TECHNICAL PAPER

HITACHI
Inspire the Next

Deploy 4,000 User Mailboxes with Microsoft Exchange 2013 on Microsoft Private Cloud using Hitachi Virtual Storage Platform G1000

Reference Architecture Guide

By Leo Nguyen

April 11, 2014

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

Deploy 4,000 User Mailboxes with Microsoft Exchange 2013 on Microsoft Private Cloud using Hitachi Virtual Storage Platform G1000

Reference Architecture Guide

The Microsoft® Private Cloud Fast Track v4 for Microsoft Exchange Server 2013 is a joint effort between Microsoft and Hitachi Data Systems to deliver private cloud solutions. Private Cloud Fast Track focuses on the new technologies and services in Microsoft System Center 2012 R2.

The key components of this solution are:

- Microsoft System Center 2012 R2 Virtual Machine Manager to deploy, configure and manage the fabric architectures (compute, network, storage, and virtualization layers).
- Hitachi Command Suite to provision and manage all Hitachi storage environments and data types.
- Hitachi Compute Systems Manager to integrate with Hitachi Command Suite to manage and monitor Hitachi compute, physical, and virtualized operating systems.

This reference architecture guide is intended for you if you are an administrator looking to deploy Microsoft Exchange 2013 in your environment. You need some familiarity with the following to benefit from this document:

- Hitachi Virtual Storage Platform G1000
- Hitachi Command Suite version 7 or later
- Brocade networking
- Microsoft Windows Server® 2012 R2
- Microsoft Hyper-V®
- Microsoft SQL Server®
- Microsoft Exchange Server 2013

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 is an overview of the Microsoft Private Cloud design architecture for Exchange Server 2013 to support 4,000 mailboxes of 2 GB in size using Hitachi compute and storage resources, and Brocade networking.

- Configure the Exchange virtual machines on a four-node Microsoft Hyper-V cluster using cluster shared volumes for high availability.
- Configure two database copies on two Hitachi Virtual Storage Platform G1000 storage systems for database resiliency.

Figure 1 on page 4 shows the physical topology of this reference architecture.

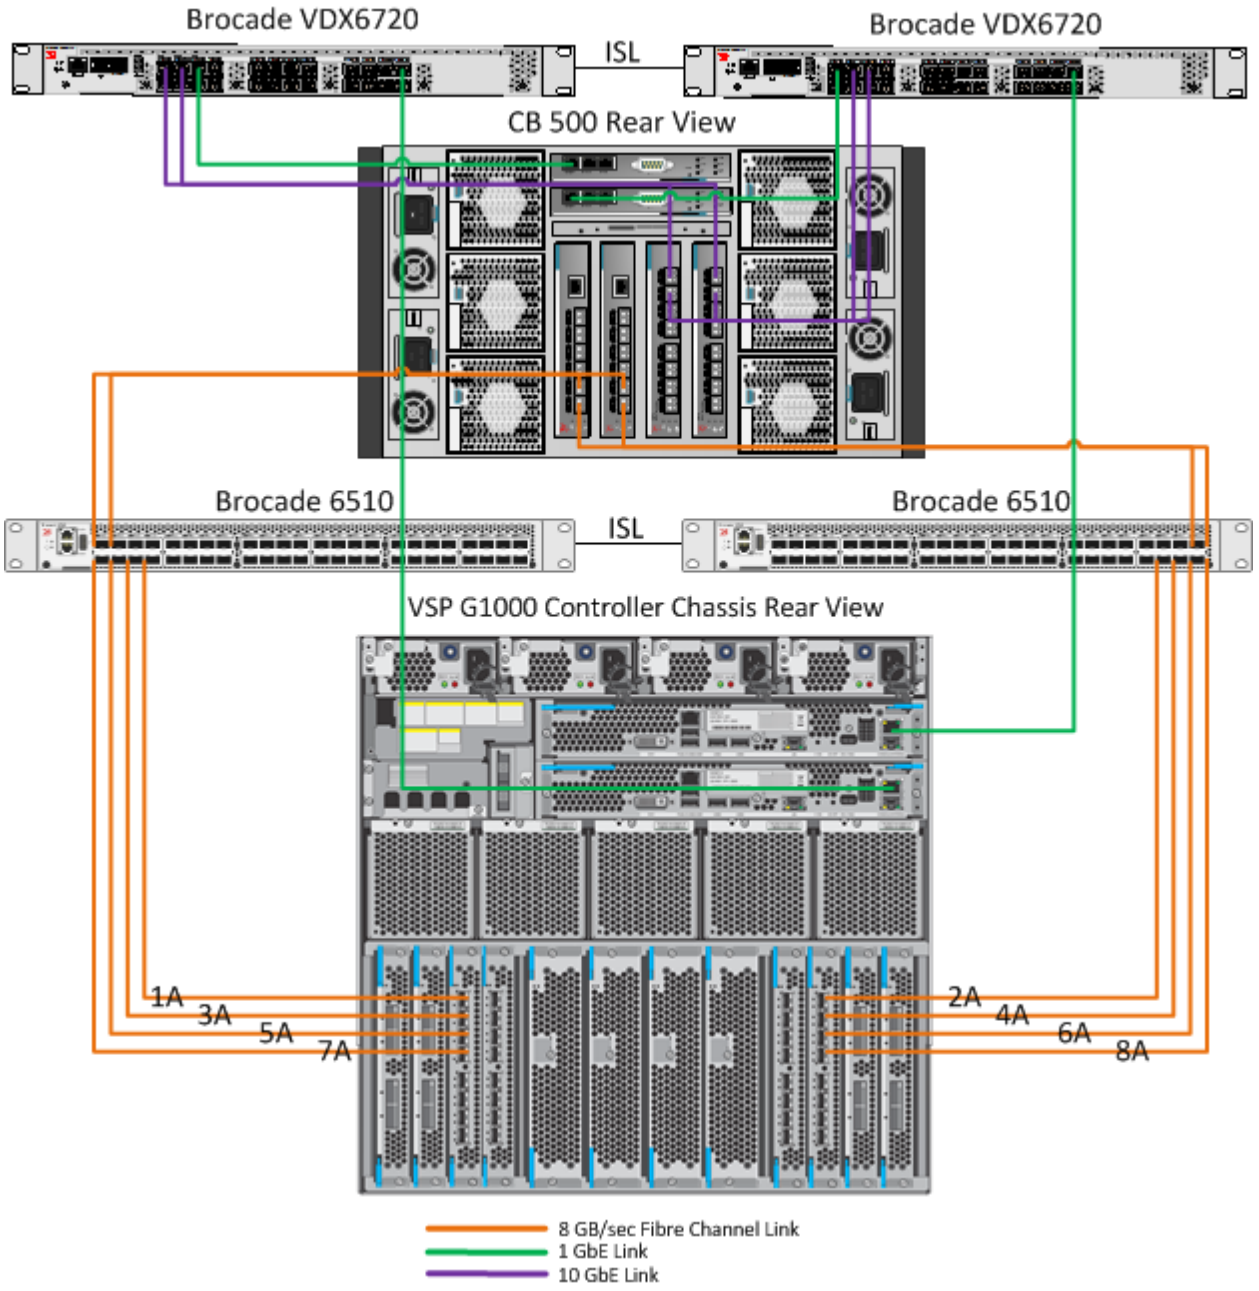
Brocade VDX6720
ISL
Brocade VDX6720
CB 500 Rear View
Brocade 6510
ISL
Brocade 6510
VSP G1000 Controller Chassis Rear View
1A
3A
5A
7A
2A
4A
6A
8A
8 GB/sec Fibre Channel Link
1 GbE Link
10 GbE Link

Figure 1

Key Solution Components

These are the key components required to deploy this solution.

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.

Virtual Storage Platform G1000 Architecture

Hitachi Virtual Storage Platform G1000 is a high performance and large capacity storage system. It has improved Hi-Star Net Architecture and an 8-core microprocessor. The storage consists of the following:

- **Controller Chassis**
  - **Channel Adapter** — This controls data transfer between the upper host and the cache memory.
  - **Disk Adapter** — This controls data transfer between the drives and cache memory.
  - **Cache Path Control Adapter** — Using PCI-Express path, this connects between the processor blades, channel adapter, disk adapter, and the cache backup module kit. It distributes data and sends hot-line signals to the processor blades.
  - **Cache Flash Memory** — This is memory to back up cache memory data when a power failure occurs.
  - **Cache Backup Module Kit** — This is a kit to back up cache memory data when a power failure occurs.

- **Processor Blades —** This consists of the DIMMs and the processor with the chip set. It controls the following using Ethernet:
  - Channel adapter
  - Disk adapter
  - PCI-Express interface
  - Local memory
  - Communication between the service processors
- **Service Processor** — This sets and modifies the storage system configuration, a device availability statistical information acquisition, and maintenance.

- **Drive Chassis** — This is an installable drive unit that connects into the controller chassis.

Virtual Storage Platform G1000 offers these features:

- **Scalability**
  - **Number of controller chassis** — 1 to 2
  - **Number of racks** — 1 to 6
  - **Number of installed channel options** — 1 to 12 sets
  - **Capacity of cache memory** — 32 GB to 2,048 GB
  - **Number of drives** — Up to the following:
    - **2.5-inch HDD** — 2,304
    - **3.5-inch HDD** — 1,152
    - **2.5-inch SSD** (flash drives) — 384
    - **FMD** (flash module drive) — 192

- **High performance**
  - Supports three kinds of high-speed disk drives at the following speeds:
    - 15k RPM
    - 10k RPM
    - 7.2k RPM
  - Supports flash drives and flash module drive with ultra-high speed response
  - Transfers high speed data between the disk adapter and drives at a rate of 6 Gb/sec with the SAS interface
  - Uses the 8-core processor on the processor blade board, doubling the processing ability
- **Large capacity**
  - Supports hard disk drives with the following capacities:
    - 300 GB
    - 600 GB
    - 900 GB
    - 1.2 TB
    - 3 TB
    - 4 TB
  - Supports flash drives with the following capacities:
    - 400 GB
    - 800 GB
  - Supports flash module drives with the following capacities:
    - 1.6 TB
    - 3.2 TB
  - Controls up to 65,280 logical volumes and up to 2,304 disk drives, providing a physical disk capacity of approximately 4,511 TB per storage system

- **Flash module drive**
  - Has a 6 Gb/sec SAS interface, the same as that the hard disk drives and solid state drives
  - Uses MLC-NAND flash memory, featuring high performance, long service life, and cost performance
- **Connectivity —** Supports the following configurations:
  - RAID-6 (6D+2P)
  - RAID-6 (14D+2P)
  - RAID-5 (3D+1P)
  - RAID-5 (7D+1P)
  - RAID-10 (2D+2D)
  - RAID-10 (4D+4D)
- **Non-disruptive service and upgrade**
  - Add, remove, and replace main components without shutting down a device while the storage system is in operation
  - Monitor the running condition of the storage system with a service processor mounted on the drive chassis
  - Enable remote maintenance by connecting the service processor with a service center
  - Upgrade the microcode without shutting down the storage system

Figure 2 on page 9 shows the controller chassis, drive chassis and its sub components for Hitachi Virtual Storage Platform G1000.

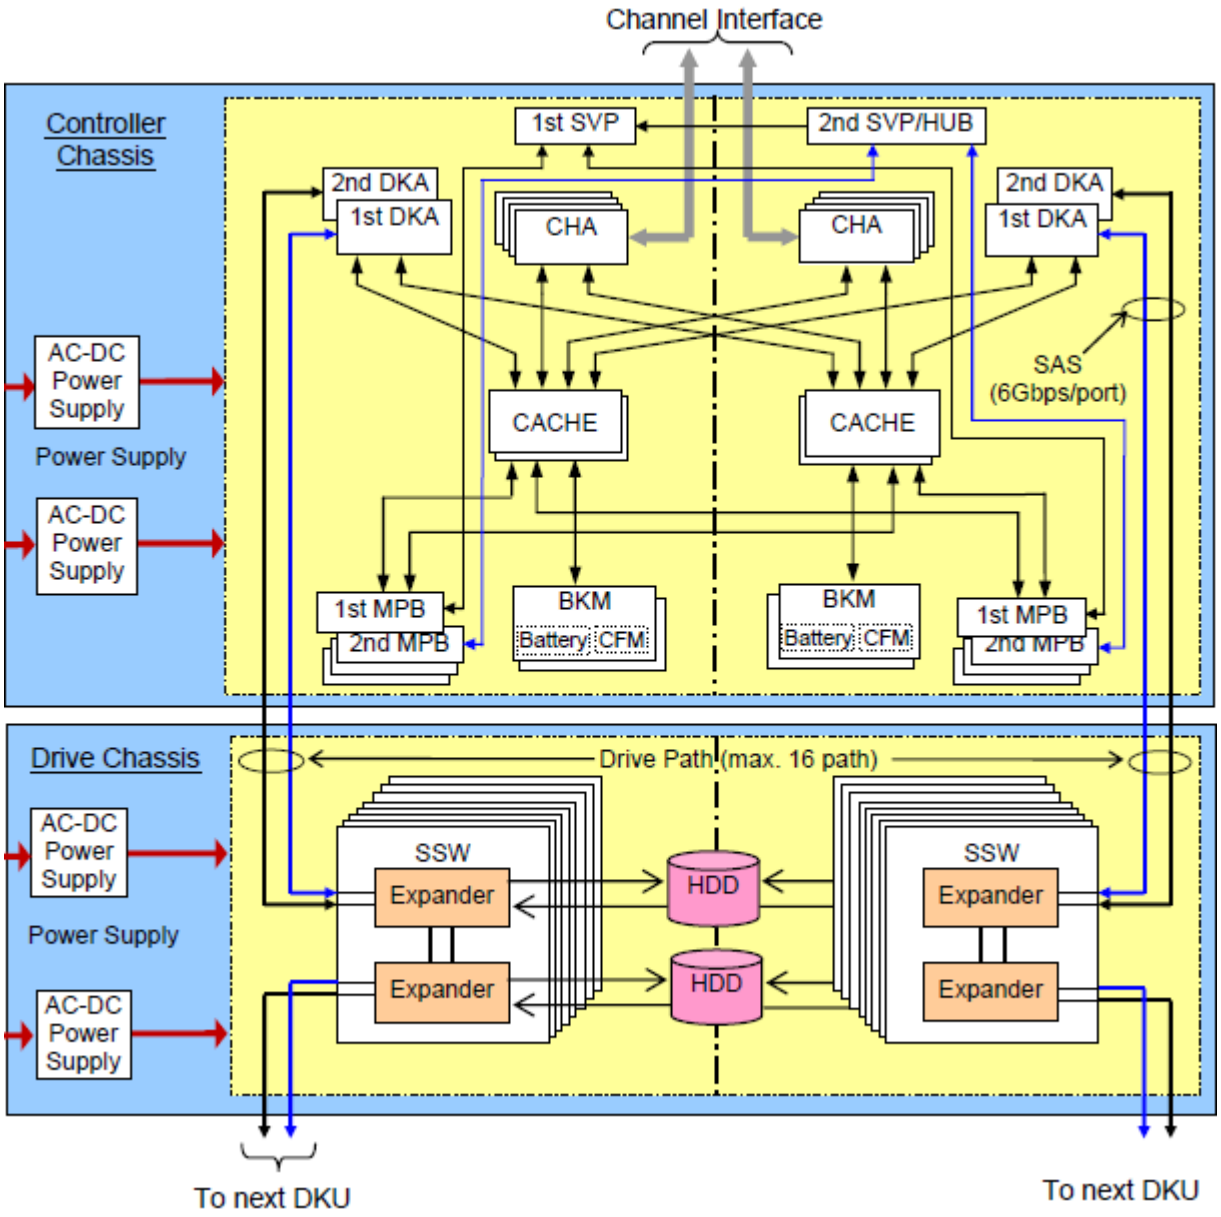

Figure 2

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 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 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 G1000, 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.

Microsoft Windows Server

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, 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

Microsoft SQL Server

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

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

- Database availability groups
- Multiple database copies supported
- Exchange Extensible Storage Engine to support slower and high capacity disks

Solution Design

This is 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 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 the Exchange virtual machines on a four-node Microsoft Hyper-V cluster using Cluster Shared Volumes for high availability. Configure four Exchange virtual machines in a database availability group for database resiliency.

Each Exchange virtual machine houses four databases. There are two active databases on storage 1 and two passive database copies on storage 2. This gives a total of 16 databases. Each database supports 500 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	A0170/00	1

Table 1. Hardware Components (Continued)

Hardware	Description	Version	Quantity
520H B2 server blade	■ Half-size blade ■ 2 × 12-core Intel Xeon E5-2697 v2 processor at 2.70 GHz ■ 192 GB RAM ■ 12 × 16 GB DIMMs ■ 2 hot-swappable 2.5 inch SAS drives ■ 1 Emulex 10 GbE onboard CNA ■ 1 Hitachi 8 Gb/sec Fibre Channel mezzanine card	04-15/10-31	4
Hitachi Virtual Storage Platform G1000	■ Dual controller ■ 32 × 8 Gb/sec Fibre Channel ports ■ 982 GB cache memory	80-01-01-00	2
DBX disk box	■ 64 × 3 TB 7.2k RPM SAS drives	n/a	2
Brocade VDX 6720-24	■ 24-port 10 GbE switch	2.0.1b	2
Brocade 6510	■ 48-port 8-16 Gb/sec Fibre Channel switch	7.0.1.a	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	8.0
Hitachi Compute Systems Manager	8.0
Microsoft Windows Server	2012 R2 Datacenter edition
Microsoft SQL Server	2012 SP1 Enterprise edition
Microsoft System Center, Virtual Machine Manager	2012 R2
Microsoft Windows Server	2012 Datacenter edition
Microsoft Exchange Server	2013 CU3 Enterprise edition

Hitachi Compute Blade 500

This is detailed information on the server chassis, server blades, and switch modules configuration.

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 non-disruptive maintenance.

Server Blade Configuration

The server blade has an on-board 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 Fibre Channel 8 Gb/sec mezzanine card is installed on the server blade at mezzanine slot 1. The card is configured for multiple paths from server to storage ports.

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
- 5A/6A
- 7A/8A

Microsoft Exchange Architecture

Microsoft Exchange Server 2013 has significantly changed its architecture from five roles (Exchange 2007 and 2010) to 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 is 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 two 12-core Intel Xeon E5-2697 v2 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 24 cores, the Processor Query Tool recommended 957 as the spec rate value. Based on the calculator results, allocate seven vCPUs with 64 GB to each of the four Microsoft Exchange virtual machines

Note — Microsoft recommends additional 10% CPU overhead for Exchange virtual machine 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 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 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 recommended 48 × 3TB SAS disks to deliver the Exchange IOPS. Add 16 additional disks for the Microsoft Windows Server operating System (VHDX) and for hot spares, for a total of 64 disks. This following storage configuration is for high availability and database resiliency.

- **Storage 1:**
  - **4 disks** — Microsoft Windows Server operating system
  - **24 disks** — Active Exchange databases and logs
  - **4 disks** — Hot spares
- **Storage 2:**
  - **24 disks** — Active 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	60	241	965	965
Total log required IOPS	13	52	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	16 MB/sec	16 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 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	114	458	458
Database space required	1123	4490	17961	17961
Log space required	55	221	883	883
Database and log volume spare required	1712	6849	27398	27398

Parity Group and LDEV Configuration

The RAID level on the parity groups are configured as RAID-10 (2D+2D) using 3 TB SAS 7.2k RPM drives. RAID-10 is recommended for performance and reliability. Two LDEVs (2734.08 GB) are created from each parity group.

Table 6 shows the parity group and LDEV configuration.

Table 6. Parity Group and LDEV

Storage Number	*Parity Group ID*	*LDEV ID*	*LDEV Size (GB)*
1	1-1	00:00:00	2734.08
		00:00:01	
	1-2	00:00:02	
		00:00:03	
	1-3	00:00:04	
		00:00:05	
	1-4	00:00:06	
		00:00:07	
	1-5	00:00:08	
		00:00:09	
	1-6	00:00:10	
		00:00:11	
	1-7	00:00:12	
		00:00:13	

Table 6. Parity Group and LDEV (Continued)

Storage Number	Parity Group ID	LDEV ID	LDEV Size (GB)
2	1-2	00:00:02	2734.08
		00:00:03	
	1-3	00:00:04	
		00:00:05	
	1-4	00:00:06	
		00:00:07	
	1-5	00:00:08	
		00:00:09	
	1-6	00:00:10	
		00:00:11	
	1-7	00:00:12	
		00:00:13	

Dynamic Pool Configuration

Create three dynamic provision pools using the following LDEV ID listed in Table 7. These pools are for the Microsoft Windows Server operating system and Exchange data.

Table 7. Dynamic Provisioning Pool

Storage Number	Dynamic Provisioning Pool	Number of Parity Groups	Number of Drives	LDEV ID	Purpose
1	0	1	4	00:00:00 to 00:00:01	Windows Server operating system (VHDX)
1	1	6	24	00:00:02 to 00:00:13	Exchange active databases and logs
2	1	6	24	00:00:02 to 00:00:13	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 8 shows the storage ports and blade server HBA configuration.

Table 8. Multipath I/O

Storage Number	*Blade Number*	*Blade HBA*	*Storage Port*	*Storage Host Group*
1	0	Blade0_HBA1_1	Storage1_1A	Blade0_port1_1_Storage1_1A
		Blade0_HBA1_2	Storage1_2A	Blade0_port1_2_Storage1_2A
	1	Blade1_HBA1_1	Storage1_3A	Blade1_port1_1_Storage1_3A
		Blade1_HBA1_2	Storage1_4A	Blade1_port1_2_Storage1_4A
	2	Blade2_HBA1_1	Storage1_5A	Blade2_port1_1_Storage1_5A
		Blade2_HBA1_2	Storage1_6A	Blade2_port1_2_Storage1_6A
	3	Blade3_HBA1_1	Storage1_7A	Blade3_port1_1_Storage1_7A
		Blade3_HBA1_2	Storage1_8A	Blade3_port1_2_Storage1_8A
2	0	Blade0_HBA1_1	Storage2_1A	Blade0_port1_1_Storage2_1A
		Blade0_HBA1_2	Storage2_2A	Blade0_port1_2_Storage2_2A
	1	Blade1_HBA1_1	Storage2_3A	Blade1_port1_1_Storage2_3A
		Blade1_HBA1_2	Storage2_4A	Blade1_port1_2_Storage2_4A
	2	Blade2_HBA1_1	Storage2_5A	Blade2_port1_1_Storage2_5A
		Blade2_HBA1_2	Storage2_6A	Blade2_port1_2_Storage2_6A
	3	Blade3_HBA1_1	Storage2_7A	Blade3_port1_1_Storage2_7A
		Blade3_HBA1_2	Storage2_8A	Blade3_port1_2_Storage2_8A

Microsoft Hyper-V and Virtual Machine Configuration

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

Hyper-V Configuration

Install Microsoft Windows Server 2012 R2 on each blade server local drives using a RAID 1 configuration. The Microsoft Hyper-V role is added.

Table 8 on page 20 shows the storage cluster shared volumes for the hosts.

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 four cluster shared volumes (CSV) and present them to the Microsoft Hyper-V hosts. Table 9 shows the following use for each CSV:

- CSV 1 and CSV 3 are used for Microsoft Windows Server operating system (VHDX)
- CSV 2 and CSV 4 are used for Microsoft Exchange data

Table 9. Cluster Shared Volume Configuration

Storage Number	Dynamic Pool No.	Hyper-V Host	CSV Number	CSV Size (TB)	Purpose
1	0	0	1	3	Windows Server OS (VHDX)
		1			
		2			
		3			
	1	0	2	18	Exchange data (active)
		1			
		2			
		3			
2	1	0	3	18	Exchange data (passive copy)
		1			
		2			
		3			

vDisk (VHDX) Configuration

Create the VHDX disks as fixed disk for performance. Use Table 10 to create the VHDX and assign to each virtual machine accordingly.

Table 10. VHDX Configuration

Hyper-V Host	Virtual Machine Name	CSV Number	vDisk VHDX (GB)	Purpose
0	HCS	1	200	Microsoft Windows Server operating system
	HCSM	1	200	
	DC1	1	200	
	DC2	1	200	
	SCVMM	1	200	
	SQL	1	200	

Table 10. VHDX Configuration (Continued)

Hyper-V Host	Virtual Machine Name	CSV Number	vDisk VHDX (GB)	Purpose
1	MBXCAS1	1	200	Microsoft Windows operating system
		2	1800	DB1 (active
		2	1800	DB2 (active
		3	1800	DB3 (passive)
		3	1800	DB4 (passive)
	MBXCAS2	1	200	Windows OS
		2	1800	DB3 (active
		2	1800	DB4 (active
		3	1800	DB5 (passive)
		3	1800	DB6 (passive)
2	MBXCAS3	1	200	Windows OS
		2	1800	DB5 (active
		2	1800	DB6 (active
		3	1800	DB7 (passive)
		3	1800	DB8 (passive)
3	MBXCAS4	1	200	Windows OS
		2	1800	DB7 (active
		2	1800	DB8 (active
		3	1800	DB1 (passive)
		3	1800	DB2 (passive)

Virtual Machine Configuration

To deploy this solution, you need to deploy the following:

- Two infrastructure virtual machines
- Four management virtual machines
- Four Microsoft Exchange virtual machines

Use Table 11 on page 24 to configure the vCPU, vRAM, and vNIC for the virtual machines.

Configure the Exchange virtual machines to use the following:

- DNS round robin for the client access role
- Self-signed certificates for the mailbox role

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

- If it exists, join to an existing Microsoft Active Directory® domain. Otherwise create a new domain.
- Use 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 11. Virtual Machine Configuration

Hyper-V Host	*Virtual Machine Name*	*vCPU*	*vRAM (GB)*	*vNIC*	*Purpose*
0	HCS	2	4	1	Storage management
	HCSM	2	4	1	Compute management
	DC1	2	4	1	Domain Controller Active Directory
	DC2	2	4	1	Domain Controller Active Directory
	SCVMM	2	32	1	Virtual Machine Manager
	SQL	4	32	1	SQL database for SCVMM
1	MBXCAS1	7	64	2	Exchange mailbox and client access
	MBXCAS2	7	64	2	Exchange mailbox and client access
2	MBXCAS3	7	64	2	Exchange mailbox and client access
3	MBXCAS4	7	64	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 four-node Microsoft Hyper-V cluster, configure the hosts as follows:

- **Host 1** — Support the infrastructure and management virtual machines
- **Host 2, Host 3, and Host 4** — Support four Microsoft Exchange virtual machines

In the event of a server blade failure, the virtual machines on that server blade automatically live migrate over to the remaining server blade hosts to support all virtual machines. Configure each server blade with two 12-core Intel E5-2697 2.70GHz v2 processors and 320 GB memory. This is enough resources to sustain a failover scenario.

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

- **Exchange Virtual Machine 1: MBXCAS1**
  - DB1-2 as active databases
  - DB3-4 as passive database copies
- **Exchange Virtual Machine 2: MBXCAS2**
  - DB3-4 as active databases
  - DB5-6 as passive database copies
- **Exchange Virtual Machine 3: MBXCAS3**
  - DB5-6 as active databases
  - DB7-8 as passive database copies
- **Exchange Virtual Machine 4: MBXCAS4**
  - DB7-8 as active databases
  - DB1-2 as passive database copies

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

Figure 3 on page 26 shows the database distribution for a four-member DAG configuration.

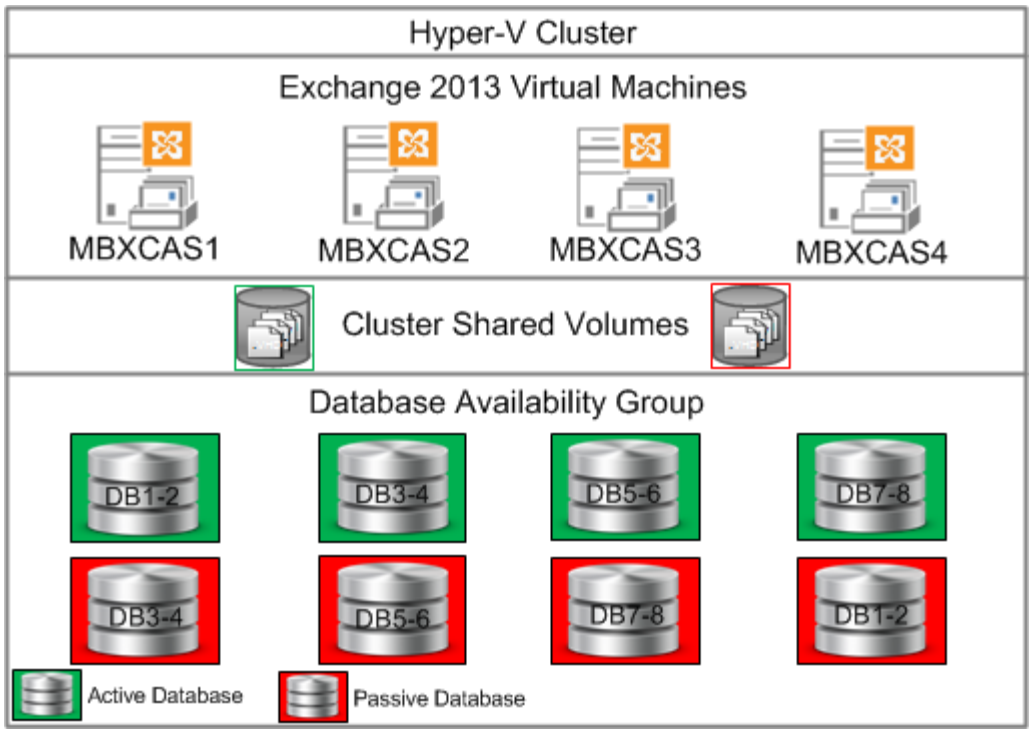
Hyper-V Cluster
Exchange 2013 Virtual Machines
MBXCAS1
MBXCAS2
MBXCAS3
MBXCAS4
Cluster Shared Volumes
Database Availability Group
DB1-2
DB3-4
DB5-6
DB7-8
DB3-4
DB5-6
DB7-8
DB1-2
Active Database
Passive Database

Figure 3

Engineering Validation

This describes the Microsoft tools, test methodology, and test results used to validate this Microsoft Private Cloud for Exchange 2013 solution.

Test Methodology

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

Jetstress Disk Subsystem and Mailbox Profile Tests

The purpose of using Microsoft Exchange Server Jetstress 2013 is to validate the storage disk subsystem performance to support 4,000 mailboxes of 2 GB in size with 150 messages sent and received per day and two database copies. This solution requires 965 IOPS for support this solution. Each of the four Microsoft Exchange virtual machines should deliver 241 IOPS or better.

This tool works with the Exchange database engine to simulate the Exchange database and log disk I/O load. Microsoft recommends verifying the performance and stability of a disk subsystem prior to putting an Exchange server into production.

Table 12 shows the Jetstress parameters configuration with 80% storage capacity. This is the Microsoft recommended test method for a worst-case scenario. The testing goal was to verify that the achieved IOPS meet or exceed the target IOPS with latency less than 20 msec.

Table 12. Jetstress Parameters

Jetstress Parameter	*Value*
Number of Exchange virtual machine	4
Number of mailboxes per host (worst-case scenario)	4,000
Mailbox size	2 GB
Number of users per database	500
IOPS per mailbox (including 20% overhead)	0.1212
Number of databases per host	4
Number of copy per database	2
Thread count	20
Target IOPS per Exchange virtual machine	241
Total IOPS per database availability group	965
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. 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 two testing objectives:

- The first objective tested the workloads on four Exchange virtual machines using 14 client virtual machines.
- The second objective tested the workloads on three Exchange virtual machines using 14 client virtual machines.
  - Exchange virtual machine 2 was powered off to simulate a database failover to Exchange virtual machine 1. This activated the passive database copy to handle all 2,000 users.

Table 13 lists the Load Generator parameters used in configuring the tests.

Table 13. Load Generator Parameters

Load Generator Parameter	*Value*
Total number of mailboxes	4,000
Number of active mailboxes per host	1,000
Number of passive mailboxes per host	1,000
Number of users per database	500
Client type	Microsoft Outlook® 2007 Online
Action profile	Microsoft Outlook 150
Mailbox size	2 GB
Length of simulation day	24 hours
Total length of simulation	2 hours

Test Results

These are the Jetstess 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 msec to support this reference architecture.

Table 14 and Table 15 show the Jetstress results.

Table 14. Disk Subsystem Jetstress Results

Database Sizing and Throughput	*MBXCAS1*	*MBXCAS2*	*MBXCAS3*	*MBXCAS4*
Achieved Transactional I/O per Second	1043.945	1043.434	906.293	907.741
Capacity Percentage	80%	80%	80%	80%
Throughput Percentage	100%	100%	100%	100%
Initial Database Size (bytes)	3586566389760	3586675441664	3584444071936	3584544735232
Final Database Size (bytes)	3597287030784	3597438025728	3593755426816	3593872867328
Database Files (Count)	4	4	4	4

Table 15. Disk Subsystem Transactional I/O Performance Jetstress Results

Virtual Machine Name	*Exchange DB Instances*	*I/O DB Reads Average Latency (msec)*	*I/O DB Writes Average Latency (msec)*	*I/O DB Reads/ sec*	*I/O DB Writes/ sec*	*I/O DB Reads Average Bytes*	*I/O DB Writes Average Bytes*
MBXCAS1	Instance1944.1	15.686	2.225	173.848	87.275	32990.962	35117.497
	Instance1944.2	15.689	1.917	173.717	86.932	32979.659	35136.610
	Instance1944.3	10.885	1.793	173.772	87.139	32992.107	35133.834
	Instance1944.4	10.825	1.567	174.001	87.260	32987.803	35113.760
MBXCAS2	Instance3096.1	15.712	2.257	173.835	87.344	32934.209	35141.682
	Instance3096.2	15.715	1.939	173.598	86.890	32944.984	35140.708
	Instance3096.3	10.803	1.745	173.850	87.040	32948.406	35115.855
	Instance3096.4	10.822	1.504	173.725	87.151	32944.705	35133.099
MBXCAS3	Instance2128.1	19.255	2.833	150.981	75.304	33181.164	35317.917
	Instance2128.2	19.306	2.565	150.927	75.233	33168.242	35285.369
	Instance2128.3	10.708	1.600	151.350	75.530	33186.376	35269.149
	Instance2128.4	10.805	1.456	151.361	75.606	33180.064	35287.384

Table 15. Disk Subsystem Transactional I/O Performance Jetstress Results (Continued)

Virtual Machine Name	Exchange DB Instances	I/O DB Reads Average Latency (msec)	I/O DB Writes Average Latency (msec)	I/O DB Reads/ sec	I/O DB Writes/ sec	I/O DB Reads Average Bytes	I/O DB Writes Average Bytes
MBXCAS4	Instance3808.1	19.334	2.791	151.351	75.380	33124.107	35296.890
	Instance3808.2	19.382	2.532	151.364	75.500	33113.849	35304.034
	Instance3808.3	10.826	1.565	151.513	75.491	33137.813	35271.565
	Instance3808.4	11.135	1.406	151.620	75.522	33139.726	35285.636

Load Generator Normal Test Results

The Load Generator normal test results show four Exchange virtual machines configured with 4,000 mailboxes of 2 GB in size successfully handled the loads.

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

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

VM Name	Number of Mailboxes	Memory % Committed	CPU % Utilization
MBXCAS1	1,000	47	9
MBXCAS2	1,000	55	11
MBXCAS3	1,000	26	18
MBXCAS4	1,000	40	8

Table 17. Simulation Statistics Load Generator Results

Type	Client Name	Tasks Completed	Task Dispatched	Total Task Completed
Remote	ws12-lg2	5372	5371	64338
Remote	ws12-lg3	5388	5388	
Remote	ws12-lg4	5404	5403	
Remote	ws12-lg5	5420	5420	
Remote	ws12-lg7	5421	5421	
Remote	ws12-lg8	5420	5418	
Remote	ws12-lg10	5420	5420	
Remote	ws12-lg11	5421	5421	
Remote	ws12-lg12	5422	5422	
Remote	ws12-lg9	5422	5422	
Remote	ws12-lg13	4811	4811	
Remote	ws12-lg14	5421	5421	

Load Generator Failover Test Results

The Load Generator failover test results show three Exchange virtual machines configured with 4,000 mailboxes of 2 GB in size successfully handled the loads.

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

Table 18. Memory Percent Committed and CPU Utilization Load Generator Results

VM Name	Number of Mailboxes	Memory % Committed	CPU % Utilization
MBXCAS1	2,000	45	45
MBXCAS3	1,000	38	38
MBXCAS4	1,000	40	40

Table 19. Simulation Statistics Load Generator Results

Type	Client Name	Tasks Completed	Task Dispatched	Total Task Completed
Remote	ws12-lg2	5370	5370	64330
Remote	ws12-lg3	5386	5386	
Remote	ws12-lg4	5402	5402	
Remote	ws12-lg5	5419	5419	
Remote	ws12-lg7	5419	5419	
Remote	ws12-lg8	5419	5419	
Remote	ws12-lg10	5420	5420	
Remote	ws12-lg11	5421	5421	
Remote	ws12-lg12	5422	5422	
Remote	ws12-lg9	5420	5420	
Remote	ws12-lg13	5421	5421	
Remote	ws12-lg14	4811	4811	

Conclusion

This reference architecture guide describes a Microsoft Private Cloud Fast Track v4 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

AS-293-00, March 2014