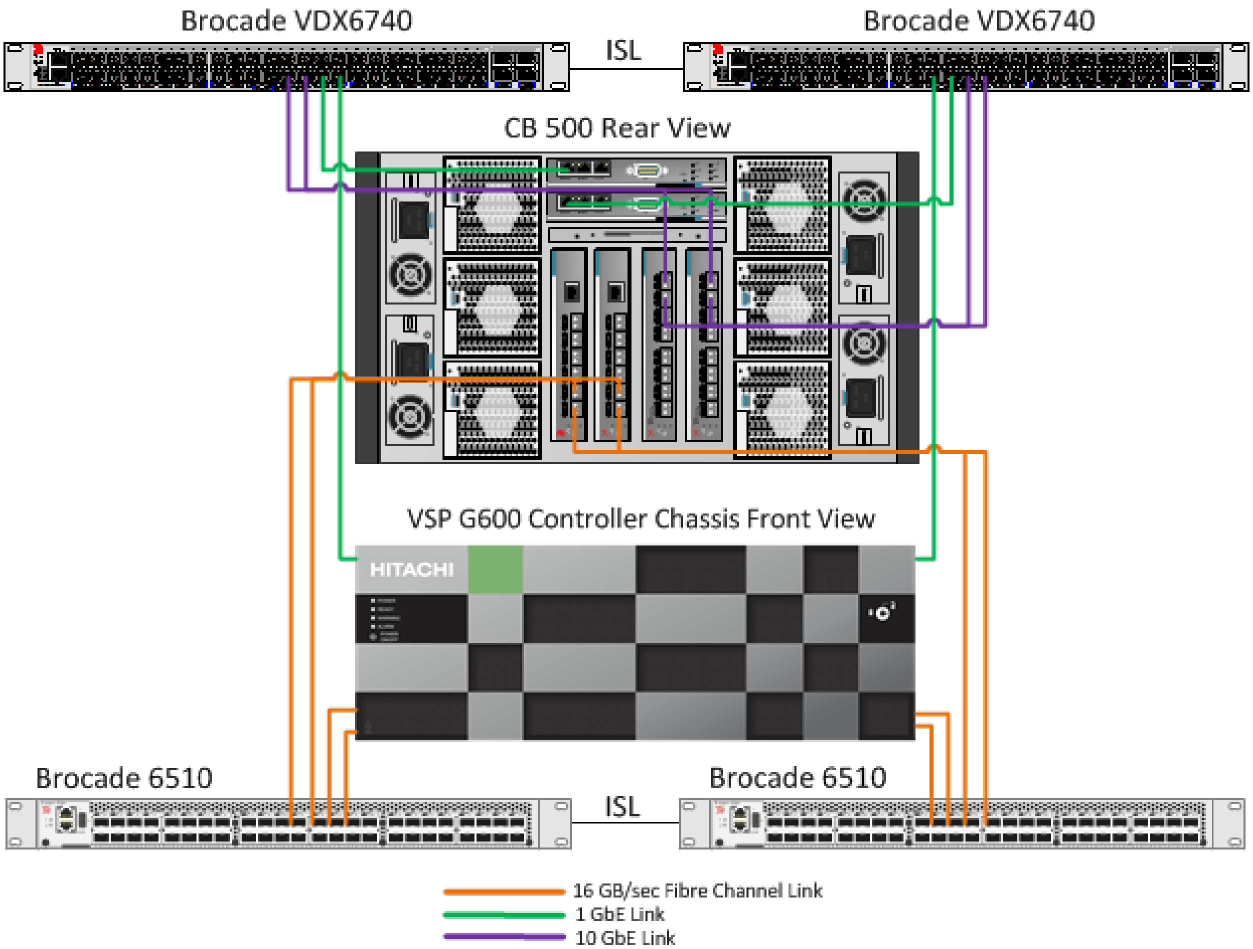

Figure 1

Key Solution Components

These are the key components required to deploy this solution.

Hitachi Virtual Storage Platform Family Systems

The Hitachi Virtual Storage Platform family systems are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform family, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all models in the Virtual Storage Platform family. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

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. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

Hitachi Dynamic Provisioning

On Hitachi storage systems, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a logical size you specify of up to 60 TB created against it without allocating any physical space initially.

Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Dynamic Provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

When used with Hitachi Virtual Storage Platform G600, Hitachi Dynamic Provisioning has the benefit of thin provisioning. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Hitachi Command Suite 8.0

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 is an optional add-on module that integrates into Command Suite to deploy, administer and manage Hitachi servers.

Microsoft Windows Server 2012 R2

Microsoft Windows Server is a multi-purpose server that increases the reliability and flexibility of your server or private cloud infrastructure.

Microsoft Hyper-V is a hypervisor-based virtualization technology that is integrated into Microsoft Windows Server. It allows for the reduction of hardware footprints and capital expenses through server consolidation.

Failover cluster provides high availability and scalability to many server workloads. In a failover cluster, if one or more of the clustered servers (nodes) fails, other nodes begin to provide service.

Microsoft System Center R2 Virtual Machine Manager

Virtual Machine Manager, a part of Microsoft System Center, is a management solution for the virtualized datacenter. It lets you configure and manage your virtualization host, networking and storage resources to create and deploy virtual machines and services to private clouds.

This solution uses SCVMM only to manage the Hyper-V private cloud hosts and guests in a single datacenter.

Microsoft SQL Server 2012 Enterprise Edition

Microsoft SQL Server is a complete set of enterprise-ready technologies and tools to derive value from information. Enjoy high levels of performance, availability and security. Employ more productive management and development tools. Deliver pervasive insight with self-service business intelligence (BI).

Microsoft Exchange Server 2013 Enterprise Edition

Microsoft Exchange Server is a messaging application that offers high availability and site resiliency at the database level.

Solution Design

This section provides detailed information on the Microsoft Private Cloud design architecture for Microsoft Exchange Server 2013 to support 4,000 mailboxes that are 2 GB in size. This architecture uses Hitachi compute and storage resources, with Brocade networking to deploy, operate and manage this solution.

Hitachi Command Suite with Hitachi storage and server integrations, combined with Microsoft Virtual Machine Manager provides the capabilities to manage the private cloud for Exchange.

Configure three Exchange virtual machines on a two-node Microsoft Hyper-V cluster using Cluster Shared Volumes for high availability and database resiliency.

Each Exchange virtual machine houses four databases. There are two active databases on storage pool 1 and two passive database copies on storage pool 2. This gives a total of 12 databases. Each database supports 667 mailboxes that are 2 GB in size.

This design balances the CPU and memory resources to provide better performance. This design architecture can sustain one Hyper-V host or one Exchange virtual machine failure without any end user disruption.

Table 1 lists the hardware components used for the solution.

Table 1. Hardware Components

Hardware	Description	Version	Quantity
Hitachi Compute Blade 500 chassis	■ Up to 8 server blades ■ 2 management modules ■ 6 cooling fan modules ■ 4 power supply modules ■ 2 Brocade 5460 Fibre Channel switch modules ■ 2 Brocade VDX 6746 DCB 10 GbE switch modules	A0241-B-9901 A0138 1019	1
CB520H B3 server blade	■ Half-size blade ■ 2 × 18-core Intel Xeon E5-2699 v3 processors at 2.30 GHz ■ 512 GB RAM ■ 16 × 32 GB DIMMs ■ 2 hot-swappable 2.5 inch SAS drives ■ 1 Emulex 10 GbE onboard CNA ■ 1 Hitachi 16 Gb/sec Fibre Channel mezzanine card	08-22	2
Hitachi Virtual Storage Platform G600	■ Dual controller ■ 32 × 8 Gb/sec Fibre Channel ports ■ 256 GB cache memory	83-01-01-40/00	1
DBX drive box	■ 48 × 4 TB 7.2k RPM SAS drives	n/a	1
Brocade VDX 6740-48	■ 48-port 10 GbE switch	4.1.2	2
Brocade 6510	■ 48-port 16 Gb/sec Fibre Channel switch	7.0.1a	2

Table 2 lists the software components used for the solution.

Table 2. Software Components

Software	Version
Hitachi Dynamic Provisioning	Microcode Dependent
Hitachi Command Suite	8.x
Microsoft Windows Server	2012 R2 Datacenter and Standard edition
Microsoft SQL Server	2012 SP1 Enterprise edition
Microsoft System Center Virtual Machine Manager	2012 R2
Microsoft Exchange Server	2013 SP1 Enterprise edition

Hitachi Compute Blade 500

This section provides detailed configuration information for the server chassis, server blades, and switch modules.

Server Chassis Configuration

The server chassis has redundant management modules to provide high availability access to manage the chassis, server blades, and switch modules. The hot-swappable power and fan modules allow for nondisruptive maintenance.

Server Blade Configuration

The server blade has a two-socket 18-core Intel Xeon E5-2699 v3 at 2.30 GHz processor with 512 GB memory. The on-board card is an Emulex 10 GbE Converged Network Adapter (CNA). The CNA card is configured into four logical NICs on channel 0 and four logical NICs on channel 1. The Hitachi dual-port Fibre Channel 16 Gb/sec mezzanine card is installed on slot 1.

Network Switch Module Configuration

The server chassis comes with two Brocade VDX 6746 10 GbE DCB switch modules installed into the chassis at slot 0/1. The uplink is configured for link aggregation and connected to two Brocade VDX 6720 switches using Brocade Virtual Cluster Switching (VCS) and Inter Switch-Link (ISL) technologies. Each switch has 24 ports with 8 external and 16 internal ports.

For performance enhancement and security, the networks are configured using the following VLANs:

- Management network using VLAN 243 — Used for server chassis, server blades, storage, and management virtual machine traffic.
- Outlook Anywhere network using VLAN 100 — Used for Exchange servers and client traffic.
- Replication network using VLAN 101 — Used for log shipping and database seeding traffic.
- Live Migration network using VLAN 102 — Used for Live Migration traffic.

SAN Switch Module Configuration

The server chassis comes with two Brocade 5460 Fibre Channel 8 Gb/sec switch modules installed into the chassis at slot 2/3. The Brocade 5460 switch has 26 ports with 6 external and 20 internal ports.

For redundancy and performance, multiple I/O paths are configured and assigned to the following storage ports:

- 1A/2A
- 3A/4A

Figure 2 shows the rear and front views of a Hitachi Compute Blade 500 chassis.

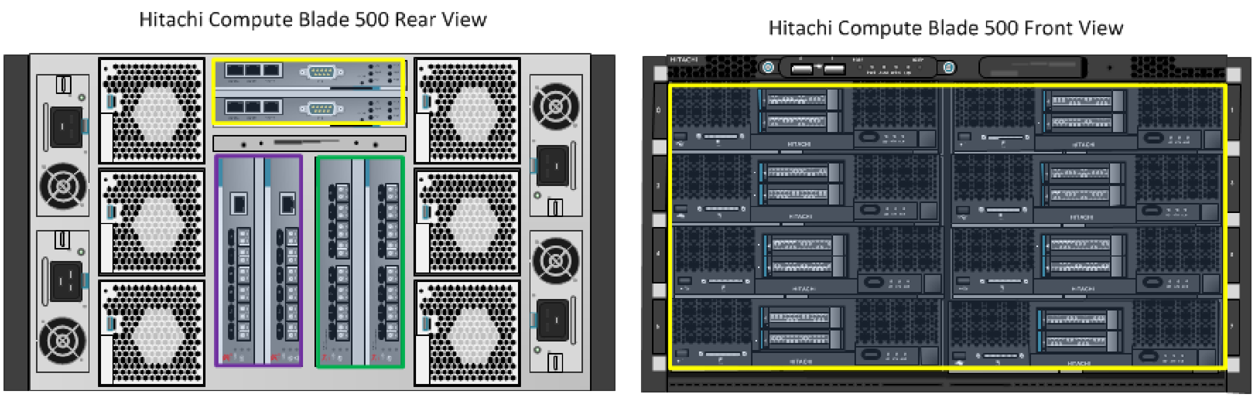

Figure 2

Microsoft Exchange Architecture

Microsoft Exchange Server 2013 has significantly changed its architecture from five roles (Exchange 2007 and 2010) to the following two roles:

- The client access server is a stateless server that handles authentication requests, and then routes the requests to the mailbox server that hosts the active copy of the mailbox database.
- The mailbox server hosts all the components and protocols that process, render and store the data. The hub transport and unified messaging server are now part of the mailbox server.

On the compute side, two times more CPU and memory resources are required to support the same workloads.

On the storage side, the IOPS and capacity requirements to support the mailbox profile are the same as Exchange Server 2010.

Compute and Storage Sizing and Design

This is the compute and storage sizing and design based on the following Exchange profile:

- 4,000 mailboxes.
- 75 KB average message size.
- 2 GB mailbox size.
- 150 messages sent and received per day.
- Two database copies.

Determining Compute CPU and Memory Requirements

The Exchange 2013 Mailbox Server Role Requirements Calculator and Processor Query Tool v1.1, using dual 18-core Intel Xeon E5-2699 v3 processors, 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 likely not justified.

Using 36 cores, the Processor Query Tool recommended 1391 as the spec rate value. Based on the calculations, allocate ten vCPUs with 96 GB to each of the three Microsoft Exchange virtual machines.

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

Determining Storage I/O Requirements

The Exchange 2013 Mailbox Server Role Requirements Calculator was used to determine the storage I/O requirements.

From Table 3, the recommended IOPS per mailbox using 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 3. Estimated IOPS Requirements

Messages sent or received per mailbox per day	*Estimated IOPS per mailbox (Active or Passive)*
50	0.034
100	0.067
150	0.101
200	0.134
250	0.168
300	0.201

Table 4 shows the IOPS requirements per database, server, database availability group, and environment. The calculator recommends 40 × 4TB SAS disks to deliver the Exchange IOPS. Add four additional disks for the Microsoft Windows Server operating System (VHDX) and four hot spares, for a total of 48 disks. This following storage configuration is for high availability and database resiliency.

Storage:

- **4 disks** — Microsoft Windows Server operating system.
- **20 disks** — Active Exchange databases and logs.
- **20 disks** — Passive Exchange databases and logs.
- **4 disks** — Hot spares.

Table 4. Calculated Host I/O and Throughput Requirements

Host I/O and Throughput Requirements	*Per Database*	*Per Server*	*Per DAG*	*Per Environment*
Total database required IOPS	80	322	965	965
Total log required IOPS	18	69	207	207
Database read I/O percentage	60%	N/A	N/A	N/A
Background database maintenance throughput requirements	1 MB/sec	4 MB/sec	12 MB/sec	12 MB/sec

Determining Storage Capacity Requirements

The Exchange 2013 Mailbox Server Role Requirements Calculator was used to determine 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 online 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. Essentially, it is 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 item search. This contributes an additional overhead of about 20 percent to the total database size.

Table 5 shows capacity requirements per database, server, database availability group, and environment.

Table 5. Calculated Disk Capacity Requirements

Disk Space Requirements	Per Database (GB)	Per Server (GB)	Per DAG (GB)	Per Environment (GB)
Transport database space required	N/A	386	1159	1159
Database space required	1497	5987	17961	17961
Log space required	74	294	883	883
Database and log volume spare required	2283	9133	27398	27398

Parity Group and LDEV Configuration

Create 11 parity groups using RAID 10 (2D+2D) from the selected 44 disks. Create 22 LDEVs from the parity groups with a capacity size of 3061.76 GB each. RAID-10 is recommended for performance and reliability for SAS 4 TB 7.2k disks.

Dynamic Pool Configuration

Create three dynamic provisioning pools. Select the correct number of LDEVs to add to the pool capacity as listed in Table 6.

Table 6. Dynamic Provisioning Pool

Dynamic Provisioning Pool	Number of Parity Groups	Number of LDEVs	Number of Drives	Purpose
0	1	2	4	Windows Server operating system (VHDX)
1	5	10	20	Exchange active databases and logs
2	5	10	20	Exchange passive databases and logs

Storage Port Configuration

Configure the storage ports as follows with multipath I/O for redundancy:

- **Security** — Enabled.
- **Type** — Fibre.
- **Fabric** — On.
- **Connection Type** — Point to point.

Table 7 shows the storage ports and blade server HBA configuration for multipath I/O.

Table 7. Multipath I/O

<table>
<tr><th>Storage Number</th><th>Blade Number</th><th>Blade HBA</th><th>Storage Port</th><th>Storage Host Group</th></tr>
<tr><td rowspan="4">1</td><td rowspan="2">0</td><td>Blade0_HBA1_1</td><td>Storage_1A</td><td>Blade0_port1_1_Storage_1A</td></tr>
<tr><td>Blade0_HBA1_2</td><td>Storage_2A</td><td>Blade0_port1_2_Storage_2A</td></tr>
<tr><td rowspan="2">1</td><td>Blade1_HBA1_1</td><td>Storage_3A</td><td>Blade1_port1_1_Storage_3A</td></tr>
<tr><td>Blade1_HBA1_2</td><td>Storage_4A</td><td>Blade1_port1_2_Storage_4A</td></tr>
</table>

Microsoft Hyper-V and Virtual Machine Configuration

This is detailed information on the Microsoft Hyper-V and virtual machine configurations.

Hyper-V Configuration

Use System Center Virtual Machine Manager to configure the Hyper-V host cluster and network. The network adapters are configured with NIC teaming using the default settings for performance and redundancy.

- **Team mode** — Switch Independent.
- **Load balancing mode** — Dynamic.
- **Standby adapter** — None (all adapter active).
- **NIC team1 using network adapter 0/1** — Management traffic.
- **NIC team2 using network adapter 2/3** — Exchange Outlook Anywhere traffic.
- **NIC team3 using network adapter 4/5** — Log shipping and database seeding traffic.
- **NIC team4 using network adapter 6/7** — Live Migration traffic.

Cluster Shared Volume Configuration

Configure three cluster shared volumes (CSV) and present them to both Microsoft Hyper-V hosts. Table 8 shows the following use for each CSV:

- CSV 0 is used for Microsoft Windows Server operating system (VHDX).
- CSV 1 and CSV 2 are used for Microsoft Exchange data.

Table 8. Cluster Shared Volume Configuration

Dynamic Pool Number	Hyper-V Host	CSV Number	CSV Size (TB)	Purpose
0	0	1	3	Windows Server OS (VHDX)
	1			
1	0	2	16	Exchange active databases and logs
	1			
2	0	3	16	Exchange passive databases and logs
	1			

vDisk (VHDX) Configuration

Create the VHDX disks as fixed disk for performance. Use Table 9 to create the VHDX disks and assign them to each virtual machine accordingly.

Table 9. VHDX Configuration

Hyper-V Host	*Virtual Machine Name*	*CSV Number*	*vDisk VHDX (GB)*	*Purpose*
0	HCS + HCSM	1	200	Windows OS
	DC1			
	DC2			
	SCVMM			
	SQL			
	MBXCAS1	1	200	Windows OS
		2	2000	DB1 (active)
				DB2 (active)
		3		DB5 (passive)
				DB6 (passive)
2	MBXCAS2	1	200	Windows OS
		2	2000	DB3 (active)
				DB4 (active)
		3		DB1 (passive)
				DB2 (passive)
	MBXCAS3	1	200	Windows OS
		2	2000	DB5 (active)
				DB6 (active)
		3		DB3 (passive)
				DB4 (passive)

Virtual Machine Configuration

Deploy and configure the following virtual machines for this solution. Use SCVMM to configure the Hyper-V hosts cluster, fabric, and network. Use the SCVMM template to rapidly deploy new virtual machines as follows:

- Two domain controllers.
- One Hitachi Command Suite server.
- One SQL server.
- One System Center Virtual Machine Manager (SCVMM) server.
- Three Exchange servers.

Use Table 10 to configure the vCPU, vRAM, and vNIC for the virtual machines.

In a production environment, it is recommended to use the following:

- If it exists, join an existing Microsoft Active Directory® domain. Otherwise, create a new domain.
- Use a hardware load balancer to properly configure all the client access servers to point to a single virtual IP (VIP) address and a fully qualified domain name (FQDN) to provide redundancy and increase performance.
- Use public certificates, as they provide the best method for securing communications between clients and servers.

Table 10. Virtual Machine Configuration

Hyper-V Host	*Virtual Machine Name*	*vCPU*	*vRAM (GB)*	*vNIC*	*Purpose*
0	HCS+HCSM	2	4	1	Storage and compute management
	DC1	2	4	1	Domain Controller Active Directory
	DC2	2	4	1	Domain Controller Active Directory
	SCVMM	2	24	1	Virtual Machine Manager
	SQL	2	24	1	SQL database for SCVMM
	MBXCAS1	10	96	2	Exchange mailbox and client access
1	MBXCAS2	10	96	2	Exchange mailbox and client access
	MBXCAS3	10	96	2	Exchange mailbox and client access

Microsoft Exchange Database Availability Group Design

This is detailed information on the Microsoft Exchange database availability group (DAG) configuration.

With a two-node Microsoft Hyper-V cluster, configure the hosts as follows:

- **Host 0** — Support the infrastructure and management virtual machines.
- **Host 1**— Support three Microsoft Exchange virtual machines.

In the event of a server blade failure, the virtual machines on that server blade automatically perform a live migrate over to the remaining server blade host to support all virtual machines. Each blade has enough resources to sustain a failover scenario.

Configure three Exchange virtual machines in a database availability group, as follows:

- **Exchange Virtual Machine 1: MBXCAS1**
  - DB1-2 as active databases.
  - DB5-6 as passive database copies.
- **Exchange Virtual Machine 2: MBXCAS2**
  - DB3-4 as active databases.
  - DB1-2 as passive database copies.
- **Exchange Virtual Machine 3: MBXCAS3**
  - DB5-6 as active databases.
  - DB3-4 as passive database copies.

Each database in this DAG configuration supports 667 mailboxes. This design balances the CPU and memory resources to provide better performance.

Figure 3 shows the database distribution for a three-member DAG configuration.

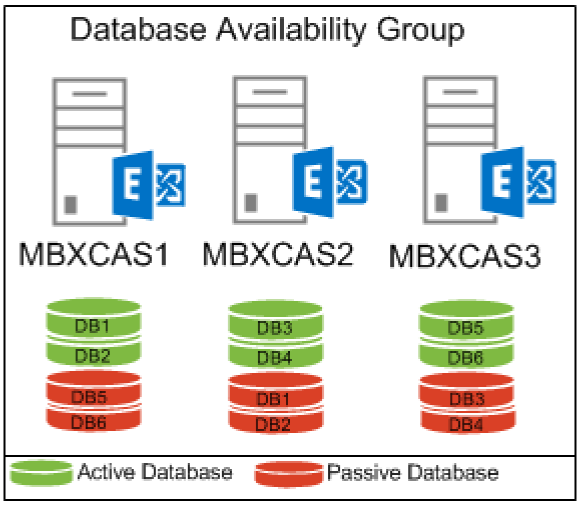

Figure 3

Engineering Validation

This describes the Microsoft tools, test methodology, and test results used to validate this solution.

Test Methodology

This is the test methodology used to validate the storage disk subsystem 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 11 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 m/sec.

Table 11. Jetstress Parameters

Jetstress Parameter	*Value*
Number of Exchange virtual machines	3
Number of databases per host	4
Number of copies per database	2
Thread count	25
Target IOPS per Exchange virtual machine	386
Total IOPS per database availability group	1159
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 nonproduction Exchange environment. These tests send multiple messaging requests from client virtual machines to the Exchange servers to see how Exchange responds to email loads.

These are the three testing objectives:

- The first objective tested the workloads on three Exchange virtual machines using two Hyper-V hosts.
- The second objective tested the workloads on three Exchange virtual machines using one Hyper-V host.
  - Hyper-V host 1 was powered off to simulate virtual machine live migration to host 0.
- The third objective tested the workloads on two Exchange virtual machines using two Hyper-V hosts.
  - One of the Exchange virtual machines was powered off to simulate a database failover to the Exchange virtual machine housing the passive databases. This activated the passive database copy to handle all 2,668 users.

Table 12 lists the Load Generator parameters used in configuring the tests.

Table 12. Load Generator Parameters

Load Generator Parameter	*Value*
Total number of mailboxes	4,000
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

These are the Jetstress and Loadgen test results.

Jetstress Disk Subsystem Throughput Test Results

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

Table 13 and Table 14 show the Jetstress results.

Table 13. Database Sizing and Throughput Jetstress Results

Database Sizing and Throughput	MBXCAS1	MBXCAS2	MBXCAS3
Achieved Transactional I/O per Second	1826	1935	1569
Capacity Percentage	100%	100%	100%
Throughput Percentage	100%	100%	100%

Table 14. Total I/O Performance Jetstress Results

Virtual Machine Name	Exchange DB Instances	I/O DB Reads Average Latency (m/sec)	I/O DB Writes Average Latency (m/sec)	I/O DB Reads/ sec	I/O DB Writes/ sec	I/O DB Reads Average Bytes	I/O DB Writes Average Bytes
MBXCAS1	Instance1624.1	9.685	4.749	313.588	141.959	33350.044	34480.912
	Instance1624.2	9.365	3.942	312.807	142.328	33340.934	34501.431
	Instance1624.3	11.204	5.564	316.134	142.438	33347.472	34492.040
	Instance1624.4	10.715	4.570	314.714	142.402	33350.172	34482.004
MBXCAS2	Instance2484.1	9.281	4.829	331.235	150.828	33335.040	34483.235
	Instance2484.2	9.308	4.141	330.470	150.476	33338.010	34487.224
	Instance2484.3	10.472	5.677	334.160	151.391	33339.188	34487.593
	Instance2484.4	10.474	4.807	334.813	151.766	33336.871	34482.055
MBXCAS3	Instance1104.1	10.516	5.831	270.463	120.432	33371.042	34685.958
	Instance1104.2	10.947	4.588	271.235	120.655	33372.783	34667.514
	Instance1104.3	11.965	6.716	271.808	121.194	33376.181	34667.938
	Instance1104.4	12.583	5.324	272.452	121.090	33384.021	34663.650

Load Generator Test Results for Three Exchange Virtual Machines on Two Hyper-V Hosts

The Load Generator normal test results show three Exchange virtual machines configured with 4,000 mailboxes on two Hyper-V hosts.

Table 15, Table 16, and Table 17 show the Load Generator test results.

Table 15. Memory Percent Committed and CPU Utilization Hyper-V Hosts

Host Name	Memory % Committed	CPU % Utilization
HYV0	39%	26%
HYV1	22%	24%

Table 16. Memory Percent Committed and CPU Utilization Load Generator Results

VM Name	Number of Mailboxes	Memory % Committed	CPU % Utilization	Test Result
MBXCAS1	1,334	41%	50%	Passed
MBXCAS2	1,334	27%	25%	Passed
MBXCAS3	1,334	27%	24%	Passed

Table 17. Simulation Statistics Load Generator Results

Type	Client Name	Tasks Completed	Task Dispatched	Total Task Completed
Remote	lg2	113846	113846	893456
	lg3	114074	114074	
	lg5	114076	114076	
	lg6	96553	96553	
	lg9	114080	114080	
	lg7	96554	96554	
	lg8	95471	95471	
	lg10	93558	93558	

Load Generator Test Results for Three Exchange Virtual Machines on One Hyper-V Host

The Load Generator normal test results show three Exchange virtual machines configured with 4,000 mailboxes on one Hyper-V host.

Table 18, Table 19, and Table 20 show the Load Generator test results.

Table 18. Memory Percent Committed and CPU Utilization Hyper-V Hosts

Host Name	Memory % Committed	CPU % Utilization
HYV0	71%	36%

Table 19. Memory Percent Committed and CPU Utilization Load Generator Results

VM Name	Number of Mailboxes	Memory % Committed	CPU % Utilization	Test Result
MBXCAS1	1,334	33%	35%	Passed
MBXCAS2	1,334	25%	25%	Passed
MBXCAS3	1,334	39%	47%	Passed

Table 20. Simulation Statistics Load Generator Results

Type	Client Name	Tasks Completed	Tasks Dispatched	Total Tasks Completed
Remote	lg5	114073	114073	909400
	lg2	113843	113843	
	lg3	114071	114071	
	lg6	114073	114073	
	lg7	114075	114075	
	lg9	114076	114076	
	lg8	114076	114076	
	lg10	111113	111113	

Load Generator Test Results for Two Virtual Machines on Two Hyper-V Hosts

The Load Generator failover test results show two Exchange virtual machines configured with 4,000 mailboxes on two Hyper-V hosts.

Table 21, Table 22, and Table 23 show the Load Generator test results.

Table 21. Memory Percent Committed and CPU Utilization Hyper-V Hosts

Host Name	Memory % Committed	CPU % Utilization
HYV0	20%	18%
HYV1	22%	37%

Table 22. Memory Percent Committed and CPU Utilization Load Generator Results

VM Name	Number of Mailboxes	Memory % Committed	CPU % Utilization	Test Result
MBXCAS1	2,668	41%	54%	Passed
MBXCAS3	1,334	35%	40%	Passed

Table 23. Simulation Statistics Load Generator Results

Type	Client Name	Tasks Completed	TasksDispatched	Total Tasks Completed
Remote	lg5	79750	79750	864009
	lg3	79830	79830	
	lg2	80515	80515	
	lg6	107960	107960	
	lg7	79383	79383	
	lg9	76859	76859	
	lg8	106974	106974	
	lg10	75523	75523	
	lg5	79750	79750	

Conclusion

This reference architecture guide describes a Microsoft Private Cloud for Microsoft Exchange Server 2013 on Hitachi compute and storage resources and Brocade networking. The solution simplifies deployment, reduces risk, and provides high availability, database resiliency, and flexible scalability.

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.

Hitachi Data Systems Corporation

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. Microsoft, Active Directory, Hyper-V, Outlook, SQL Server, and Windows 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.

AS-380-01, April 2015