TECHNICAL PAPER

HITACHI
Inspire the Next

Deploy Oracle Multitenant on Hitachi Unified Compute Platform for Oracle Database

Reference Architecture Guide

By Nathan Tran

April 2, 2015

ORACLE Diamond Partner

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 Oracle Multitenant on Hitachi Unified Compute Platform for Oracle Database

Reference Architecture Guide

This reference architecture guide describes the design and implementation of a solution for Oracle Database 12c environment with Oracle Multitenant on Hitachi Unified Compute Platform for Oracle Database.

Hitachi Unified Compute Platform is a family of completely integrated and flexible solutions. Each solution is configured for immediate deployment to run top-tier infrastructure applications without over-purchasing or provisioning unnecessary equipment. Each custom-built solution has its entire solution stack certified. There are no compatibility issues.

Oracle Multitenant, an Oracle Database 12c Enterprise Edition option, introduces an architecture to consolidate multiple databases without changing their applications. Manage many databases as one, yet retaining the isolation and resource prioritization of separate databases.

Supporting rapid provisioning and upgrades, Oracle Multitenant fully complements other Oracle Database options. These other options include Oracle Real Application Clusters and Oracle Active Data Guard.

These are benefits of using Oracle Database 12c with Oracle Multitenant:

- **Cost reduction**

  Consolidating hardware and sharing database memory and files reduces costs for hardware, storage, availability, and labor. For example, 100 pluggable databases on a single server can share one database instance and one set of database files to require less hardware and fewer personnel.

- **Easier and more rapid movement of data and code**

  By design, you can quickly plug a pluggable database into a container database, unplug the pluggable database from the container database, and then plug this pluggable database into a different container database. The implementation technique for plugging and unplugging is similar to the transportable tablespace technique.

- **Ease of performance tuning**

  Performance metrics can be collected at the container database level. Tuning can be set that could benefit all Oracle pluggable databases. It is easier to size one system global area than 100 system global areas.

- **Ease of management**

  Support, manage, and backup of many pluggable databases within a single container database.

Figure 1 illustrates how Hitachi Unified Compute Platform is used for Oracle Multitenant deployment.

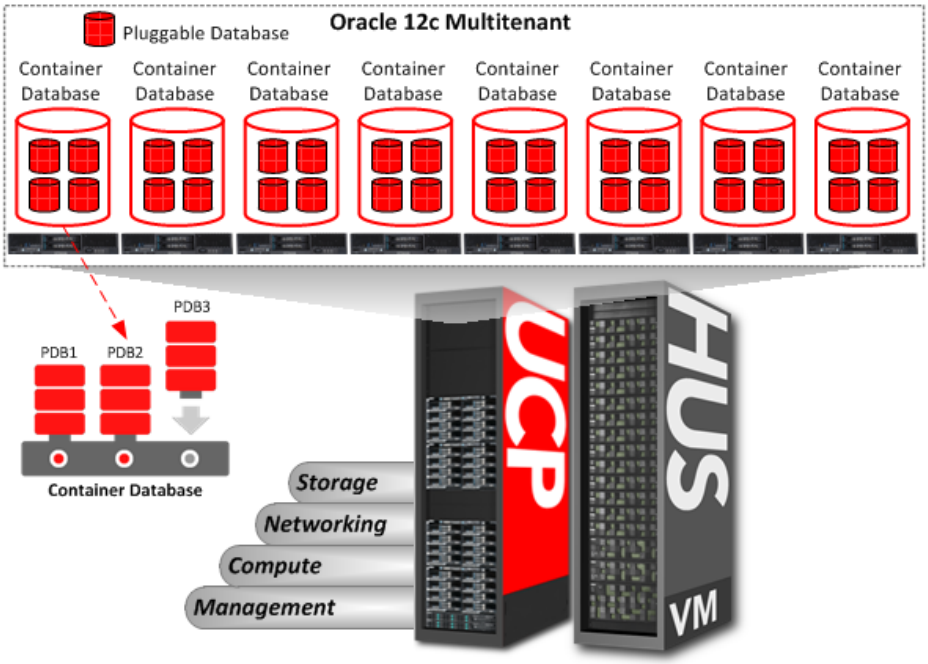

Figure 1

In addition to the general benefits of using Oracle Database 12, the benefits of this solution include the following:

- Faster deployment
- Reduced risk
- Predictable performance
- Ability to scale out
- Lower total cost of ownership

This reference architecture guide is for IT administrators, architects, database administrators, chief information officers, chief technology officers, or business intelligence users with an interest or responsibility for planning an Oracle Database 12 solution using Oracle Multitenant.

To use this reference architecture guide, you need familiarity with the following:

- Hitachi Unified Compute Platform
- Hitachi Unified Storage VM
- Hitachi Compute Blade 500
- Storage area networks
- Oracle Database 12c
- Red Hat Enterprise Linux

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

Use of Hitachi Unified Compute Platform for Oracle Database with Oracle Database 12c using Oracle Multitenant provides a tightly integrated hardware stack that is highly available and scalable. It meets the needs of a shared database environment, taking advantage of flexible resources for sharing and cost savings.

The following components create this Unified Compute Platform for Oracle Database solution:

- **Hitachi Compute Blade 500** — Enterprise-class server platform, containing internal Fibre Channel and network switch modules, that provides dense compute resources and high I/O throughput
- **Hitachi Unified Storage VM** — Midrange storage system with dual active-active controllers for data redundancy that is designed for all data types and flexibly adapts for performance, capacity, and multi-vendor storage
- **Hitachi Compute Rack 210H** — 1U rack-mountable servers for management environment
- **Hitachi Unified Compute Platform Director** — Provide management and provisioning functions for Hitachi Unified Compute Platform
- **Brocade Fibre Channel Switch Modules** — Provides redundant Fibre Channel connectivity for storage access
- **Brocade Ethernet Switch Modules** — Provides redundant network connectivity for management and client access
- **VMware vSphere 5** — Enterprise virtualization solution to create a dynamic and flexible data center with integrated management and reporting capability for a high level of server and service uptime. This provides management and provisioning function as part of Hitachi Unified Compute Platform for Oracle Database.

Figure 2 on page 5 shows the high-level diagram of this reference architecture.

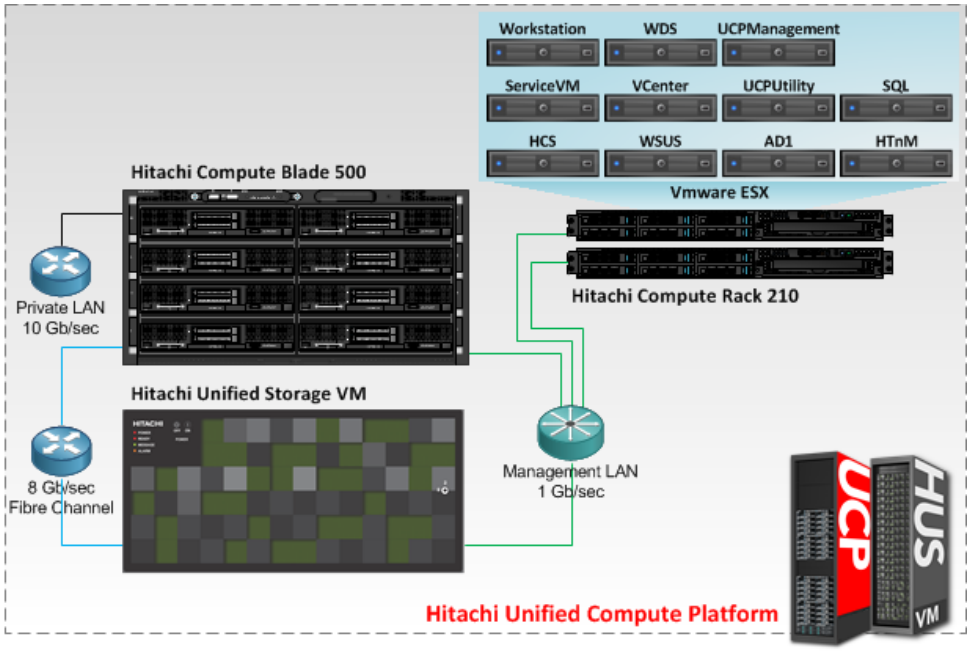
Workstation
WDS
UCPManagement
ServiceVM
VCenter
UCPUtility
SQL
HCS
WSUS
AD1
HTnM
Vmware ESX
Hitachi Compute Blade 500
Hitachi Compute Rack 210
Private LAN
10 Gb/sec
Hitachi Unified Storage VM
8 Gb/sec
Fibre Channel
Management LAN
1 Gb/sec
UCP
HUS
VM
Hitachi Unified Compute Platform

Figure 2

Key Solution Components

These are the key hardware and software components used to deploy this Hitachi Unified Compute Platform for Oracle Database with Oracle Multitenant reference solution.

Hardware Components

Table 1 lists the detailed information about the hardware components used in the Hitachi Data Systems lab to validate this solution.

Table 1. Hardware Components

Hardware	*Description*	*Version*	*Quantity*

Hitachi Unified Storage VM	■ 4 × 8 Gb/sec Fibre Channel front-end connectivity modules
	■ 4 × 6 Gb/sec SAS back-end links back-end connectivity modules
	■ 2 MP storage blades, each with two 8-core Intel Xeon processors, 2.33 GHz
	■ 128 GB cache memory
	■ 34 × 200 GB SSD disks
	  ■ 2 hot spares
	■ 146 × 600 GB 10k RPM SAS disks, 2.5 inch SFF
	  ■ 2 hot spares
		73-03-04-00	1

Hitachi Compute Blade 500	■ Server blade chassis ■ 2 Brocade 5460 Fibre Channel switch modules, each with 6 × 8 Gb/sec uplink ports ■ 2 Brocade VDX 6746 Ethernet switch modules, each with 8 × 10 Gb/sec uplink ports ■ 2 management modules ■ 6 cooling fan modules ■ 4 power supply modules ■ Chassis supports up to eight 520H B2 server blades, although this solution only uses two 520H B2 server blades	Firmware: 04-15 BMC: 04-11 EFI: 10-23	1

Table 1. Hardware Components (Continued)

Hardware	Description	Version	Quantity
520H B2 server blade	■ Half blade ■ 2 × 12-core Intel Xeon E5-2697 processor, 2.70 GHz ■ 224 GB RAM ■ 14 × 16 GB DIMMs	Firmware: 04-15 BMC: 04-11 EFI: 10-23	2
Hitachi Compute Rack 210H	■ Dual 8-core Intel Xeon E5-2670 processor, 2.60 GHz CPU ■ 192 GB RAM ■ 2 × 300 GB 10k RPM SAS drives	01-09-00	2
Brocade 6510 Fibre Channel Switch	■ 48-port enterprise-class Fibre Channel switch in a 1U enclosure ■ 4/8/16 Gb/sec Fibre Channel ports ■ 24, 36, and 48-port configurations	7.2.0a	2
Brocade VDX 6720 Switch	■ 10 Gb/sec LAN switch ■ Top-of-Rack 60 ports switch ■ Fibre Channel over Ethernet (FCoE), iSCSI, and NAS support	3.0.1aa	2
Brocade FCX 648 Switch	■ Enterprise-class stackable Layer 2/3 edge switches ■ 10/100/1000 Mb/sec RJ-45 ports ■ 4 Optional 10 Gb/sec SFP+ ports	7.4.0	2

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 migration 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 more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions

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.

The Hitachi Compute Blade 500 chassis contains internal Fibre Channel and network switches for the high availability requirements for this solution.

As part of this Hitachi Unified Compute Platform for Oracle Database solution, Hitachi Compute Blade 500 provides a reliable and powerful compute environment for Oracle Container Database.

Hitachi Compute Rack 210H

Hitachi Compute Rack 210H is midrange rack mountable server platform, providing advanced systems management and redundancy options. It is data center friendly, with a 1U footprint, while delivering the performance that is required to meet enterprise-level challenges.

The benefits of Hitachi Compute Rack 210H are the following:

- Web-based management interface
- RAID level configuration, with up to six 2.5 inch internal drives
- Sustainable power-saving capabilities
- Configuration flexibility to meet business needs
- Dense 1U rack mountable design

The two Hitachi Compute Rack 210H servers are used for a high availability management component for Hitachi Unified Compute Platform for Oracle Database.

Software Components

Table 2 describes the software components deployed for this reference architecture on Hitachi Unified Compute Platform for Oracle Database.

Table 2. Software Components

Software		*Version*	*Function*
Hitachi Storage Navigator		Microcode Dependent	Storage management software
Hitachi Dynamic Provisioning		Microcode Dependent	Storage side feature
Hitachi Command Suite		V7.6	Unified management
Red Hat Enterprise Linux		6.4	Operating system
Oracle Database		12c release 1	Database engine
Unified Compute Platform Management	Hitachi Unified Compute Platform Director	3.02	Unified Compute Platform management
	VMware ESX Server	5.5.0	VMware host
	VMware vSphere	5.5.0	VMware management
	Microsoft® Active Directory®	Component of Microsoft® Windows Server® 2012	Identity server
	Microsoft Windows Server	2012 Enterprise x64	Operating system

Hitachi Storage Navigator

Hitachi Storage Navigator enables essential management and optimization functions. Using Java agents, Storage Navigator runs on most browsers. A command line interface is available.

Use Storage Navigator for the following:

- Pool creation and expansion
- LUN creation and expansion
- Online microcode updates and other system maintenance functions
- Performance metrics

You need Storage Navigator to take advantage of the full features of Hitachi Unified Storage VM.

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 specified logical size. Up to 60 TB can be created against it without initially allocating any physical space.

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. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Use Hitachi Dynamic Provisioning to create many disk pools for placing the Oracle container database files with dedicated disk resource to ensure maximum performance during peak operation.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the thin provisioning and wide striping features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42MB page level.

Breaking the volume into pages, Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time. It makes decisions on moving data between the available storage tiers based on actual use. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Hitachi Dynamic Tiering on Hitachi Virtual Storage Platform allows a single pool to contain tiers made up of differently-arranged RAID groups using any type of disk. It manages data migration between the various tiers within a pool automatically. This eliminates most user management of storage tiers within a storage system, and maintains peak performance under dynamic conditions without storage administrator intervention.

This solution uses Hitachi Dynamic Tiering to automate the placement of frequently use Oracle 12c pluggable database (PDB) on SSD disks to achieving the fastest access of data.

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.

Red Hat Enterprise Linux 6.4

Using the stability and flexibility of Red Hat Enterprise Linux, reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability on military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shirt and technologies evolve.

Oracle Database 12c

Oracle Database is optimized for use with Oracle products. It uses Oracle Database Automatic Storage Management, combining the features of a volume manager and an application-optimized file system for database files. ASM is part of the grid infrastructure component in Oracle Database.

VMware vSphere 5

VMware vSphere 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 has the following components:

- **ESXi** — 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.

The Hitachi Unified Compute Platform management environment for this solution is provided by two Hitachi Compute Rack 210H servers in an active-active VMware vSphere cluster configuration.

Solution Design

This is the detailed description of the reference architecture environment implementing an Oracle 12c multitenant environment using Hitachi Unified Compute Platform.

- **Compute Environment** — This solution uses 520B H2 half-size blades in the Hitachi Compute Blade 500 chassis for Oracle 12c container database
- **Storage System** — Hitachi Unified Storage VM with Hitachi Dynamic Tiering to create a multi tiers disk pool for Oracle data.
- **SAN Fabric** — The Fibre Channel fabric for this solution consist of two Brocade 5460 Fibre Channel switch modules and two Brocade 6510 Fibre Channel switches.

Compute Environment

Hitachi Compute Blade 500 chassis provides a highly redundant and flexible compute environment for Oracle Multitenant deployment.

Hitachi Compute Blade 500 supports up to eight blade servers using half size blades. Intermixing of half server blades and full server blades is supported by Hitachi Compute Blade 500 to achieve different node count configurations. It is out of scope for this reference architecture to consider mixing half size server blades and full size server blades.

This reference architecture focuses only on a half-size blade configuration using 520H B2 server blades to provide up to eight separate Oracle container databases.

Figure 3 shows how Hitachi Compute Rack 500 supports eight server blades using half size blades within the same chassis.

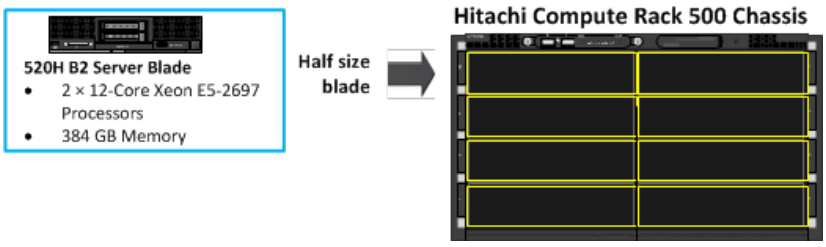

Figure 3

Provisioning of the server blades uses service template for Red Hat Enterprise Linux on Hitachi Unified Compute Platform Director software. Each server blade has a unique server Identity that is used by the Unified Compute Platform Director software for bare-metal operating system deployment. Figure 4 illustrates how to use the service template to provision the operating system onto the server blade.

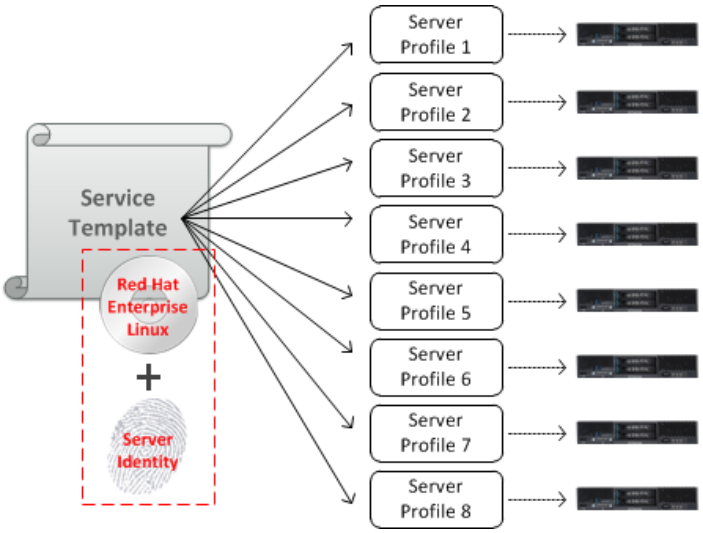

Figure 4

Management, provisioning, and orchestration services for this solution are provided by two Hitachi Compute Rack 210H servers.

Storage Architecture

The storage implementation for this reference architecture takes into consideration recommended practices by Hitachi Data Systems and Oracle for the deployment of a storage design for a multitenant database solution.

This solution uses Hitachi Unified Storage VM for the underlining storage. It provides fast and reliable storage resources for Oracle database file placement.

Hitachi Dynamic Tiering automates data placement on disk for the Oracle data pool. This ensures that high access data are placed on the tier-1 storage, consisting of SSD disks.

Create the Oracle log and Oracle fast recovery area from two separate dynamic provisioning pools. Store the log in a dedicated pool to achieve optimal performance for Oracle log access.

Hitachi Unified Storage VM requires the flash acceleration and flash optimization license options.

Figure 5 shows the high-level storage design for this solution.

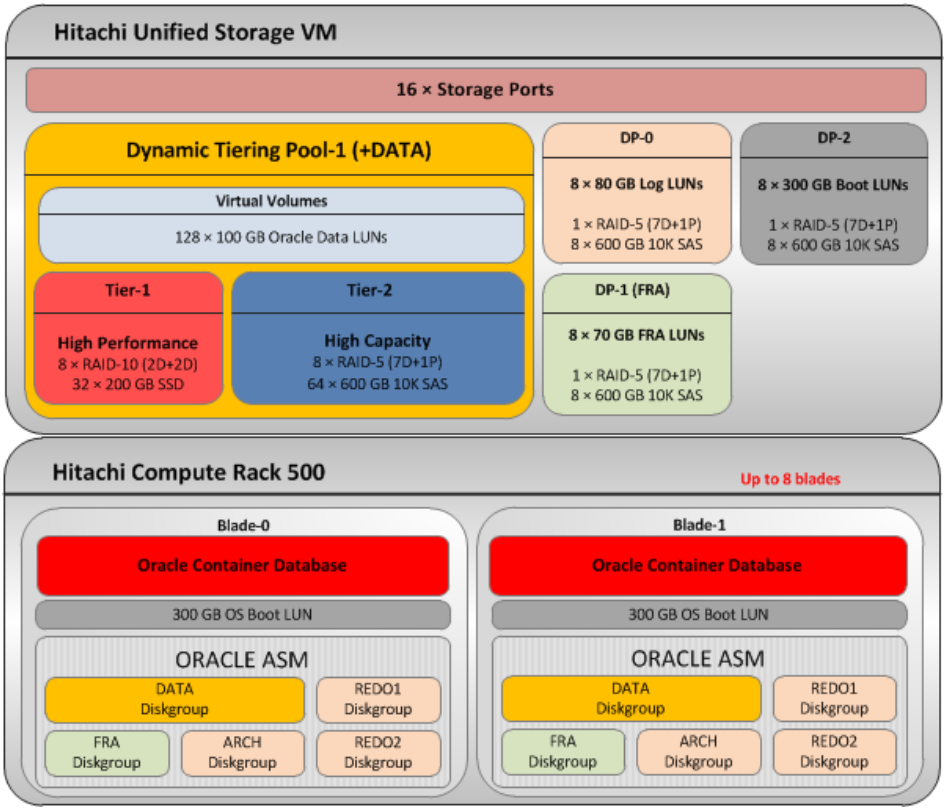

Figure 5

RAID Configuration

The storage implementation for this solution uses Hitachi Dynamic Tiering and Hitachi Dynamic Provisioning to create pools on Hitachi Unified Storage VM.

Place Oracle container database files in a dynamic tiering pool consisting of two disk tiers:

- T1 uses SSD disks for high performance data access
- T2 uses SAS drives for high capacity data storage

Hitachi Dynamic Tiering automates the task of placing frequently use data into the T1 disk tier for fast access to data while keeping other data on T2. Data movement is in pages of 42 MB.

There are three dynamic provisioning pools for the following:

- Oracle logs
- Oracle fast recovery area
- Hitachi Unified Compute platform management file stores

The three dynamic provisioning pools provide dedicated disk for specific needs without affecting performance of critical functions, such as Oracle archive logs access.

Table 3 describes the different RAID groups in this reference architecture.

Table 3. RAID Configuration

Purpose	*RAID Group Configuration*	*PG*	*Total PG*
Oracle data Tier-1 disks	RAID-10 (2D+2D)	1-8	8
Oracle data Tier-2 disks	RAID-5 (7D+1P)	9-23	15
Oracle logs disks	RAID-5 (7D+1P)	24	1
Oracle fast recovery area disks	RAID-5 (7D+1P)	25-26	2
Unified Compute Platform management disks	RAID-6 (6D+2P)	27	1

Dynamic Tiering Pool

Use Hitachi Dynamic Tiering to create a dynamic tiering pool for container database file placement. It automates placing high access database files on Tier-1 disks. These are SSD for faster access to data.

Figure 6 shows the components that make up the dynamic tiering pool.

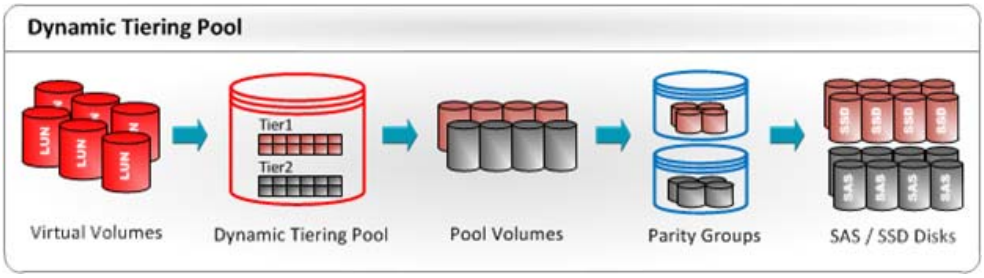

Figure 6

Table 4 describes the required dynamic tiering pool and the disk tiers that for this solution.

Table 4. Dynamic Tiering Pool

Pool ID	*Tier*	*RAID Level*	*Disk Type*	*Capacity*	*Disk Count*	*Purpose*
DT-1	T1	RAID 10 (2D+2D)	200 GB SSD	2.86 TB	32	High performance disk tier
	T2	RAID 5 (7D+1P)	600 GB 10k RPM SAS	29.33 TB	64	High capacity disk tier

Dynamic Provisioning Pools

This describes the dynamic provisioning pool this solution uses.

Figure 7 shows the relationship between the disk and the dynamic provisioning pool, and how to create volumes.

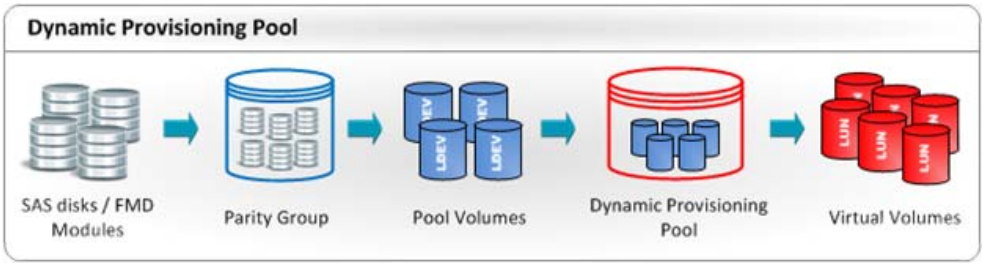

Figure 7

Table 5 lists the dynamic provisioning pools used in this solution.

Table 5. Dynamic Provisioning Pools

Dynamic Provisioning Pool	*RAID Level*	*Disk Count*	*Capacity*	*Purpose*
DP-1	RAID-6 (6D+2P)	8	3.14 TB	Hitachi Unified Compute Platform Management Pool
DP-2	RAID-5 (7D+1P)	8	3.66 TB	Oracle Logs
DP-3	RAID-5 (7D+1P)	16	7.33 TB	Oracle fast recovery area

Oracle Automatic Storage Management Configuration

Oracle Automatic Storage Management (ASM) is the recommended storage management solution from Oracle that provides an alternative to conventional volume managers, file systems, and raw devices. The ASM disk group consists of pool volumes created from each dynamic provisioning or dynamic tiering pool, and then configured as ASM disk.

Table 6 lists the ASM disk group each storage blade uses in this solution to support of a container database.

Table 6. ASM Configuration

ASM Disk Group	*Purpose*	*Storage Pool*	*LUN Count*	*LUN Capacity*	*Path*	*Multipath*
DATA	Oracle data diskgroup	DT-1	32	100 GB	4	DM
REDO1	Redo log 1 diskgroup	DP-2	4	80 GB	4	DM
REDO2	Redo log 2 diskgroup	DP-2	4	80 GB	4	DM

Table 6. ASM Configuration (Continued)

ASM Disk Group	Purpose	Storage Pool	LUN Count	LUN Capacity	Path	Multipath
ARCH	Archive log diskgroup	DP-2	4	80 GB	4	DM
FRA	Fast recovery area diskgroup	DP-3	8	70 GB	4	DM

Database Layout

The database layout design uses recommended practices from Hitachi Data Systems for an Oracle container database to support the consolidation of many Oracle pluggable databases.

The layout also takes into account Oracle Automatic Storage Management best practices when using Hitachi storage.

Base the storage design for database layout needs on the requirements of a specific application implementation. The design can vary greatly from one implementation to another. The components in this solution set have the flexibility for use in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

- **Data and Indexes Tablespace**
  - Place DATA tablespace in the Data ASM diskgroup.
  - Set the tablespace to a small initial size with auto extend enabled to maximize storage utilization from thin provisioning.
- **Temp Tablespace**
  - Place TEMP tablespace in the Data ASM diskgroup.
  - Limit the size of each small tempfile to 30 GB.
- **UNDO Tablespace**
  - Place UNDO tablespace in the Data ASM diskgroup.
  - Limit the size of each small undo datafile to 30 GB.
- **Online Redo Logs**
  - Recommend using two ASM diskgroups for redo logs redundancy.
  - For this solution, use REDO1 diskgroup and REDO2 diskgroup for online redo logs.
- **Data Block Size**
  - Set the database block size to 8 KB.
  - Set the ASM allocation unit to 1 MB.

SAN Architecture

The SAN architecture for this solution consists of two in-chassis Brocade 5460 Fibre Channel switch modules and two external Brocade 6510 switches. Adding two external Brocade 6510 switches permit expansion beyond a single Hitachi Compute Rack 500 chassis.

Figure 8 shows the Fibre Channel physical connections for this solution.

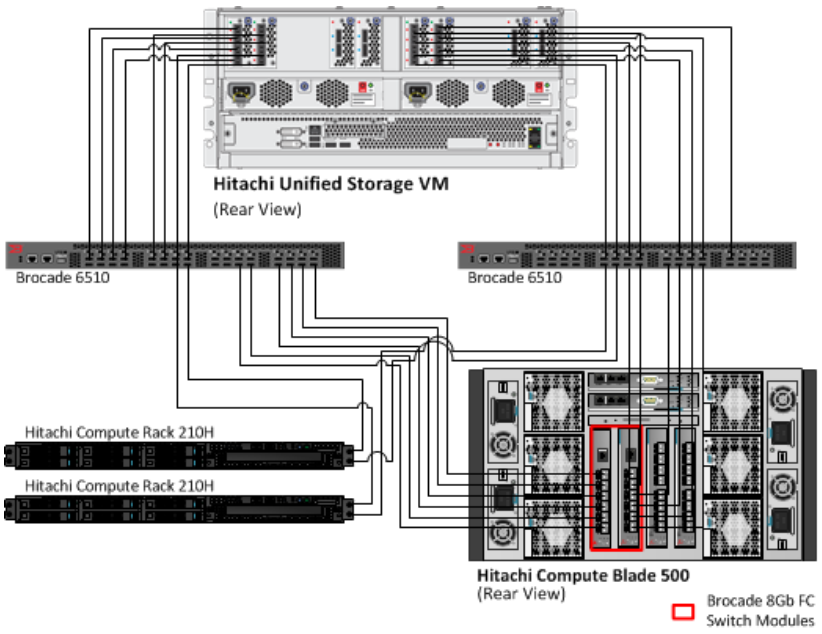

Figure 8

Fibre Channel Zoning

Configure Fibre Channel zoning for this solution using Hitachi Unified Compute Platform Director. Dynamically map each host HBA port to four storage ports on Hitachi Unified Storage VM.

Use round-robin load balancing on host-side HBA port with the least active path polity. This ensures distributing the storage resource load evenly for optimal database performance.

Multipath I/O Options

This solution requires using multipathing I/O software to provide link redundancy and performance advantages of load balancing of the I/O between all the Fibre Channel links.

This solution uses the Red Hat Enterprise Linux native device mapper multipathing. Device mapper multipathing (DM-Multipath) allows you to configure multiple I/O paths between server nodes and storage arrays into a single device. These I/O paths are physical SAN connections that can include separate cables, switches, and controllers. Multipathing aggregates the I/O paths, creating a new device that consists of the aggregated paths.

Network Architecture

The network architecture for this solution uses the following:

- 1 Gb/sec connections for the management network
- 10 Gb/sec connections for server blade network communications

The server blade network connections are provided with a pair of Brocade 10 Gb/sec Ethernet switch modules that are uplinked to the Brocade top-of-rack LAN switches.

Figure 9 shows the physical network implementation for this solution.

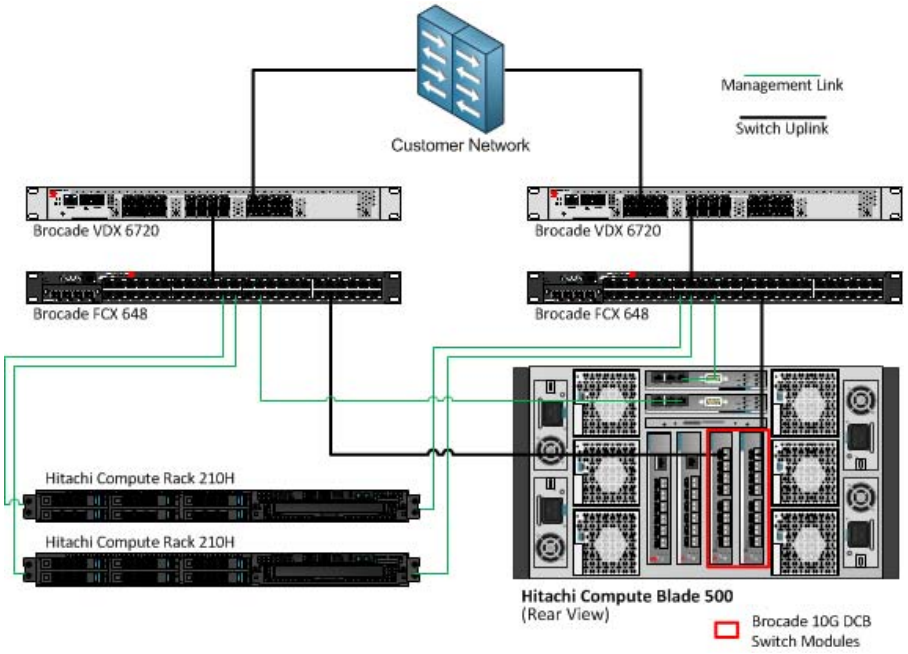

Figure 9

Server Blade Networking

Each server blade is configured with a dual-port 10 Gb/sec LAN mezzanine card that provides redundant network connections for the Linux operating system in a NIC bond. The CNA ports are connected internally to the Brocade 10 Gb/sec DCB switch modules. Figure 10 shows the internal links for each CNA port.

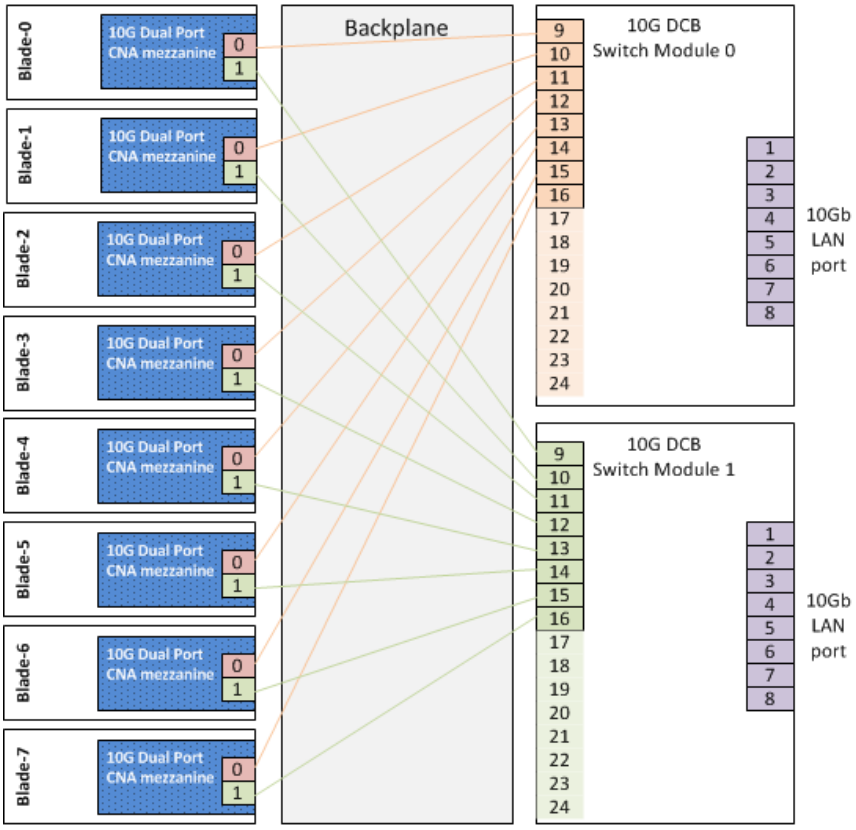

Figure 10

Engineering Validation

This section describes the test methodology used to test this reference architecture.

Test Methodology

The test methodology validates the compute, storage, and Oracle database performance advantage of using this reference architecture from Hitachi Data Systems. These scenarios also provide data to understand the overall capability of the Oracle container database and pluggable databases.

Storage Benchmark

To ensure that the storage system and architecture meet the requirement to support of an Oracle Multitenant database workload, this solution was tested using VDbench.

The VDbench test the following:

- Storage IOPS capability for small random I/O
- Storage aggravated throughput capability for large sequential I/O
- I/O latency

Oracle Database Workload

To better understand the capability of this solution for Oracle container database implementation the Swingbench test suite were used to simulate the load of many order entry system running as pluggable databases within the Oracle container database.

Test Results

This section describes the results that were observed during the testing of this solution. To better understand the performance of each area and component of this architecture, each component was evaluated separately to ensure that optimal performance was achieved when the solution was under stress.

These test results are organized into the following categories:

- “Storage Performance” on page 22
- “Swingbench Workload Performance” on page 22

Storage Performance

This section describes how the Hitachi Unified Storage VM performed during the test period.

Table 7 lists the performance results for Hitachi Unified Storage VM.

Table 7. Storage IOPS Performance Test Results

Metric	Block Size	Average Response Time (msec)	Result
Maximum storage IOPS	8 KB	0.99	510,026 IOPS
Maximum per blade IOPS	8 KB	1	128,879 IOPS

Table 8 lists the throughput results for Hitachi Unified Storage VM.

Table 8. Storage Throughput Test Results

Metric	Block Size	Average Response Time	Result
Maximum solution throughput	64K	8	7,520 MB/sec
Maximum per blade throughput	64K	8	940 MB/sec

Swingbench Workload Performance

These are the test results for the Oracle database performance using Swingbench for load simulation. This test simulates an online order entry system to observe the capability and performance of an OLTP database workload using an Oracle pluggable database.

Table 9 lists results for the Oracle database performance using one container database on a single server blade.

Table 9. Oracle Container Database Workload Test Results

Metric	Result
Maximum TPS	3764
Insert Latency	4 msec
Update Latency	1 msec
Read Latency	2 msec
Write Latency	5 msec
Total Transaction during test run	7905241
Failed Transaction	0

Table 10 lists the details of the order entry system transactions.

Table 10. Order Entry System Transaction Details

Metric	*Result*
Test Duration	30 minutes
Select Statements	39524283
Insert Statements	17819624
Update Statements	8120049
Commit Statements	10422898

Conclusion

This Hitachi Unified Compute Platform for Oracle Database solution using Oracle Multitenant delivers a number of benefits, including the following:

- Increased scalability
- Highly available
- Reduced deployment time with simplified management and administration

Using Hitachi Dynamic Tiering to create a pool with many disk tiers allows Oracle Database to store and retrieve data at a much faster rate. It moves hot data onto SSD disks while leaving history data not often accessed on less expensive SAS disks.

The Hitachi Data Systems compute and storage products provide the following in this reference architecture:

- Optimized performance for transactional applications with minimal latency
- Balanced architecture approach that maximizes the performance of all hardware and software components of the design
- Validated configuration that meets or exceeds industry best practices for a share database environment using Oracle Multitenant

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-376-00, April 2015