TECHNICAL PAPER

Deploy SAP ERP 3-tier Using Hitachi Unified Storage VM in a Scalable Environment

Reference Architecture Guide

By Prasad Patkar

May 21, 2013

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 SAP ERP 3-tier using Hitachi Unified Storage VM in a Scalable Environment

Reference Architecture Guide

This is a reference guide for SAP ERP three tier architecture using Hitachi Unified Storage VM. It contains advice on how to build a virtual infrastructure that meets the unique requirements of your organization, providing the flexibility to scale out as organizational needs grow.

The benefits of this solution include the following:

- Faster deployment
- Reduced risk
- Predictability
- Ability to scale out
- Lower cost of ownership

This guide documents how to deploy this configuration using the following:

- Hitachi Compute Blade 2000
- Hitachi Unified Storage VM
- Hitachi Dynamic Provisioning
- VMware vSphere 5.1

Use this document to support sales, support, and appliance building by understanding the SAP ERP architecture and deployment.

This solution supports three different configurations. The SAP quick sizer and benchmarking tools are used as a base to identify the number of SAP Application Performance Standard (SAPS) and to determine the size of the configuration. Table 1 lists the SAP ERP configuration sizes for three tier architecture.

Table 1. SAP ERP Configuration Sizes

Configuration Size	***Maximum supported SAPS per node***	***Number of SD Users supported per node***	***Number of CB2000 X57A2 blades per node***	***SMP connector***
Small Node	18000	3400	1	None
Medium Node	36000	6800	2	2-blade SMP per node
Large Node	72000	13600	4	4-blade SMP per node

This technical paper contains advice on how to build a virtual infrastructure for a small node configuration. It assumes you have familiarity with the following:

- Storage area network-based storage systems
- General storage concepts
- General network knowledge
- Common IT storage practices

Note — Testing of the small node configuration was in a lab environment. The results obtained from the small node configuration tests are used for sizing the medium and large node configurations. 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 uses a VMware infrastructure supported by Hitachi hardware. The following components create this SAP ERP solution:

- **Hitachi Compute Blade 2000** — An enterprise-class server platform
- **Hitachi Unified Storage VM** — Hitachi Unified Storage VM storage virtualization system is designed for organizations that need to manage their storage assets more efficiently
- **Hitachi Dynamic Provisioning** — Provides wide striping and thin provisioning functionalities for greater operational and storage efficiency
- **VMware vSphere 5.1** — Virtualization technology providing the infrastructure for the data center
- **Emulex dual port Fibre Channel Host Bus Adapters** — Provides SAN connectivity to the servers

Figure 1 illustrates the high-level logical design of this reference architecture for a small node configuration on Hitachi Unified Storage VM and Hitachi Compute Blade 2000.

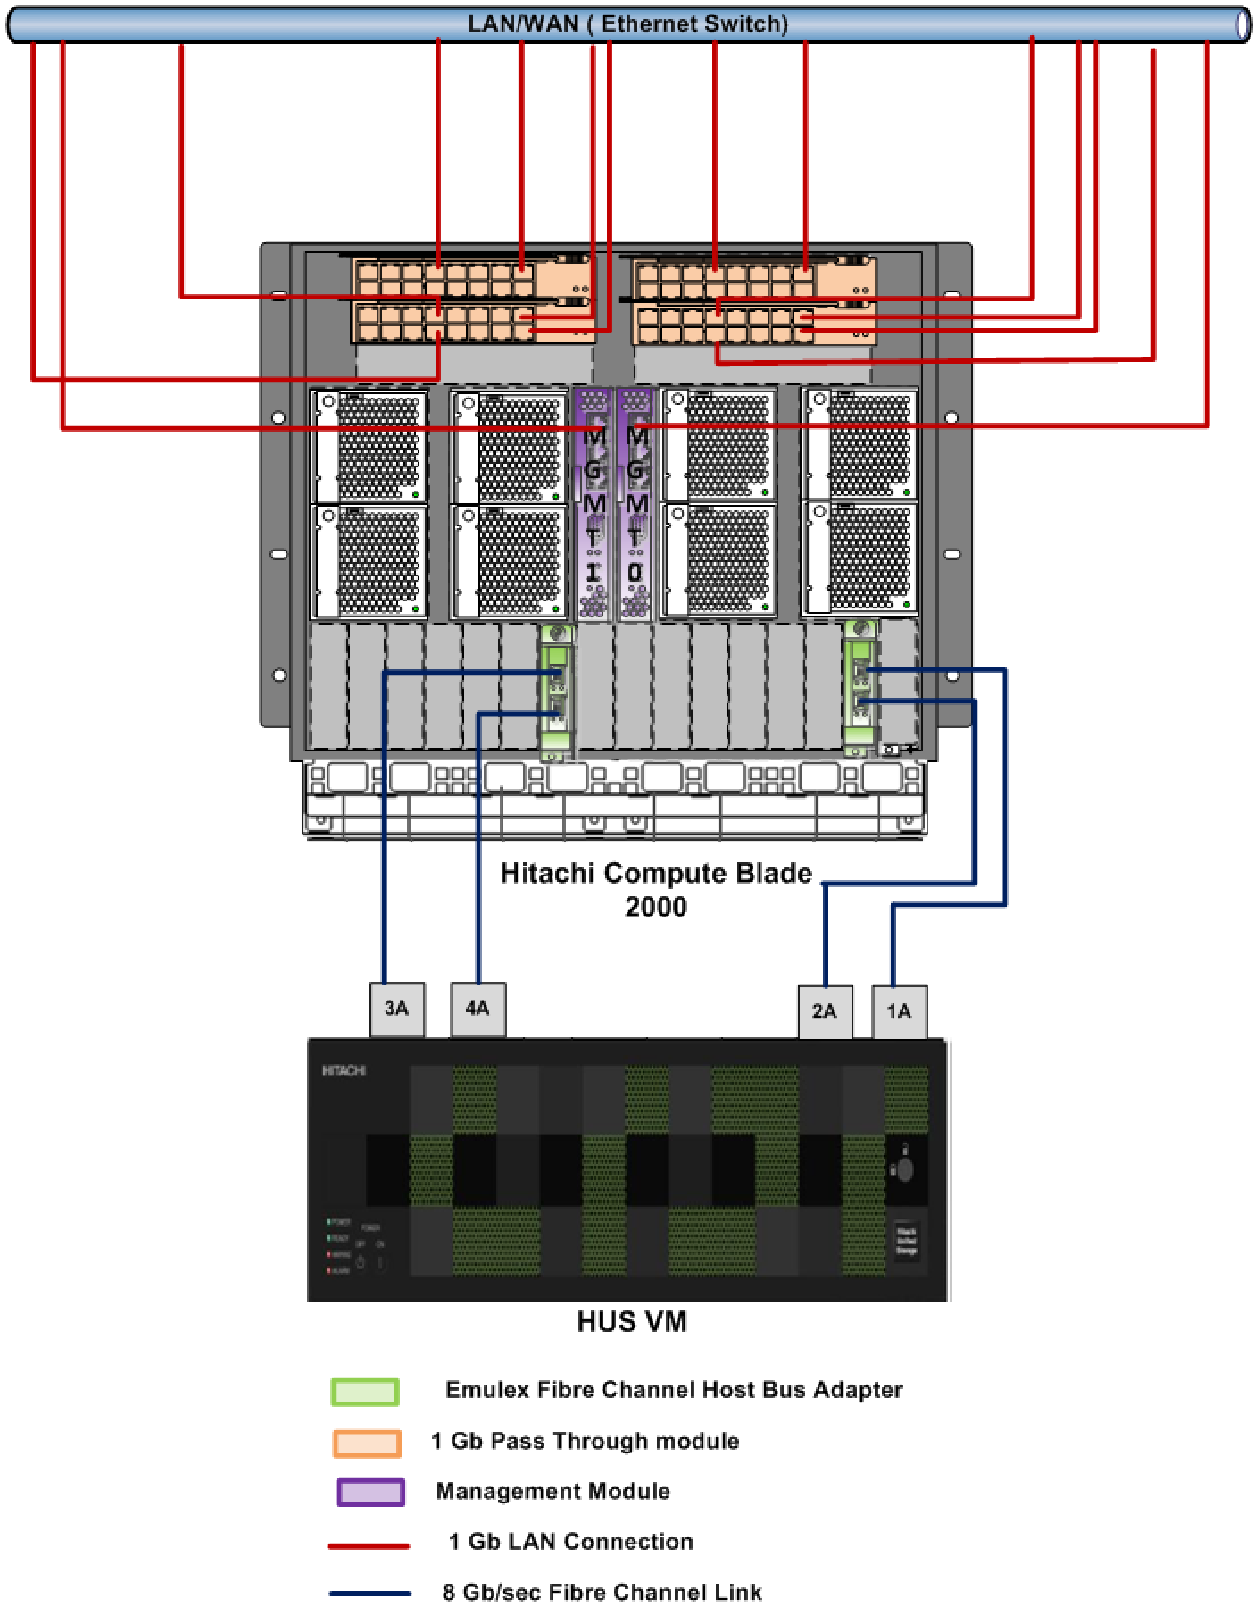
LAN/WAN (Ethernet Switch)
M G M T 1
M G M T 0
Hitachi Compute Blade 2000
3A
4A
2A
1A
HITACHI
HUS VM
Emulex Fibre Channel Host Bus Adapter
1 Gb Pass Through module
Management Module
1 Gb LAN Connection
8 Gb/sec Fibre Channel Link

Figure 1

Key Solution Elements

These are the key hardware and software elements used in this reference architecture.

Hardware Elements

Table 2 lists the detailed information about the hardware components used in the Hitachi Data Systems lab to validate the **small node configuration**.

Table 2. Hardware Elements

Hardware	Quantity	Configuration	Role
Global Solution Rack	1	▪ 4 x 1U PDUs for chassis ▪ 2 x side PDUs for HUS VM	Rack
Hitachi Compute Blade 2000 chassis	1	▪ 8-blade chassis ▪ 2 management modules ▪ 8 cooling fan modules ▪ 1 × 1 Gb/sec LAN pass-through module per chassis	Server blade chassis
X57A2 server blade	2	▪ 2 × 10-core processors ▪ 256 GB RAM	2 Nodes- 1 node for production and 1 node for non-production
Emulex HBA	8	▪ 8 Gb/sec dual port Fibre Channel HBA	Host bus adapters
1GbE 4-port LAN mezzanine card	2	▪ Slot 0 of each blade	Network connectivity
HUS VM	1	▪ 64 GB cache ▪ 8 x 8Gb FC ports ▪ 6 expansion trays ▪ 168 x 600 GB 10k SAS drives	Primary storage

Hitachi Compute Blade 2000

Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

The small, medium, and large node configurations use two, four and eight X57A2 server blades respectively in the Hitachi Compute Blade chassis.

Table 3 has the specifications for the Hitachi Compute Blade 2000 used in this solution.

Table 3. X57A2 Server Blade Configuration

Feature	Configuration
Processors	▪ Intel Xeon processor E7-8800 ▪ 2 processors per server blade
Processor SKU	▪ Intel Xeon processor E7-8870
Processor frequency	▪ 2.4 GHz
Processor cores	▪ 10 cores
Memory DIMM slots	▪ 32
Memory	▪ 256 GB RAM ▪ 8 GB DIMMs
Network ports	▪ 2 × 1Gb Ethernet
Other interfaces	▪ 2 USB 2.0 port ▪ 1 serial port

Hitachi Unified Storage VM

Hitachi Unified Storage VM is an entry-level enterprise storage platform. It combines storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a storage virtualization system to provide central storage services to existing storage assets.

Unified management delivers end-to-end central storage management of all virtualized internal and external storage on Unified Storage VM. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, as well as virtual NAS functionality, without compromising performance or scalability.

The benefits of Unified Storage VM are the following:

- Enables the move to a new storage platform with less effort and cost when compared to the industry average
- Increases performance and lowers operating cost with automated data placement
- Supports scalable management for growing and complex storage environment while using fewer resources
- Achieves better power efficiency and with more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions
- Consolidates management with end-to-end virtualization to prevent virtual server sprawl

Software Elements

Table 4 describes the software products used to deploy this reference architecture.

Table 4. Software Elements

Software	*Version*
Hitachi Storage Navigator Modular 2	Microcode Dependent
Hitachi Dynamic Provisioning	Microcode Dependent
VMware vCenter server	5.1.0
VMware Virtual Infrastructure Client	5.1.0
VMware ESXi	5.1.0
RedHat Enterprise Linux	6.2
Oracle	11.2.0.3
SAP ERP	ECC 6.0 EhP5 SPS08

Hitachi Storage Navigator Modular 2

Hitachi Storage Navigator Modular 2 provides essential management and optimization of storage system functions. Using Java agents, Storage Navigator Modular 2 runs on most browsers. A command line interface is available.

Use Storage Navigator Modular 2 for the following:

- RAID-level configurations
- LUN creation and expansion
- Online microcode updates and other system maintenance functions
- Performance metrics

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 Unified Storage VM, Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 1 GB chunks, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can re-balance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

VMware vSphere 5.1

VMware vSphere 5.1 is a virtualization platform that provides a data center infrastructure. It features vSphere Distributed Resource Scheduler (DRS), high availability, and fault tolerance.

VMware vSphere 5 has the following components:

- **ESXi 5.1** — This is a hypervisor that loads directly on a physical server. It partitions one physical machine into many virtual machines that share hardware resources.
- **vCenter Server 5.1** — This allows management of the vSphere environment through a single user interface. With vCenter, there are features available such as vMotion, Storage vMotion, Storage Distributed Resource Scheduler, High Availability, and Fault Tolerance.

SAP ERP

Use SAP Enterprise Resource Planning (ERP) to secure a sound foundation to compete in the global marketplace with efficient support for your specific industry's business processes and operations. ERP software is a proven foundation to support and streamline your business processes, no matter what the size of your enterprise. View Solutions for Enterprise Resource Planning to see different ERP applications.

A 3-tier configuration has separate operating systems for presentation, business logic, and database. The operating system can run on a physical machine or a virtual machine. Alternately, a 3-tier configuration can be a single system with separate operating systems when it is not possible to run one operating system on the whole system. This is different from a 2-tier solution, which executes on one system and has the capability to run under one operating system. See SAP Standard Application Benchmark Publication Process (PDF) for more information about benchmark definitions and standards.

Solution Design

This is detailed information on the SAP ERP reference solution. It includes information required to build the basic infrastructure for the virtualized data center environment.

This reference architecture guide includes the following:

- Hitachi Compute Blade 2000 Chassis Configuration
- Direct Connect Fibre Channel Architecture
- Storage Architecture
- Virtual machine configuration
- SAP ERP Configuration
- Network Architecture

Hitachi Compute Blade 2000 Chassis Configuration

Following are the three different configurations:

- **Small node configuration**— It consists of a single blade per node. Blade 0 is the production node and Blade 4 is the non-production node. Each node has a 2 x 10 core processor and 256 GB of memory.
- **Medium node configuration**— It consists of two blades per node. Blades 0 and 1 are configured using a 2-blade SMP connector to turn into one production node. Blades 4 and 5 are configured using a 2-blade SMP connector to turn into one non-production node. Each node has a 4 x 10 core processor with 512 GB of memory.
- **Large node configuration**— It consists of four blades per node. Blades 0, 1, 2 and 3 are configured using a 4-blade SMP connector to turn into one production node. Blades 4, 5, 6 and 7 are configured using a 4-blade SMP connector to turn into one non-production node. Each node has an 8 x 10 core processor and 1 GB of memory.

Figure 2 shows the front view of the Hitachi Compute Blade 2000 chassis for small, medium, and large node configurations.

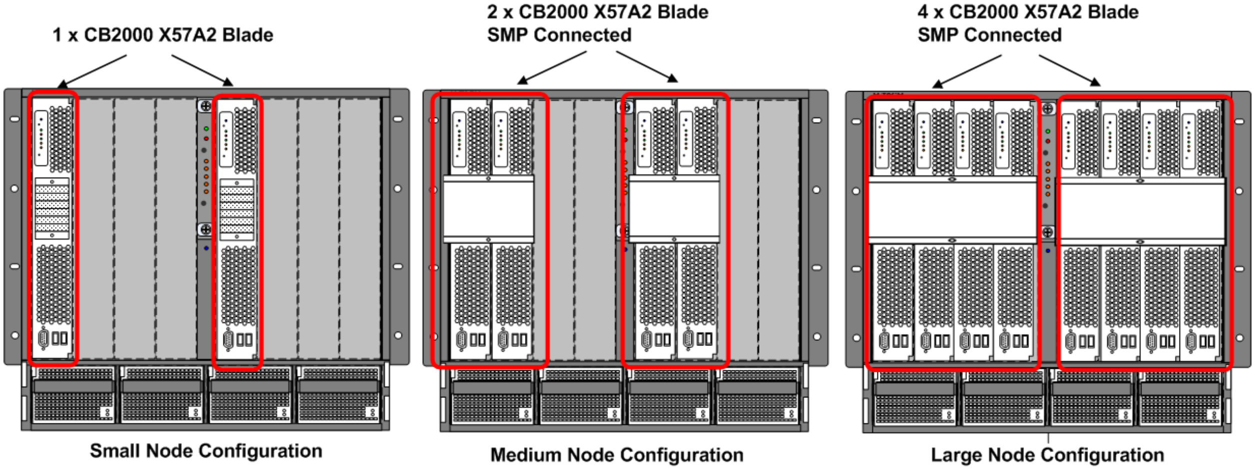

Figure 2

This design provides the flexibility to scale out as organizational needs grow.

Figure 3 shows the back view of the Hitachi Compute Blade 2000 chassis for small, medium, and large node configurations.

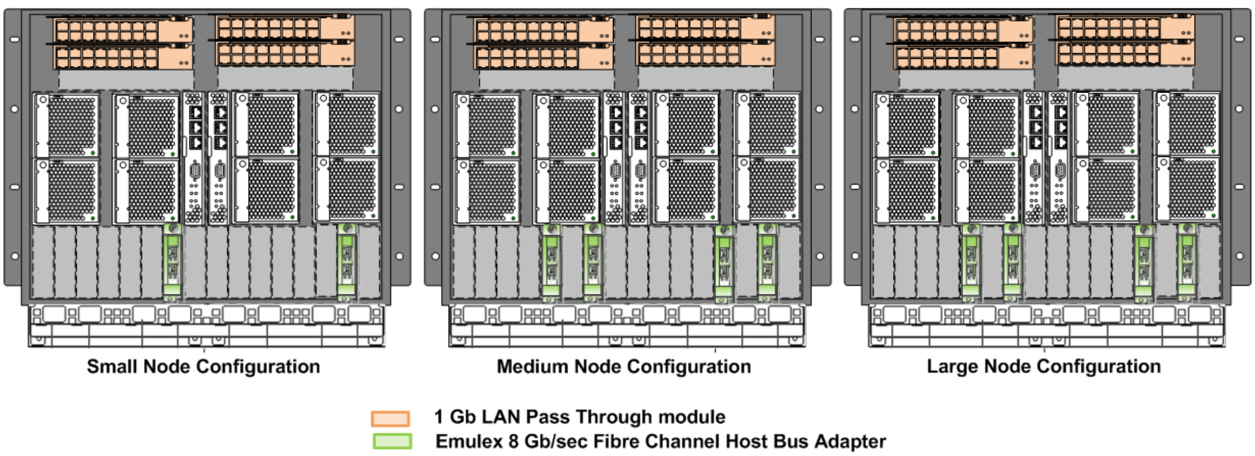

Figure 3

Use one LAN pass through module each in Switch Module 0, 1, 2 and 3. There are two PCIe slots available for each blade.

- **Small node configuration**— The right PCIe slot of blade 0 and blade 4 has one Emulex 8 Gb/sec dual port host bus adapter.
- **Medium node configuration**— The right PCIe slot of blade 0, blade 1, blade 4, and blade 5 has one Emulex 8 Gb/sec dual port host bus adapter.
- **Large node configuration**— The right PCIe slot of blade 0, blade 1, and blade 4 and blade 5 has one Emulex 8 Gb/sec dual port host bus adapter.

Direct Connect Fibre Channel Architecture

The direct connect Fibre Channel architecture has one Emulex Fibre Channel host bus adapter on each of the following for a direct connection to Hitachi Unified Storage VM.

- **Small node configuration**— The right PCIe slot of blade 0 and blade 4 has one Emulex 8 Gb/sec dual port host bus adapter.
- **Medium node configuration**— The right PCIe slot of blade 0, blade 1, blade 4, and blade 5 has one Emulex 8 Gb/sec dual port host bus adapter.
- **Large node configuration**— The right PCIe slot of of blade 0, blade 1, and blade 4 and blade 5 has one Emulex 8 Gb/sec dual port host bus adapter.

This direct-attached storage configuration provides better performance with the direct connection of Hitachi Unified Storage VM and the server blades, compared to a Fibre Channel switch connection.

This solution for a small node configuration uses Storage Port 1A, Storage Port 2A, Storage Port 3A, and Storage Port 4A on Hitachi Unified Storage VM.

- Port 1A and Port 2A connect to the Emulex Fibre Channel host bus adapter in the left PCIe slot of Server Blade 0 (production node) The Emulex Fibre Channel host bus adapter in the left PCIe slot of the Server Blade 0 connects to the following:
  - Port 1A connects to the top port
  - Port 2A connects to the bottom port
- Port 3A and Port 4A connect to the Emulex Fibre Channel host bus adapter in the left PCIe slot of Server Blade 4 (non-production node). The Emulex Fibre Channel host bus adapter in the left PCIe slot of Server Blade 4 connects to the following:
  - Port 3A connects to the top port
  - Port 4A connects to the bottom port

This configuration supports high availability by providing multiple paths from the hosts within Hitachi Compute Blade 2000 to multiple ports on Hitachi Unified Storage VM. In case of an Emulex HBA port failure, this redundancy gives the SAP server additional paths to Hitachi Unified Storage VM.

For the direct connection between Hitachi Compute Blade 2000 and Hitachi Unified Storage VM, set the Hitachi Unified Storage VM Fibre Channel ports to **loop topology**.

Table 5 shows the storage port mapping for a small node configuration.

Table 5. Storage Port Mapping for Small Node Configuration

Blade, Slot, Port	*Value*
Blade 0, Slot 1, Port 0	1A
Blade 0, Slot 1, Port 1	2A
Blade 4, Slot 1, Port 0	3A
Blade 4, Slot 1, Port 1	4A

Set the port properties for the direct connection between Hitachi Compute Blade 2000 and Hitachi Unified Storage VM as shown in Table 6.

Table 6. Port Properties

Property	*Value*
Port Attribute	Target
Port Security	Disabled
Port Speed	Auto (8Gbps)
Fabric	Off
Connection Type	FC-AL

Figure 4 shows the direct connect Fibre Channel architecture for a small node configuration.

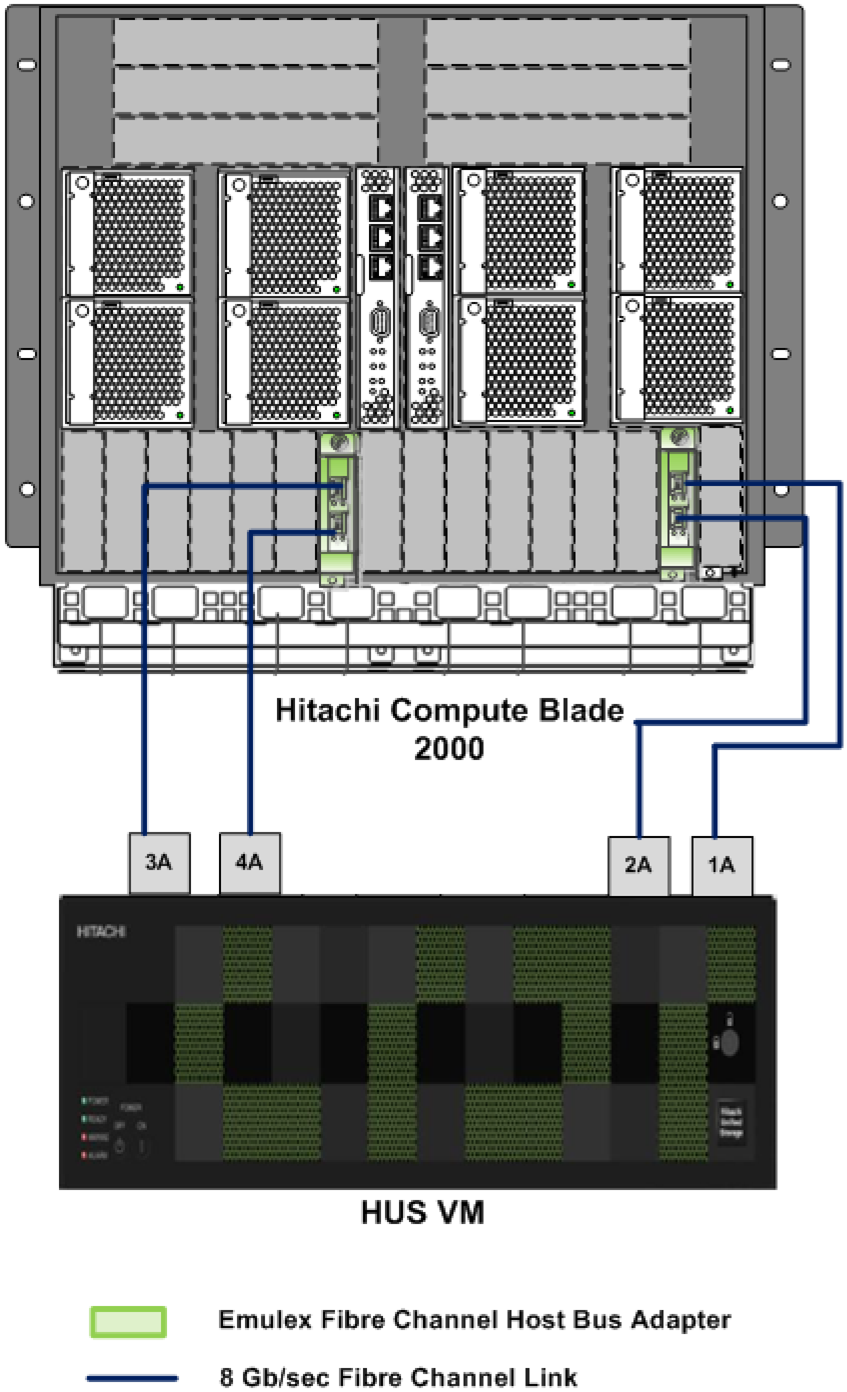

Figure 4

Storage Architecture

Table 7 shows the Hitachi Unified Storage VM components

Table 7. Hitachi Unified Storage Platform VM Components

Storage System	*Hitachi Unified Storage VM*
Microcode Level	73-02-01-00/00
Cache Memory	64 GB
Number of ports	8
CHB	2 Pairs
DKB	2 Pairs
RAID Group Type	RAID-5 (7D+1P) -OS RAID-5 (3D+1P)- Binaries and Log RAID-10 (2D+2D)- Data
Number of Drives	168
Drive Capacity	600 GB
Drive Type	SAS 10K RPM

Many factors drive the sizing and configuring of storage. This includes I/O and capacity requirements. The following describe how the storage sizing for this reference architecture was determined:

- Parity Group Configuration
- LDEV Configuration
- Storage Requirements

Parity Group Configuration

This reference architecture uses the following RAID configuration on Hitachi Unified Storage VM.

- Two RAID-5 (7D+1P) parity group created using sixteen 600 GB SAS 10k RPM drives.
- Seven RAID-5 (3D+1P) parity groups created using twenty eight 600 GB SAS 10k RPM drives.
- Thirty RAID-10 (2D+2D) parity groups created using one hundred twenty 600 GB SAS 10k RPM drives.
- Four 600 GB SAS 10k RPM drives as spare drives.

Table 8 has the configuration for each parity group.

Table 8. Parity Groups

Parity Group	RAID Level	Drive Size	Drive Speed	Usable Total Capacity	Usage
1	RAID-5 (7D+1P)	600 GB	10k RPM	3.7 TB	Operating system
2	RAID-5 (7D+1P)	600 GB	10k RPM	3.7 TB	Data store-destination storage for virtual machines
3	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Production server SAP binaries
4	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Production server Oracle binaries
5 and 6	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Production server Logs
7	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Non-Production server SAP binaries
8	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Non-Production server Oracle binaries
9	RAID-5 (3D+1P)	600 GB	10k RPM	1.6 TB	Non-Production server Log
10 to 29	RAID-10 (2D+2D)	600 GB	10k RPM	1 TB	Production server Data
30 to 39	RAID-10 (2D+2D)	600 GB	10k RPM	1 TB	Non-production server Data

LDEV Configuration

This reference architecture contains the following:

- Six 200 GB LDEVs to host boot operating system for six virtual machines
- One 3 TB LDEV to host the data store for the destination storage of virtual machines
- One 1.6 TB LDEV to host the production server SAP binaries
- One 1.6 TB LDEV to host the production server Oracle binaries
- Two 1.6 TB LDEVs to host the production server log volumes
- One 1.6 TB LDEV to host the non-production server SAP binaries
- One 1.6 TB LDEV to host the non-production server Oracle binaries
- One 1.6 TB LDEV to host the non-production server log volume
- Twenty LDEVs with capacity of 1 TB each to host the Hitachi Dynamic pool volume to store the production server data
- Ten LDEVs with capacity of 1 TB each to host the Hitachi Dynamic pool volume to store the non-production server data

Table 9 shows the LDEV allocation, volume group and file system for each parity group.

Table 9. LDEV Allocation, Volume Group and File System

Parity Group	LDEV	LDEV Size	Volume Group	File System, Size, Type	LDEV usage
1	1-2	200 GB	None- Mapped as raw device	root, 512 MB, ext3 swap, 132 GB, swap home, 20 GB, ext3	OS for Virtual Machine- Production servers- SAP CI and SAP DB
	3-6	200 GB	None- Mapped as raw device	root, 512 MB, ext3 swap, 66 GB, swap home, 20 GB, ext3	OS for Virtual Machine- Production servers- SAP application server 1, 2, 3 and 4
	7-10	200 GB	None-Mapped as raw device	root, 512 MB, ext3 swap, 66 GB, swap home, 20 GB, ext3	OS for Virtual Machines- Non Production servers- DEV CI, DEV DB, QA CI, QA DB
	11-12	200 GB	None-Mapped as raw device	root, 512 MB, ext3 swap, 33 GB, swap home, 20 GB, ext3	OS for Virtual Machines- Non Production servers- SAP QA application server 1, SAP QA application server 2

Table 9. LDEV Allocation, Volume Group and File System (Continued)

Parity Group	LDEV	LDEV Size	Volume Group	File System, Size, Type	LDEV usage
2	13	3 TB	None- Mapped as raw device	data store, 3 TB, VMFS5	Destination storage for virtual machines
3	14	1.6 TB	VG_BIN	lv_usrsap, 50 GB, ext3 lv_sapmnt, 50 GB, ext3 lv_trans, 200 GB, ext3 lv_media, 500 GB, ext3	SAP binaries Shared file system Transport directory Media share
4	15	1.6 GB	VG_ORACLE	lv_oracle, 100 GB, ext3 lv_oraarch, 1.5 TB, ext3	Oracle binaries Oraarch files
5	16	1.6 TB	VG_LOG1	lv_origlogA, 200 GB, ext3 lv_mirrlogA, 200 GB, ext3	Orig log A Mirror log A
6	17	1.6 TB	VG_LOG2	lv_origlogB, 200 GB, ext3 lv_mirrlogB, 200 GB, ext3	Orig log B Mirror log B
7	18	1.6 TB	VG_BIN_NP	lv_usrsap_dev, 50 GB, ext3 lv_sapmnt_dev, 50 GB, ext3 lv_trans_np, 200 GB, ext3 lv_usrsap_qa, 50 GB, ext3 lv_sapmnt_qa, 50 GB, ext3	Dev SAP binaries Dev Shared file system Dev/QA Transport directory QA SAP binaries QA Shared file system
8	19	1.6 GB	VG_ORACLE_NP	lv_oracle_dev, 100 GB, ext3 lv_oraarch_dev, 1.5 TB, ext3 lv_oracle_qa, 100 GB, ext3 lv_oraarch_qa, 1.5 TB, ext3	Dev Oracle binaries Dev Oraarch files QA Oracle binaries QA Oraarch files

Table 9. LDEV Allocation, Volume Group and File System (Continued)

Parity Group	LDEV	LDEV Size	Volume Group	File System, Size, Type	LDEV usage
9	20	1.6 TB	VG_LOG_NP	lv_origlogA_dev, 200 GB, ext3 lv_mirrlogA_dev, 200 GB, ext3 lv_origlogB_dev, 200 GB, ext3 lv_mirrlogB_dev, 200 GB, ext3 lv_origlogA_qa, 200 GB, ext3 lv_mirrlogA_qa, 200 GB, ext3 lv_origlogB_qa, 200 GB, ext3 lv_mirrlogB_qa, 200 GB, ext3	Dev Orig log A Dev Mirror log A Dev Orig log B Dev Mirror log B QA Orig log A QA Mirror log A QA Orig log B QA Mirror log B
10 to 29	21 to 40	1 TB	VG_DATA	lv_data1 to lv_data20, 1 TB, ext3	HDP pool for PRD data
30 to 39	41 to 50	1 TB	VG_DATA_NP	lv_data21 to lv_data30, 1 TB, ext3	HDP pool for Non PRD data

Hitachi Dynamic Provisioning is used to create a dynamic provisioning pool on the Hitachi Unified Storage VM for storing data. LDEVs 21 to 40 are used as pool volumes for production data pool and LDEVs 41 to 50 are used as pool volumes for non-production data pool. Each of these pools has virtual volumes. Logical Volume Manager is used to create a file system for data on these virtual volumes.

Production server LDEVs are assigned to Storage Port 1A and Storage Port 2A. Non production server LDEVs are assigned to Storage Port 3A, and Storage Port4A on Hitachi Unified Storage VM. These LDEVs are assigned to the virtual machines using raw device mapping.

Virtual Machine Configuration

With hyper-threading enabled on the 2 × 10 Core Intel Xeon E7-8870 processors, 40 physical CPUs are available for each node. There were 6 virtual machines configured with 40 virtual CPUs on the production server blade.

Table 10 has the specifications for the virtual machine used in this solution for production node.

Table 10. Virtual Machine Configuration for Production Node

Production Instance	Number of virtual machines	Number of vCPU per virtual machine	vRAM per virtual machine	Purpose
Central Instance (CI)	1	8	64 GB	Primary application server
Database Instance (DB)	1	8	64 GB	Database Server
Dialog Instance (DI)	4	6	32 GB	Additional application server

Table 11 has the specifications for the virtual machine used in this solution for non-production node.

Table 11. Virtual Machine Configuration for Non-Production Node

Production Instance	Number of virtual machines	Number of vCPU per virtual machine	vRAM per virtual machine	Purpose
Central Instance (CI) for development	1	2	32 GB	Primary application server
Database Instance (DB) for development	1	2	32 GB	Database Server
Central Instance (CI) for quality assurance	1	2	32 GB	Primary application server
Database Instance (DB) for quality assurance	1	2	32 GB	Database Server
Dialog Instance (DI)	2	2	16 GB	Additional application server

Note- More virtual machines can be created on the non-production node with varied configuration depending on the requirements.

SAP ERP Configuration

This explains the SAP ERP configuration.

The Logical Volume Manager for the Linux operating system is used to configure the SAP ERP file system.

SAP ERP Data Volume Configuration

Hitachi Dynamic Provisioning is used to create a dynamic provisioning pool on the Hitachi Unified Storage VM. There is one pool "SAP_PRD_DATA" dedicated for production server data and another pool "SAP_NP_DATA" dedicated for non-production server data. Each of these pools has virtual volumes. Logical Volume Manager is used to create the file system for data on these virtual volumes.

SAP ERP Software Installation

After configuring the file system for the SAP ERP file system, the latest version of SAP ERP is installed.

The following profile parameters are set on the application servers to get optimal performance.

- Default profile

```
rdisp/tm_max_no=800

rdisp/TRACE=1

rdisp/bufrefmode=sendon,exeoff

rdisp/vb_delete_after_execution=0

rdisp/vbmail=0

rdisp/vb_dispatching=0

rdisp/delete_ddlog=0

rdisp/accept_remote_trace_level=0

rdisp/appc_ca_bmk_no=1500

rdisp/autoabaptime=0

rdisp/bufreftime=20000

rdisp/elem_per_queue=4000

rdisp/max_comm_entries=1200

rdisp/max_wprun_time=0

rdisp/ROLL_SHM=32768

rdisp/ROLL_MAXFS=32768

rdisp/PG_SHM=32768

rdisp/GP_MAXFS=32768

rdisp/start_icman=FALSE

rdisp/version_check=off

rdisp/wp_ca_blk_no=800
```

```
ipc/shm_protect_disabled=true
abap/buffersize=500000
abap/pxa=shared,unprotect
abap/initrc_degree=0
abap/no_sapgui_rfc=0
em/initial_size_MB=4608
em/max_size_MB=4608
em/mem_reset=off
es/use_shared_memory=TRUE
es/implementation=std
es/use_mprotect=FALSE
login/multi_login_users=sap_perf
login/end_of_license=0
itsp/enable=0
nobuf/max_no_buffer_entries=5000
icm/ccms_monitoring=false
gw/max_conn=800
rsdb/ntab/entrycount=25000
rsdb/ntab/ftabsize=35000
rsdb/ntab/sntabsize=1024
rsdb/ntab/irdbsize=5000
rsdb/otr/buffersize_kb=2048
rsdb/esm/buffersize_kb=2048
rsdb/obj/buffersize=80000
rsdb/obj/max_objects=20000
rsdb/max_blocking_factor=40
rsdb/max_in_blocking_factor=40
rsdb/min_blocking_factor=5
rsdb/min_in_blocking_factor=5
rsdb/prefer_fix_blocking=0
rsdb/prefer_in_itab_opt=0
```

```
rsdb/prefer_union_all=1
rtbb/buffer_length=3072
zcsa/db_max_buftab=30000
zcsa/table_buffer_area=50000000
zcsa/presentation_buffer_area=5000000
zcsa/calendar_area=250000
ztta/roll_area=3000000
ztta/roll_extension_dia=350000000
ztta/dynpro_area=800000
```

- Central Instance- Instance Profile

```
ipc/shm_psize_01=-40
ipc/shm_psize_02=-40
ipc/shm_psize_03=-40
ipc/shm_psize_04=-40
ipc/shm_psize_05=-40
ipc/shm_psize_06=-40
ipc/shm_psize_07=-40
ipc/shm_psize_08=-40
ipc/shm_psize_09=-40
ipc/shm_psize_10=136000000
ipc/shm_psize_18=-40
ipc/shm_psize_19=-40
ipc/shm_psize_30=-40
ipc/shm_psize_31=-40
ipc/shm_psize_33=-40
ipc/shm_psize_34=-40
ipc/shm_psize_40 = 112000000
ipc/shm_psize_41=-40
ipc/shm_psize_51=-40
ipc/shm_psize_52=-40
ipc/shm_psize_54=-40
```

```
ipc/shm_psize_55=-40
ipc/shm_psize_57=-40
ipc/shm_psize_58=-40
ipc/shm_psize_62=-40
ipc/shm_psize_63=-40
ipc/shm_psize_64=-40
ipc/shm_psize_65=-40
ipc/shm_psize_81=-40
ipc/shm_psize_1002=-40
ipc/shm_psize_58900100=-40
ipc/shm_psize_58900102=-40
em/largepages=TRUE
```

- Dialog Instance- Instance Profile

```
ipc/shm_psize_10 = 136000000
ipc/shm_psize_40 = 112000000
ipc/shm_psize_01=-40
ipc/shm_psize_02=-40
ipc/shm_psize_03=-40
ipc/shm_psize_04=-40
ipc/shm_psize_05=-40
ipc/shm_psize_06=-40
ipc/shm_psize_07=-40
ipc/shm_psize_08=-40
ipc/shm_psize_09=-40
ipc/shm_psize_10=136000000
ipc/shm_psize_18=-40
ipc/shm_psize_19=-40
ipc/shm_psize_30=-40
ipc/shm_psize_31=-40
ipc/shm_psize_33=-40
ipc/shm_psize_34=-40
```

```
ipc/shm_psize_40 = 112000000
ipc/shm_psize_41=-40
ipc/shm_psize_51=-40
ipc/shm_psize_52=-40
ipc/shm_psize_54=-40
ipc/shm_psize_55=-40
ipc/shm_psize_57=-40
ipc/shm_psize_58=-40
ipc/shm_psize_62=-40
ipc/shm_psize_63=-40
ipc/shm_psize_64=-40
ipc/shm_psize_65=-40
ipc/shm_psize_81=-40
ipc/shm_psize_1002=-40
ipc/shm_psize_58900100=-40
ipc/shm_psize_58900102=-40
```

Note — The parameters listed above are from the small node configuration tests performed in the 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.

Network Architecture

Hitachi Compute Blade 2000 contains the network hardware as shown in Table 12.

Table 12. Hitachi Compute Blade 2000 Network Hardware

NICs (per Blade)	■ 2 Onboard Intel 82576 Gigabit Ethernet ports
Mezzanine Slot 0 (per Blade)	■ 1 Ethernet mezzanine card ■ 4 × 1 Gb/sec ports
Switch Bay 0	■ 1 Gb LAN pass-through module ■ 16 × 1 Gb/sec ports
Switch Bay 1	■ 1 Gb LAN pass-through module ■ 16 × 1 Gb/sec ports
Switch Bay 2	■ 1 Gb LAN pass-through module ■ 16 × 1 Gb/sec ports
Switch Bay 3	■ 1 Gb LAN pass-through module ■ 16 × 1 Gb/sec ports

There are 1 Gb/sec pass-through modules installed in Switch Bay 0, Switch Bay 1, Switch Bay 2, and Switch Bay 3 of the Hitachi Compute Blade 2000 chassis.

Each blade has two LAN on motherboard NIC ports and connects through the chassis mid-plane to the internal ports of the LAN pass-through module in Switch Bay 0 and Switch Bay 1. There is one 4-port 1 Gb/sec LAN Mezzanine card on Slot 0 of each blade. Thus there are six network ports per blade.

Hitachi Compute Blade 2000 chassis has two management modules for redundancy. Each module supports an independent management LAN interface from the data network for remote and secure management of the chassis and all blades. Each module supports a serial command line interface and a web interface. It also supports SNMP and email alerts. Each module is hot-swappable and supports live firmware updates without the need for shutting down the blades.

Figure 5 shows the network connections for a small node configuration.

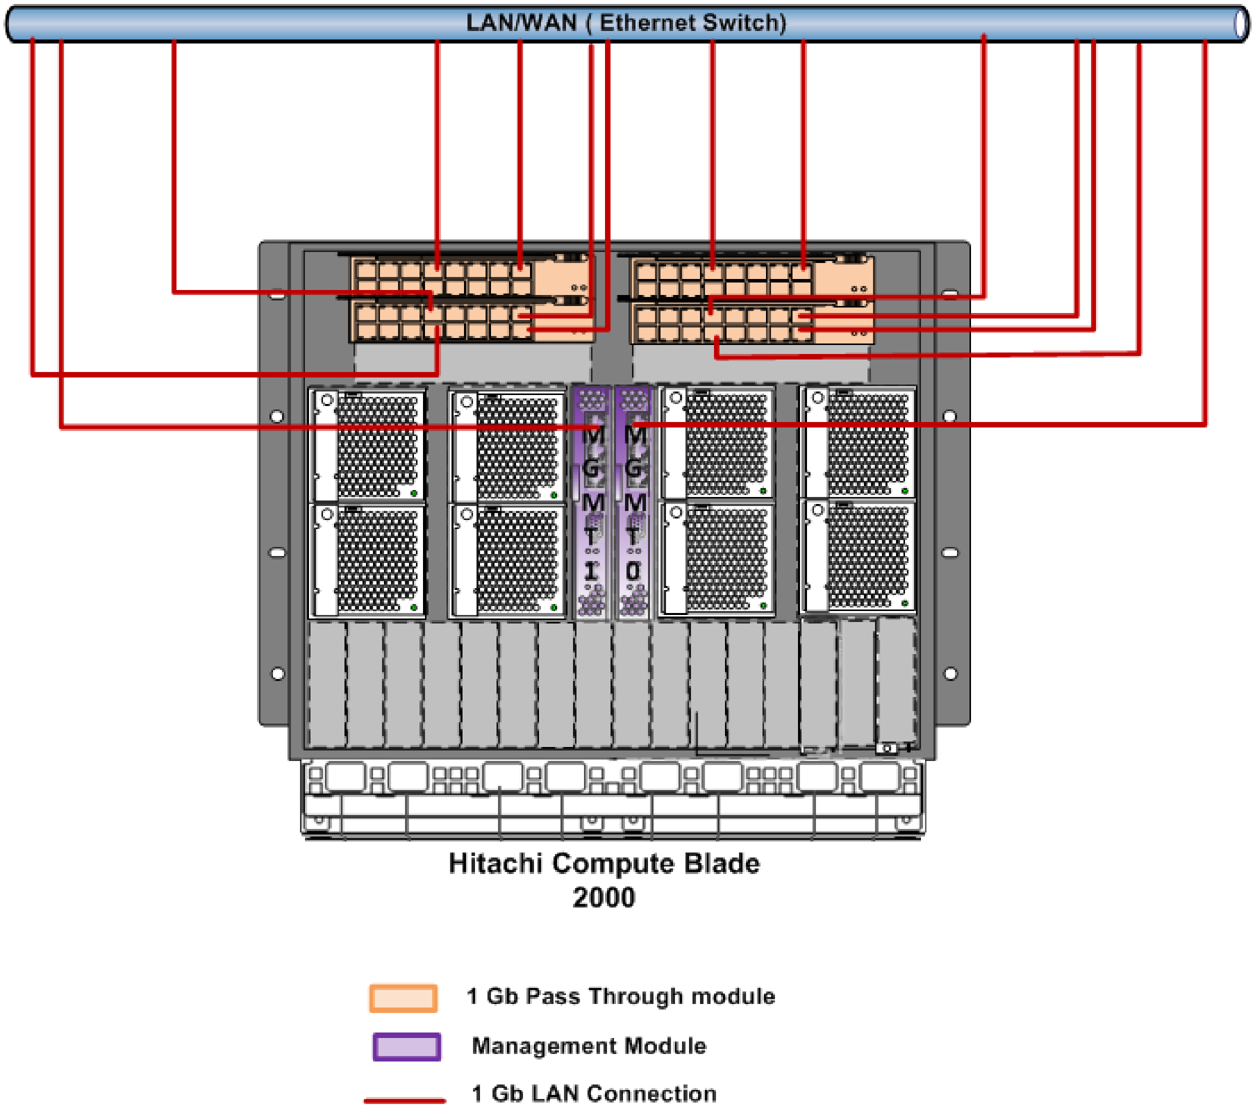

Figure 5

Set up an IPV4 address on the network adapter of the ESX host and set a public IP. Create one virtual network adapter on each VM. Assign the public network to this virtual adapter.

Engineering Validation

Validation of the SAP ERP reference solution was conducted in the Hitachi Data Systems laboratory. The validation testing includes KPI performance test cases using the SAP SD benchmarking tool kit designed and executed by Hitachi Data System.

Test Methodology

Because of intellectual property limitations, this paper does not include the test methodology.

Test Results

Because of intellectual property limitations, this paper does not include the test results or an analysis of those results.

Conclusion

This reference architecture guide discusses how to design an SAP ERP solution with Hitachi Unified Storage VM. The purpose of the SAP benchmark testing was to provide general guidance on the optimal resources available with this solution.

Each implementation has its own unique set of application requirements. Design your implementation of this environment by understanding the I/O workload and the SAP Application Performance Standard (SAPS) in your environment. Creating an environment that meets your unique needs results in increased ROI from avoiding over or under provisioning resources.

Having the capability to add additional blades to an existing node allows for non-disruptive upgrades to the underlying infrastructure. This provides immediate benefits to your environment that gives you the flexibility to scale out as your organization needs grow.

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-224-00, May 2013