SAP Business Suite on HANA on Hitachi Unified Compute Platform 4000 for VMware vSphere

TECHNICAL PAPER

Reference Architecture Guide

By Mark Walker and Milind Pathak

August 2015

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.

Contents

SAP Business Suite on HANA on Hitachi Unified Compute Platform 4000 for VMware vSphere

Reference Architecture Guide

This reference architecture provides a design concept for running SAP Business Suite on Hitachi Unified Compute Platform 4000 for VMware vSphere.

The Hitachi Unified Compute Platform solution and its tight integration with VMware vSphere allows for complete application landscapes, such as SAP Business Suite, to be designed and described at the virtualization level. As a result, it is no longer necessary for customers to be concerned with the underlying hardware configuration of Hitachi Unified Compute Platform for VMware vSphere, because all application interaction and management is provided at the virtualization level. Thus, this reference architecture should rather be considered a “virtual” reference architecture that can be applied to any UCP 4000 for VMware vSphere solution, independent of hardware specifics.

Solution Overview

The reference architecture uses a concept based on simplified building blocks (Virtual Machines), each of which hosts a single 3-Tier SAP system. These building blocks provide a simple and elastic basis for efficient management and operation of SAP Landscapes running on Hitachi Unified Compute Platform. The reference architecture paper also goes into some detail to explain how to size these building blocks based on typical SAP System sizes. Although advanced topics, such as High-Availability Clusters, Data Protection etc., are beyond the scope of this specific paper, the SAP Adaptive Computing configuration used within the reference architecture should provide a good foundation to incorporate such functionality into the design. In order to provide engineering validation points for this virtual reference architecture, the SAP ERP (Enterprise Resource Planning) version ECC6 EHP 7 component of SAP Business Suite was installed and tested using the design concept proposed in this reference architecture paper.

This reference architecture uses a single Hitachi Compute Blade 500 (CB 500) chassis, two Hitachi 520H B3 server blades and two Hitachi Compute Rack 210H (CR 210H) management servers, as shown in Figure 1.

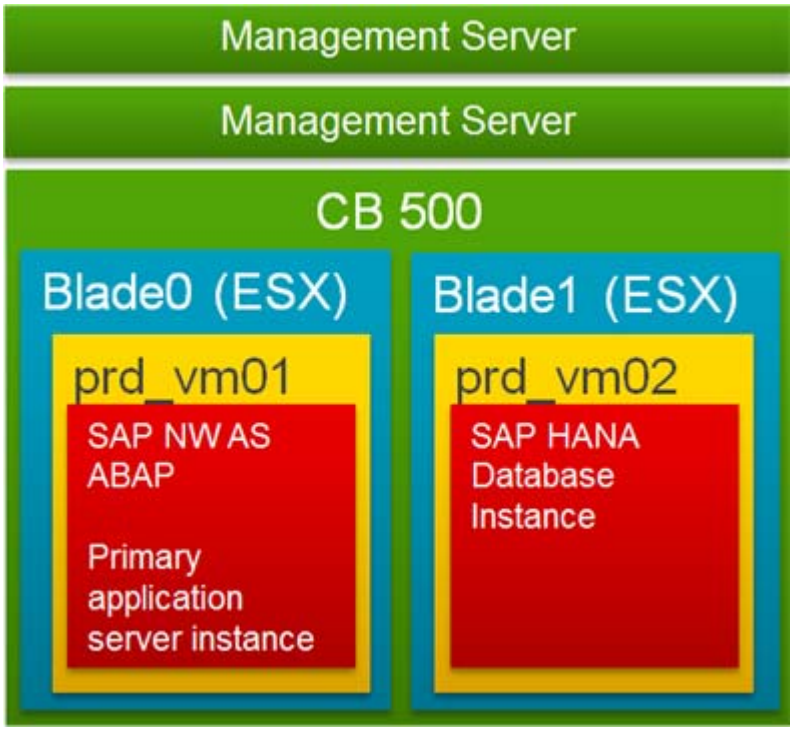

Figure 1

Networking is handled by Brocade FCX648 switches for the chassis' management network, and Brocade 6740 switches for the connection to the customer network. The Fibre Channel switches used for the SAN fabric are Brocade 6510 switches.

The SAP ERP was installed on one VMware virtual machine (VM) running on one 520H B3 server blade, while the underlying SAP HANA database was running on another VMware virtual machine that was running on the second 520H B3 server blade in the chassis.

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 Components

These are the major components used in this solution.

Hardware Components

Hitachi Unified Compute Platform 4000

Hitachi Unified Compute Platform for VMware vSphere offers the following:

- Full parity across the RESTful API
- Command line interface
- Graphical user interface

Hitachi Unified Compute Platform Director software on Unified Compute Platform integrates directly into VMware vSphere. It provides unified end-to-end infrastructure orchestration within a single interface.

Unified Compute Platform for VMware vSphere leverages your existing storage in one of two ways:

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

Unified Compute Platform for VMware vSphere provides the following benefits:

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

Hitachi 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 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 Hitachi 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 a complex storage environment while using fewer resources
- Achieves better power efficiency 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

Brocade Switches

Brocade and Hitachi Data Systems partner to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

The specific Unified Compute Platform in this engineering validation environment used the following Brocade products:

- Two Brocade 6510 8 Gb/sec Fibre Channel switches. There are two switches are at the back of each CB 500 chassis in slot 2 and slot 3.
- Two Brocade 6740 10 Gb/sec Ethernet switches connected to the CB 500 server blades and the customer network. There are two switches are at the back of each CB 500 chassis in slot 0 and slot 1.
- Two Brocade FCX648 1 Gb/sec Ethernet switches for the management network of the chassis

Software Components

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 42 MB pages, 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. An expanded pool can be rebalanced across the current and newly added RAID groups for an even striping of the data and the workload.

SAP Business Suite ECC6 EHP 7

SAP Business Suite provides customers with an end-to-end integrated platform to run all key business processes that occur on a day-to-day basis. The standard SAP-provided business processes cover all core commercial functions: financial, accounting, manufacturing, logistics, sales, marketing, and human resources. A particular specialty is the detail with which SAP provides localization, such as regional legislation and business practices. Industry specific applications such as banking, automotive and chemical further supplement the generic business processes. In addition, the SAP NetWeaver platform provides a powerful engine for SAP customers to develop and deploy their own business processes.

SAP Business Suite ECC6 EHP 7 consists of the following five core applications:

- SAP CRM (Customer Relationship Management) 7.0 EHP 3
- SAP ERP (Enterprise Resource Planning) 6.0 EHP 7
- SAP PLM (Product Lifecycle Management) 7.02
- SAP SCM (Supply Chain Management) 7.0 EHP 3
- SAP SRM (Supplier Relationship Management) 7.0 EHP 3

SAP NetWeaver is the underlying technology platform to power SAP Business Suite. SAP NetWeaver is a three-layer application stack: Presentation, Application Server, and Database. This three-layer application stack has one deployment scenario when used with the SAP HANA database: “3-Tier.”

In a 3-Tier scenario, the Database and Application run on separate servers, while the presentation layer runs on a separate operating system, e.g. the client.

SAP HANA Database

SAP HANA is an in-memory columnar database provided by SAP. Its in-memory characteristic offers unparalleled speed, while the columnar storage and data compression together keep the database size down to a minimum, allowing for all or nearly all of the data to be kept in main memory. SAP HANA's tight integration with the SAP Business Suite software provides an optimal system for running SAP's Business Suite - indeed several features of the latest versions of the Business Suite are only available for the HANA platform, leading to SAP renaming the Business Suite to "S/4 HANA" when used on top of the HANA database and when using the simplified data structures and user interface provided.

SUSE Linux Enterprise Server (SLES) 11 SP3

SLES Linux was the first operating system to be certified for use with SAP HANA and remains today the OS of choice for the vast majority of SAP HANA customers.

VMware vSphere 5.5

VMware vSphere 5.5 is a virtualization platform that provides a datacenter infrastructure. It features vSphere Distributed Resource Scheduler (DRS), High Availability, and fault tolerance.

VMware vSphere 5.5 has the following components:

- ESXi 5.5 — A hypervisor that loads directly on a physical server. It partitions one physical machine into many virtual machines that share hardware resources.
- vCenter Server — 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.

Solution Design

This section contains details concerning the following:

- SAP Architecture
- Sizing Unified Compute Platform for a SAP Business Suite Landscape
- Physical Server Configuration
- Network Architecture
- SUSE Linux Enterprise Server (SLES) Configuration
- VMware Configuration
- SAN Architecture
- Storage Architecture

SAP Architecture

In keeping with the Hitachi Unified Compute Platform value proposition of providing simplified infrastructure management, this reference architecture also aims for simplified management of SAP components. To achieve this, the reference architecture employs a building block concept leveraging the following SAP operational concepts:

- One 3-Tier SAP system with two virtual machines
  - Meaning
    - One virtual machine for the HANA database
    - One virtual machine for the SAP Business Suite
  - Reasoning
    - SAP HANA has very exclusive needs on its host machine
    - Allows a simplified mount point structure to support the SAP Adaptive Computing configuration
    - Troubleshooting & debugging of one SAP System without risk of affecting other SAP Systems in the landscape
    - Easier SAP and database patch management for the SAP System because each component is on a different system
- SAP Adaptive Computing configuration
  - Meaning
    - The SAP system is installed on a virtual IP address/DNS name, which is independent of the IP address/ hostname statically associated with the virtual machine and its operating system
    - The SAP Systems (database and application servers) are installed on shared storage, which is decoupled from the operating system and is able to be moved (detached and attached) from one virtual machine to another virtual machine
  - Reasoning
    - Flexibility to easily move (relocate) a SAP System to a new virtual machine, or even a physical server, if required
    - Simplified and lower risk OS upgrade path since the SAP System can simply be moved (relocated) to a virtual machine/blade already running the upgraded OS version
    - In combination with VMware physical RDM storage, the SAP System can be moved (relocated) into different Hypervisor technologies

Unified Compute Platform reference architecture with 520H B3 server blades allows for the building blocks to host 3-Tier SAP Systems up to 8000 SAPS.

Using this solution, it is possible to host several different HANA-based S-size Business Suite applications, such as CRM, ERP, and SRM. Common building blocks, for example production systems, can be grouped together into VMware clusters to leverage capabilities such as vMotion to move between host blades or high-availability failover software to offer redundancy between compute blades and/or chassis - see the Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere document for more details.

Sizing Unified Compute Platform for a SAP Business Suite Landscape

This reference architecture was sized as follows:

- SAP 3-Tier configuration
- SUSE Linux Enterprise Server (SLES) 11 SP3
- One VM for HANA SPS09 - with SAP System identifier (SID) “HIT”
  - 60 virtual cores (vCPU) mapped directly to 30 physical cores (pCPU)
  - Hyper-Threading On (60 logical CPU)
  - NUMA optimizations off
  - 240 GB RAM (30 × 8 GB per core) – see the Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere document for more details
- One VM for the SAP Business Suite System - SID “HS4”
  - ERP 6.0 EHP 7
  - 60 virtual cores (vCPU) mapped directly to 30 physical cores (pCPU)
  - Hyper-Threading On (60 logical CPU)
  - NUMA optimizations off
  - 240 GB RAM (30 × 8 GB per core) – see the Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere document for more details

In addition to SAP, storage for the operating systems also has to be considered. As described in the Storage Architecture section, this reference architecture makes use of VMware Datastores for the operating system.

Help Sizing SAP Systems on Hitachi Unified Compute Platform

Assistance is available from your HDS pre-sales representative and/or by contacting Hitachi Data Systems via email: sap@hds.com.

Physical Server Configuration

Hitachi Unified Compute Platform (UCP) uses two CR 210H rack mount servers, each running VMware ESX Hypervisors, to host the management and automation software to run UCP itself, and this is defined as the “UCP Management Datacenter” within VMware vCenter. See Table 1 for technical details.

For the application workload (SAP Business Suite applications), UCP can be configured with up to 16 CB 500 blade chassis. Each blade chassis can accommodate up to eight 520H B3 server blades. For technical details, see Table 2 and Table 3 respectively.

The actual number of blades required depends on both the number and size of SAP systems to be hosted. The CPU and memory per blade defined in Table 3 are already balanced and do not require modification.

Table 1. Unified Compute Platform Management Servers

Feature	Configuration
Server Type	■ CR 210H 1U Rack mount
CPU	■ Two Intel Xeon E5-2670 ■ Total 16 physical cores (8 per CPU) ■ 2.6 GHz
Memory	■ 192 GB (12 × 16 GB DIMMs)
Network	■ 2 Port 1 Gb/sec Ethernet onboard card ■ 2 Port 10 Gb/sec Ethernet (Emulex) PCI-Ex card
Fibre Channel	■ Hitachi 8Gb/sec 2-port Fibre Channel mezzanine card (FIVE-EX based Fibre Channel to PCIe HBA)

Table 2. Unified Compute Platform Compute Chassis

Feature	Configuration
Chassis Type	■ Hitachi Compute Blade 500 chassis ■ Scalable to Eight 520H B1/B2/B3 Server Blades
Power and cooling	■ Four 2142W Power Supplies ■ Six Fan Modules
Admin access	■ Two SVP Modules
Network Switches	■ Two Brocade 6740 - 24 Port 10 Gb/sec Ethernet Switch Modules (16 internal, 8 external)
Management Switch	■ Two Brocade FCX648 - 48 Port 1 Gb/sec Ethernet Switch Modules
Fibre Channel Switches	■ Two Brocade 6510 - 36 or 48 Port Fibre Channel Switch Modules

Table 3. Unified Compute Platform Compute Blades

Feature	Configuration
Server Blade	■ 520H B3
CPU	■ Two Intel Xeon E5-2699v3 ■ Total 36 physical cores (18 per CPU) ■ 3 GHz

Table 3. Unified Compute Platform Compute Blades (Continued)

Feature	Configuration
Memory	■ 384 GB (24 × 16 GB)
Network	■ Onboard 2 Port 10 Gb/sec Ethernet
Fibre Channel	■ Hitachi 8Gb/sec 2-port Fibre Channel mezzanine card (FIVE-EX based Fibre Channel to PCIe HBA)

Network Architecture

Hitachi Unified Compute Platform used in this reference architecture provides two physical networks, both with redundant switching connectivity. In the Brocade Networking model, UCP utilizes Brocade Virtual Cluster Switching (VCS) Technology to eliminate spanning tree, and optimize for traffic. All VDX switches in a single UCP instance are configured as VCS Fabric Cluster mode, and form a single VCS fabric. UCP also utilizes Virtual LAG (vLAG) technology which provides multi-path networking between VCS fabric and the customer network. A high level architecture diagram for the base rack is shown in Figure 2.

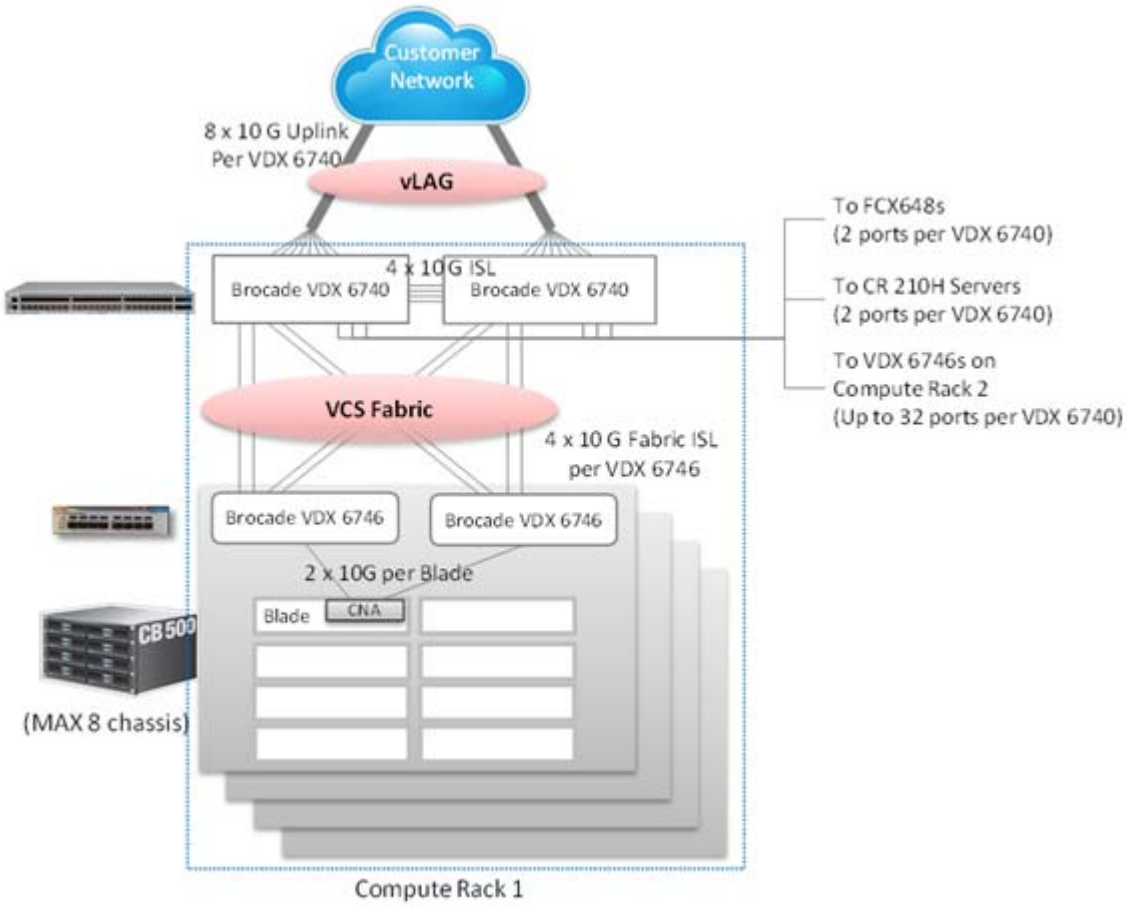

Figure 2

The network architecture can be divided as follows:

- Management Network
- Compute Network

Management Network

1 Gb/sec Management Network using two Brocade FCX648 LAN switches consists of the following:

- There are two 1 G connections from each CR210 H server (one to each FCX648 switch).
- There are two connections from each management module of the CB 500 chassis (one to each FCX648 switch).

The physical 1 Gb/sec Management network provides two uplinks into the customer management Ethernet network for redundancy. By default, the physical management network is divided into two VLANs for network connectivity:

- Management VLAN: Used for all UCP management traffic including Element Management, VMware vSphere management, and Service Provisioning.
- Migration VLAN: Used for vMotion traffic between two CR 210H management servers.

This is shown in Table 4 (actual VLAN IDs and subnets are customizable to fit customer environments).

Table 4. Default Unified Compute Platform Management Network VLANs

Detail	Management VM Network	Migration Network
Default VLAN	101	102
Default subnet	192.168.101.xx	192.168.102.xx
Default subnet mask	255.255.255.0	255.255.255.0
Default gateway	192.168.101.1	192.168.102.1

Unified Compute Platform predefines the physical LAN architecture and therefore no customer-specific customization is required.

Figure 3 shows an overview of the management network connectivity for the single-rack Unified Compute Platform solution used for engineering validation, including Hitachi Compute Rack 210 rack mount servers, which host Unified Compute Platform management software (see Physical Server Configuration on page 8 for details). Note that all connectivity fabrics for both compute and storage components are redundantly connected.

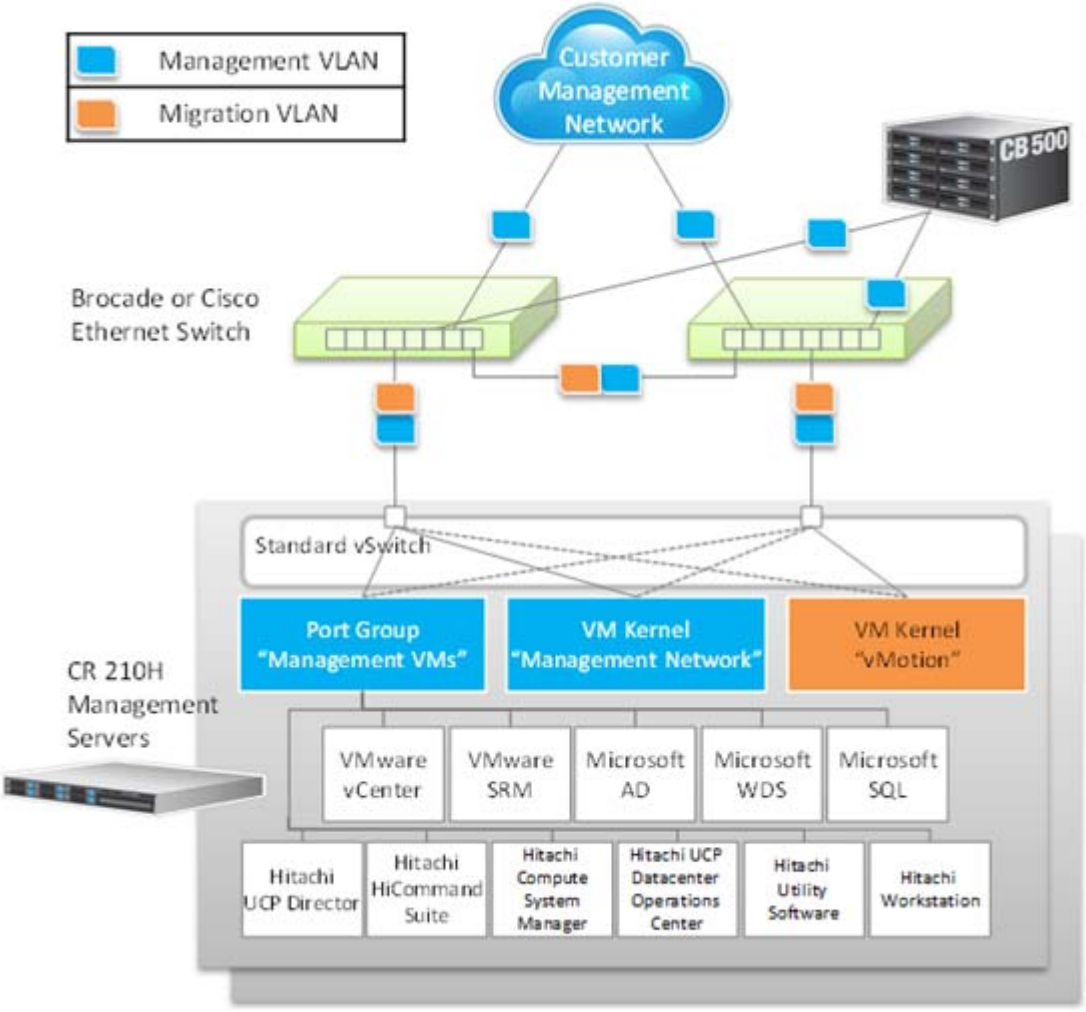

Figure 3

Compute Network

Compute network is a 10 Gb/sec network using two Brocade 6740 LAN switches per rack and two Brocade 6476 Ethernet switches per CB 500 chassis. There are four 10Gb/sec connections from the first Brocade 6476 switch to the first Brocade 6740 switch, and four 10Gb/sec connections from the second Brocade 6476 switch to the second Brocade 6740 switch. Therefore, 80 Gb/sec network bandwidth is available per chassis. UCP supports Standard vSwitch and vSphere Distributed Switch (VDS). For engineering validation of this solution, ESXi host was configured with one Standard vSwitch with two physical adapters (vmnic) for redundancy. Users can configure an arbitrary Port Group or an Uplink Port Group with a VLAN range for the user's virtual machine and other traffic. Additionally, there are two connections from each CR 210H (one to each Brocade 6740 switch). The data traffic of the service virtual machines hosted on CR210H servers goes through the compute network. The compute network architecture is shown in Figure 4.

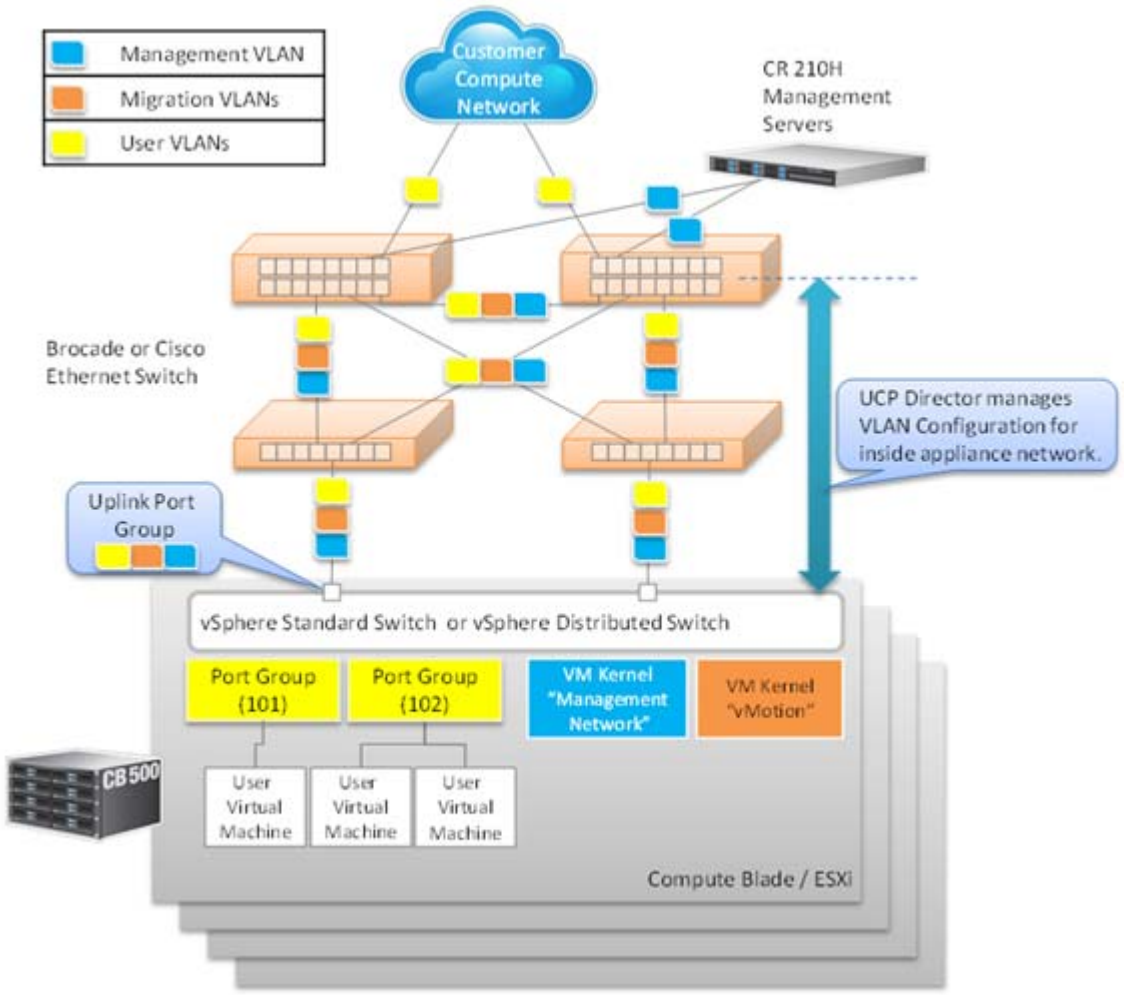

Figure 4

The physical 10 Gb/sec Ethernet compute network can be divided into multiple VLANs as required. The default Unified Compute Platform compute network VLANs are shown in Table 5. As with the management VLANs, the actual VLAN IDs and subnets can be customized to fit customer environments. Up to sixteen 10 Gb/sec Ethernet uplinks (eight per switch) are available for connection to the customer Ethernet network to deliver high bandwidth, redundant connectivity.

Access to specific Compute Networks (VLANs) from VMware clusters or VMware VMs is automatically configured via Hitachi Unified Compute Platform Director, which then configures VMware Distributed Switches as required.

Table 5. Default Unified Compute Platform Compute Network VLANs

Detail	Compute Network	Compute Network	Compute Network
Default VLAN	103	104	105
Default subnet	192.168.103.xx	192.168.104.xx	192.168.105.xx
Default subnet mask	255.255.255.0	255.255.255.0	255.255.255.0
Default gateway	192.168.103.1	192.168.104.1	192.168.105.1

Unified Compute Platform predefines the physical LAN architecture and therefore no customer-specific customization is required.

SUSE Linux Enterprise Server (SLES) Configuration

This reference architecture uses the SLES for SAP Applications (SUSE Linux Enterprise Server for SAP Applications) version 11 SP3 configured according to SAP recommendations.

VMware Configuration

The VMware configuration of this reference architecture follows the official SAP on VMware recommendations given in SAP Note 1122388 - Linux: VMware vSphere Configuration Guidelines as follows:

- Do not over commit memory
- Installing VMware Tools is mandatory

For storage connectivity, this reference architecture uses VMware physical RDM (Raw Device Mapping) in order to achieve the following:

- Allow easy movement between VMware, physical, and other Hypervisors (for example HVM/LPAR)
- Enable native backend storage capabilities such as cloning and snapshots for individual SAP Systems

Hitachi Unified Compute Platform uses vSphere Distributed Switches (VDS) for configuring network connectivity. The actual configuration of the VDS and Distributed Ports Groups (dvPG) is taken care of automatically by Hitachi Unified Compute Platform Director.

SAN Architecture

Hitachi Unified Compute Platform offers two SAN connectivity alternatives, one uses Brocade switches and the other uses Cisco switches. Although the focus of this reference architecture is on using Brocade components, the virtualized nature of the reference architecture equally applies to a Cisco-based Unified Compute Platform.

The specific Hitachi Unified Compute Platform used for this reference architecture uses two redundant SAN fabrics. There is one 8 Gb/sec Brocade 6510 Fibre Channel switch and one 8 Gb/sec Brocade 5460 Fibre Channel switch in every blade chassis. Each 520H B3 server blade has two redundant 8 Gb/sec Fibre Channel connections, one in each SAN fabric (see Figure 5 for details).

Host Groups are used to control access to storage LUNs within the Unified Compute Platform SAN fabrics. Hitachi Unified Compute Platform Director automatically creates and assigns appropriate Host Groups to LUNs when creating/adding/removing LUNs via Unified Compute Platform Director.

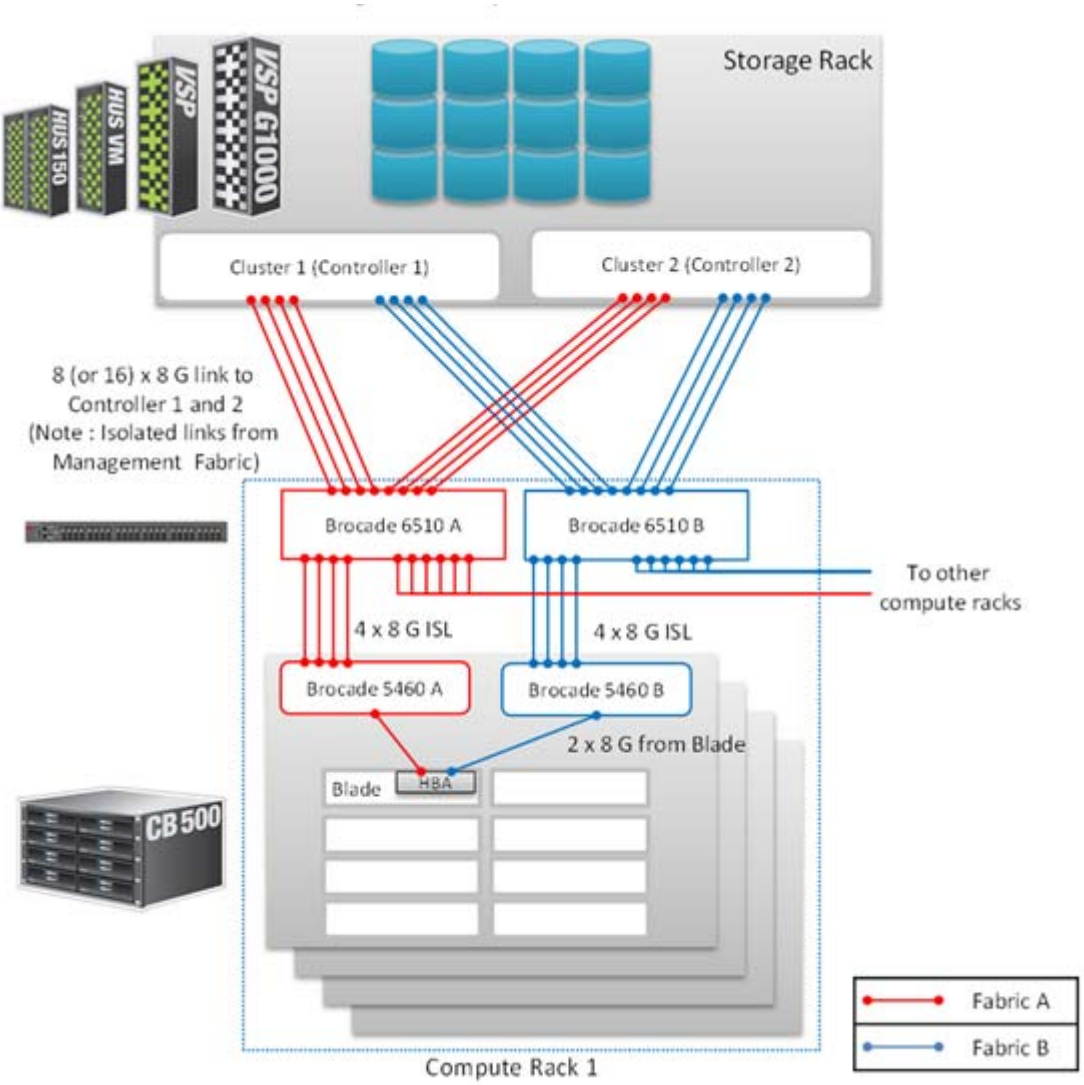

Figure 5

Storage Architecture

The storage requirements for this reference architecture are the following:

- 1 × RAID-6 (6D+2P) Parity Group mapped to 1 HDP Pool with one 3 TB LDEV for HANA logs
- 3 × RAID-6 (6D+2P) Parity Groups mapped to 1 HDP Pool with six 1.6 TB LDEVs for HANA data
- 1 × RAID-6 (6D+2P) Parity Group mapped to 1 HDP Pool with one 3 TB LDEV for:
  - The operating system LUN for the HANA VM
  - The operating system LUN for the Business Suite VM
  - The /hana/shared partition
  - The /sapmnt partition for Business Suite
  - The /usr/sap/SID partition for Business Suite
  - The /usr/sap/trans partition for the Business Suite transport files

The reference architecture distinguishes between three groups of storage based on purpose:

- UCP Management Storage
- UCP Compute: SAP System Independent Storage
- UCP Compute: SAP System Specific Storage

The two “Unified Compute Platform Compute” groups are customer definable. In order to ensure the flexibility of Unified Compute Platform, this reference architecture recommends using a minimum variation of basic storage unit types (parity groups). Since Hitachi Unified Compute Platform leverages Hitachi Dynamic Provisioning (HDP) for storage management, all parity groups can be flexibly combined into Dynamic Provisioning pools to evenly spread data and workload across all parity groups within the Dynamic Provisioning pool.

Unified Compute Platform Management Storage

Hitachi Unified Compute Platform reserves a single RAID-6 (6D+2P) parity group to host Unified Compute Platform Management software. This storage pool is exclusively for Unified Compute Platform Management purposes. See Table 6 for details.

Table 6. Unified Compute Platform Management Storage

Parity Group	LDEV	HDP Pool	RAID Level	Drive Size	Drive Type/ Speed	Usable Capacity per Parity Group	Total Usable Capacity	Purpose
1-1	00:00:00 00:00:01	UCP-Mgmt.	RAID-6 (6D+2P)	600 GB	SAS 10k	3.6 TB	3.6 TB	Unified Compute Platform Management Server

Unified Compute Platform: SAP System Independent Storage

Within this reference architecture, a single VMware Datastore is created to hold each of the following for the SAP HANA VM:

- HANA data volume
- HANA log volume
- HANA SUSE Linux Enterprise Server (SLES) OS
- HANA shared file system

Similarly, one datastore is created for each of the following for the Business Suite VM:

- Business Suite OS
- /sapmnt directory
- /usr/sap/SID directory
- /usr/sap/trans directory for Business Suite transports

Unified Compute Platform: SAP System Specific Storage

This reference architecture uses SAS RAID-6 (6D+2P) HDP Pools for storage LUNs. Figure 6, Figure 7 and Figure 8 show a graphical overview of the storage layout.

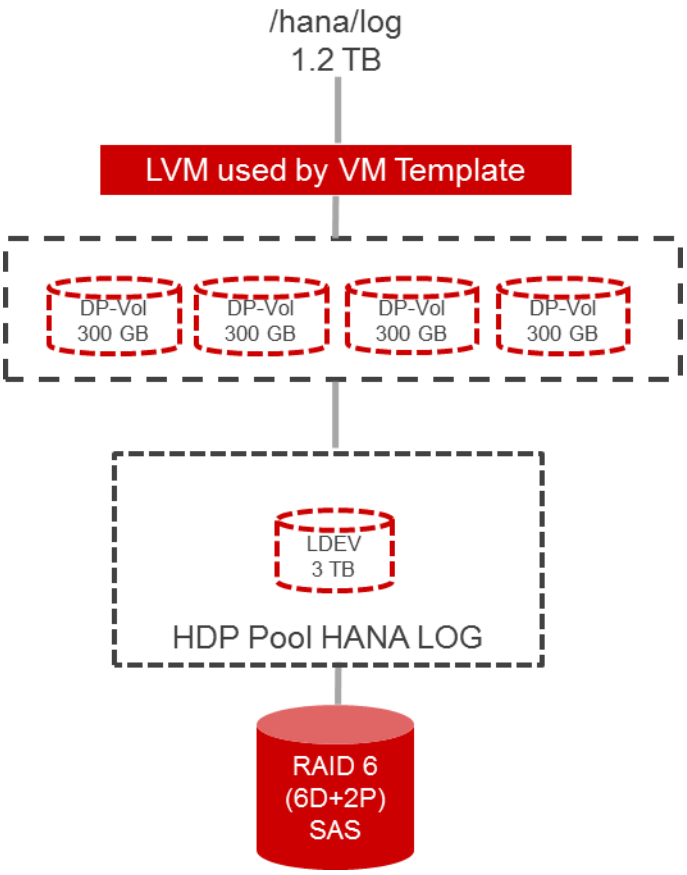

Figure 6

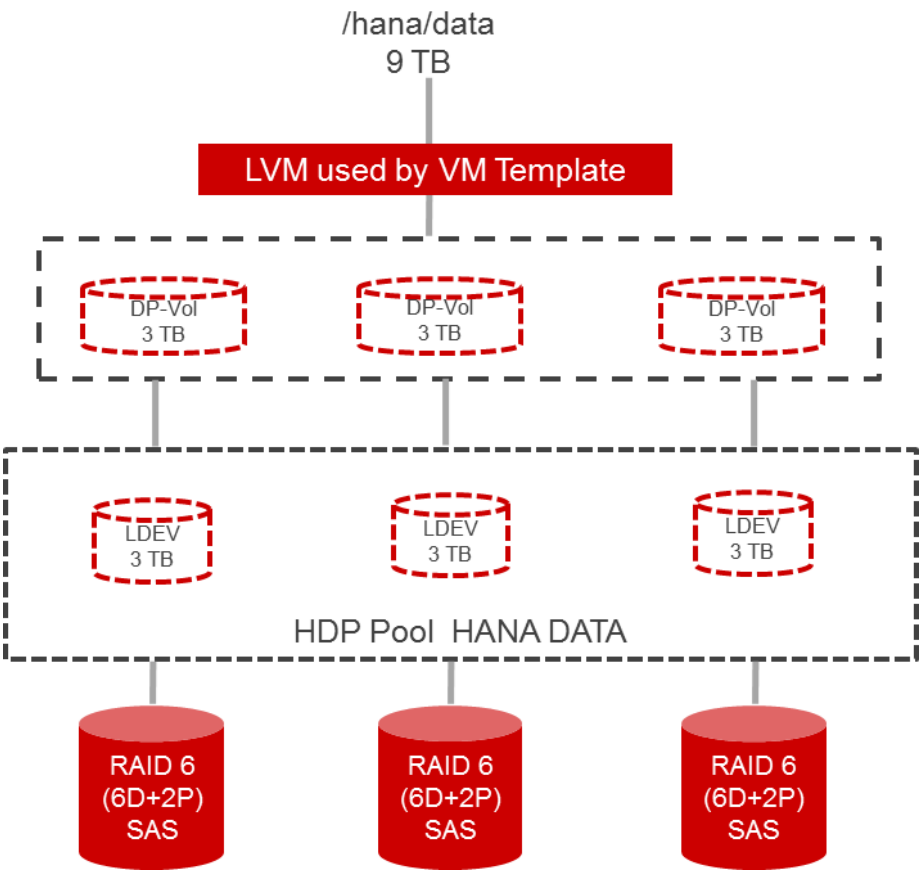
/hana/data
9 TB
LVM used by VM Template
DP-Vol
3 TB
DP-Vol
3 TB
DP-Vol
3 TB
LDEV
3 TB
LDEV
3 TB
LDEV
3 TB
HDP Pool HANA DATA
RAID 6
(6D+2P)
SAS
RAID 6
(6D+2P)
SAS
RAID 6
(6D+2P)
SAS

Figure 7

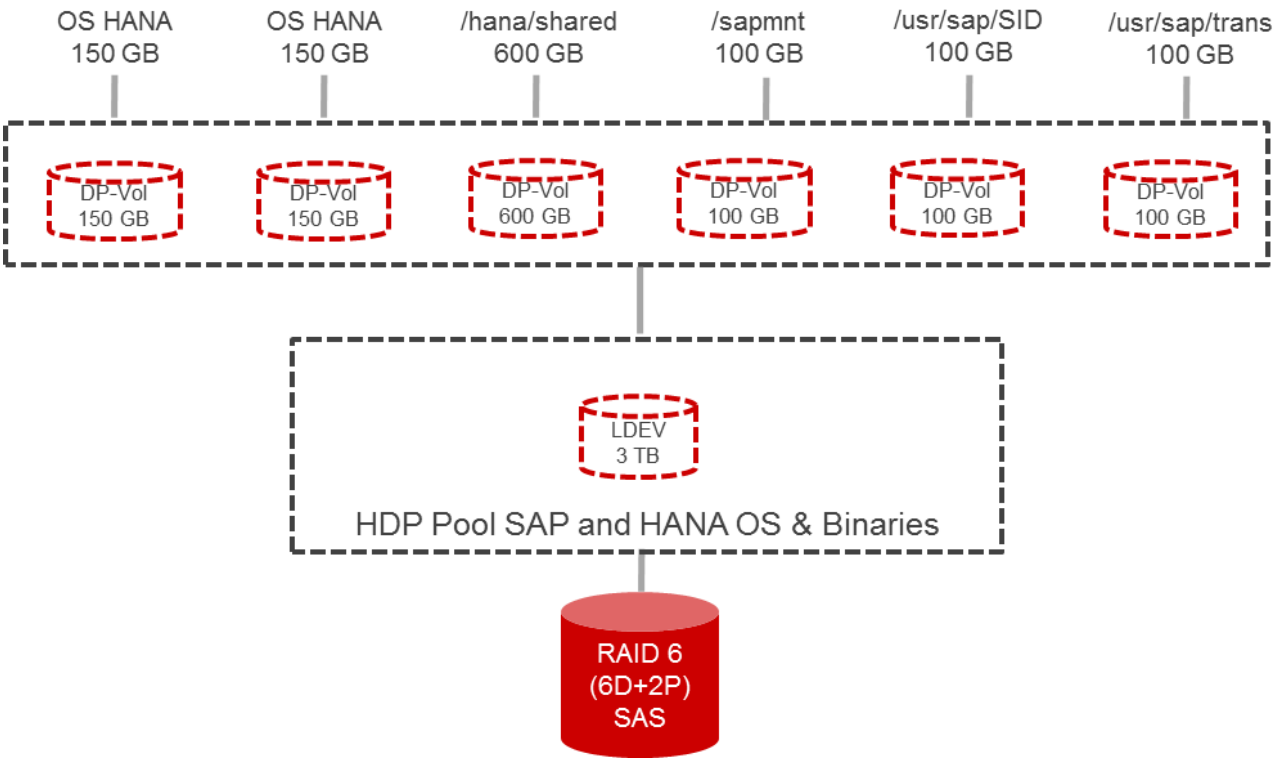
OS HANA
150 GB
OS HANA
150 GB
/hana/shared
600 GB
/sapmnt
100 GB
/usr/sap/SID
100 GB
/usr/sap/trans
100 GB
DP-Vol
150 GB
DP-Vol
150 GB
DP-Vol
600 GB
DP-Vol
100 GB
DP-Vol
100 GB
DP-Vol
100 GB
LDEV
3 TB
HDP Pool SAP and HANA OS & Binaries
RAID 6
(6D+2P)
SAS

Figure 8

Of particular interest is the use of Hitachi Dynamic Provisioning to stripe sapdata file systems across parity groups and not the Linux LVM (Logical Volume Manager), because Dynamic Provisioning dynamically rebalances DP-Vols across new parity groups as they are added. For an example storage layout, see Table 7 and Table 8.

Table 7. Example Parity Group and Hitachi Dynamic Provisioning Pool Configuration for SAP Business Suite

HDP Pool ID	Parity Group RAID Level and Disks	LDEV ID	LUN ID	LDEV Size	Description	Datastore Name
0	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:00:3F	1	100 GB	/sapmnt	sapmnt
		00:00:3E	2	160 GB	ESXi_operating_system_ERP_Node	ossap
		00:00:40	3	100 GB	/usr/sap	usrsap
		00:00:41	4	100 GB	saptrans	transdir
		00:00:36	5	Memory Size	SAP_shared	hanashared

Table 8. Example Parity Group and Hitachi Dynamic Provisioning Pool Configuration for SAP HANA

HDP Pool ID	Parity Group RAID Level and Disks	LDEV ID	LUN ID	LDEV Size	Description	Datastore Name
1	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:35	1	160 GB	ESXi_operating_system_HANA_Node	OSHANA
		00:36	2	240 GB	HANA_shared	hanashared
2	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:37	3	300 GB	LOG_1	hanalog1
		00:38	4	300 GB	LOG_2	hanalog2
		00:39	5	300 GB	LOG_3	hanalog3
		00:3A	6	300 GB	LOG_4	hanalog4
3	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	00:3B	7	3000 GB	DATA_1	hanadata1
		00:3C	9	3000 GB	DATA_2	hanadata2
		00:3D	8	3000 GB	DATA_3	hanadata3

Engineering Validation

In order to validate this reference architecture, a standard SAP benchmark was executed on the system.

At this time, the SAP SD benchmark for HANA (which would be the ideal benchmark system to use for this Business Suite setup) was not available. It was therefore decided to use the BW-EML benchmark, which provides detailed throughput information for the underlying NetWeaver architecture. This benchmark system gives full-stack NetWeaver, SAP, and HANA benchmarking scores, where other benchmarks would not provide such a full report.

The benchmark was set up according to the SAP Implementation Guide: Benchmark Tools and Appl. Benchmarks document. Minimal tuning was done to the Business Suite system profile, according to benchmarking best practices, to obtain the best performance from the system.

The benchmark run consists of the EML users "logging-on" to the system, and running predefined (but somewhat randomized) queries on the BW subsystem. Once the defined number of users have logged on, a "high-load" phase, with a duration of one hour, begins. During this phase, the EML users continue to solicit the system, and several BW data loads are run at the same time. The BW data loads are timed to take exactly one hour. In order for the benchmark score to be considered valid, the EML users should receive no "system overloaded" messages, and the data loads must finish within the "high-load" phase.

The BW-EML benchmark uses several different concepts that are all linked together:

- Number of dialog (DIA) processes - the quantity of available dialog processes in the BW system. These are used by the EML users performing their query runs. Each dialog process uses memory, and will add some load to both the NetWeaver server and the underlying database if used
- Number of batch (BTC) processes - the number of batch (background) processes in the system. These are used by the data load process. Again, each process uses memory, and adds load to the system and database when used
- Number of EML users - the number of EML users who will log on to the system and perform work. The more EML users, the more work performed, and hence the higher score. However, if there are too many EML users, then the system will become overloaded, and either swap, or return errors, rendering the benchmark invalid
- Number of EML loops - the number of times each individual EML user will perform navigation in the BW queries. This parameter will only influence the length of time each EML user will work, in order to respect the requirement that the system be fully utilized during the high-load phase.

The parameters used during the engineering validation of this reference architecture, and the resulting BW-EML benchmark scores obtained, are detailed in Table 9. There were 50 dialog processes and 10 batch processes used to run the benchmark tests.

Table 9. BW-EML Benchmark Score Details

Number of EML users	Number of EML Loops	HANA CPU Usage	HANA Memory Usage	NetWeaver CPU Usage	NetWeaver Memory Usage	Benchmark Score
190	10	25%-35%	239GB max	35%-55%	55GB Max	62154
250	10	30%-45%	239GB max	40%-60%	62GB Max	81788
250	10	30%-45%	239GB max	40%-60%	62GB Max	81878
350	10	50%-70%	223GB Max	45%-70%	73GB Max	97684

Conclusion

This reference architecture gives guidance on how the converged Hitachi Unified Compute Platform for VMware vSphere can host a complete SAP Business Suite Landscape. By using a straightforward SAP 3-Tier Building Block approach, the reference architecture proposes a solution that is both flexible and easy to manage in daily operations.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.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-412-00 August 2015.