

# Deploying SAP with Oracle in Linux Environments Using Hitachi Dynamic Tiering on the Hitachi Virtual Storage Platform

Implementation Guide

*By Anantha Adiga*

November 2010



## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). Be sure to include the title of this white paper in your email message.

# Table of Contents

<b>Solution Overview .....</b>	<b>4</b>
Hitachi Virtual Storage Platform .....	5
Hitachi Dynamic Provisioning Software .....	5
Hitachi Dynamic Tiering Software .....	6
<b>Tested Solution Components.....</b>	<b>7</b>
Tested Deployment Hardware.....	7
Tested Deployment Software .....	9
<b>Solution Implementation.....</b>	<b>9</b>
Configure the SAN .....	10
Configure Servers and the Operating Environment .....	11
Configure Storage .....	12
Install, Configure and Verify SAP Software.....	21
<b>Appendix — Tested Deployment Oracle Parameter File .....</b>	<b>22</b>

# Deploying SAP in Oracle and Linux Environments Using Hitachi Dynamic Tiering Virtual Storage Platform

## Implementation Guide

SAP is the leader in the field of business process management solutions. Organizations rely on SAP's software to automate mission-critical operations. Very often SAP applications must be available 24 hours a day, seven days a week, meaning a loss of application availability can have a devastating effect on businesses. That is just as true for a small business as it is for a global enterprise.

SAP's Enterprise Resource Planning (ERP) software integrates business processes and information technologies. Generally, SAP ERP implementations are enterprise-wide and integrate a variety of SAP business modules.

The mission-critical nature of SAP deployments requires an IT infrastructure that delivers high performance and is highly available. The infrastructure must be easy to deploy, manage and scale. It is crucial that the storage area network (SAN), which houses database files, not only is capable of handling current loads, but is also capable of scaling to meet future transaction loads.

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage.

Hitachi Virtual Storage Platform enables companies to speed SAP deployments, to grow SAP application environments without limits and to reduce risk and cost. The storage virtualization features of the Virtual Storage Platform are designed to meet the increasing demands of server virtualization. The first requirement for support of server virtualization is the ability to scale, which has been greatly enhanced. The second is to provide integration with virtual server APIs to make the virtualization of servers more efficient. The third is to simplify the management and provide application and server visibility into the storage virtualization infrastructure for openness, communication and accountability.

Hitachi Dynamic Tiering software is a revolutionary solution that eliminates the time-consuming manual processes of data classification and movement between storage tiers, thus optimizing tiered storage usage while improving performance. It is only available on the Hitachi Virtual Storage Platform.

This white paper provides deployment guidelines for a successful implementation of SAP in Oracle and Linux environments using the Hitachi Virtual Storage Platform. It describes the hardware and software required to build the solution and provides links to supporting documentation needed to build, test and validate the solution. Although this document does not provide step-by-step detailed instructions for each and every task required to deploy the solution, it does provide a consolidated resource where administrators can easily locate related materials needed to construct a functional solution.

This white paper is written for SAP and storage administrators for businesses of all sizes who are deploying SAP on the Hitachi Virtual Storage Platform. It assumes familiarity with SAN technologies and tools and SAP products and architecture.

## Solution Overview

This white paper describes a solution that includes an SAP client that accesses the SAP application infrastructure, which in turn stores data in and accesses data from a Hitachi Virtual Storage Platform, as shown in Figure 1.




Figure 1

## Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the new Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. The Virtual Storage Platform provides virtual storage that meets the growing demands of server virtualization.

The trend in server virtualization is to consolidate the I/O workload of many servers onto a single storage system. As more virtual machines are consolidated onto a physical host, storage systems must be able to dynamically add more storage resources to keep up with I/O demand. The 3D scaling capability of the Virtual Storage Platform meets that requirement.

Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. With the Hitachi Virtual Storage Platform, you can increase performance, capacity and connectivity by adding cache, processors, connections and disks to the base system. A virtual server that accesses the storage system can use all these resources, which act as one system managed as a common pool of resources.

Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. By scaling out you can support increased resource needs in virtualized server environments.

Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform's pool of storage resources. Once virtualized, external data can then be migrated, tiered, replicated and managed by the Virtual Storage Platform. In this manner, older data storage systems can gain a longer useful life. You can extend distance replication for business continuity to lower-cost, lower-function storage systems by virtualizing them behind a Virtual Storage Platform.

The switch matrix architecture of the Virtual Storage Platform makes all of this possible. It connects the basic components, front-end directors, back-end directors, global cache modules and virtual storage directors. You can add redundant pairs of directors and cache modules as required without disruption to connected host servers. All these resources are tightly coupled through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

For more information about the Hitachi Virtual Storage Platform, see the Hitachi Data Systems [web site](#).

## Hitachi Dynamic Provisioning Software

On the Virtual Storage Platform, Hitachi Dynamic Provisioning software provides wide striping and thin provisioning functionalities. In the most basic sense, Hitachi Dynamic Provisioning software is similar to the use of a host-based logical volume manager (LVM), but with several additional features available within the Hitachi Virtual Storage Platform and without the need to install software on the host or incur host processing overhead. Hitachi Dynamic Provisioning software provides for one or more pools of wide striping across many RAID groups within a Virtual Storage Platform. One or more Dynamic Provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB (with no initial physical space allocated) are created against each pool.

Primarily, you deploy Hitachi Dynamic Provisioning software to avoid the routine issue of hot spots that occur on logical devices (LDEVs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many RAID groups as members of a striped Dynamic Provisioning pool underneath the virtual or logical volumes seen by the hosts, a host workload is distributed across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning software also carries the side benefit of thin provisioning, where physical space is only assigned from the pool to the DP-VOL as needed using 42MB pages, up to the logical size specified for each DP-VOL. A pool can also be dynamically expanded by adding more capacity or reduced by withdrawing pool capacity. Either operation is performed without disruption or requiring downtime. Upon expansion, a pool can be rebalanced so that the data and workload are wide striped evenly across the current and newly added RAID groups that make up the pool.

Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration, as follows:

- Optimizes or “right-sizes” storage performance and capacity based on business or application requirements.
- Supports deferring storage capacity upgrades to align with actual business usage.
- Simplifies the storage administration process.
- Provides performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.
- Eliminates hot spots across the different RAID groups by smoothing the combined workload.
- Significantly improves capacity utilization.

For more information, see the Hitachi Dynamic Provisioning software [datasheet](#).

## Hitachi Dynamic Tiering Software

Hitachi Dynamic Tiering software simplifies storage administration, improves performance and reduces costs in a tiered storage environment. Through automation, it eliminates the complexity of managing multiple tiers and moving data around to align storage tiers with business needs.

Most data is rarely accessed after it is created. As a result, it should not be stored on your most expensive tier of storage, but instead moved to a lower less expensive storage tier. Defining where and for how long data should reside at any point in its lifecycle can be complex and problematic.

Many organizations use a data tiering approach to manage application performance, manually provisioning space from several storage technologies with different performance and cost characteristics. Using this approach, data specialists typically look to past usage patterns to determine how to manually configure tiering, making the storage infrastructure unable to effectively respond to dynamic application and data use. If usage patterns change rapidly, manually tiered storage systems produce less than optimal results.

Hitachi Dynamic Tiering takes the automation of tiered storage to a new level. It enables the management of multiple storage tiers as a single entity. By leveraging the existing features of Hitachi Dynamic Provisioning software, Hitachi Dynamic Tiering presents a new kind of virtual volume with embedded smart tiering that monitors access and moves data at the 42MB page level. Hitachi Dynamic Provisioning software breaks the volume into pages and 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. It allows you to ensure that the right data is in the right place at the right time.

Hitachi Dynamic Tiering automatically moves pages of data within virtual volumes configured on a Dynamic Provisioning pool to the most appropriate media according to workload. This maximizes service levels and minimizes total cost of storage ownership. After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time and makes decisions on moving data between the available storage tiers based on actual use. If a page on a lower tier is accessed frequently, Hitachi Dynamic Tiering moves it a higher tier. Using this approach, Hitachi Dynamic Tiering enables you to improve the availability and performance of your storage systems and the applications using that storage.

Previously, each Dynamic Provisioning pool had to be created using one RAID level and one disk type. Hitachi Dynamic Tiering on the Virtual Storage Platform allows a single pool to contain tiers made up of multiple types of RAID groups and any type of disk. Hitachi Dynamic Tiering manages the various tiers within a Dynamic Provisioning 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.

Hitachi Dynamic Tiering is ideal for database applications and can provide improved performance for databases stored in a file system by keeping the frequently referenced data, such as indices, on tier 1 storage while moving less frequently referenced pages to a lower tier of storage. Hitachi Dynamic Tiering is also ideal for file and content data that requires high performance up front but becomes inactive over time. It is ideal for Exchange 2010 which provides very large mail boxes which are replicated for fast recovery. The mailboxes can be allocated to Dynamic Provisioning pools where recent mail resides on tier 1 and inactive mail is automatically migrated to lower tiers. The self-adjusting system can optimize performance based on demand, enabling content service providers to manage their storage infrastructure more effectively and allowing the use of less expensive, higher capacity disk drives such as SATA.

## Tested Solution Components

The following sections describe the hardware and software components required to deploy this solution.

### Tested Deployment Hardware

In the tested deployment, the SAP client tier consists of Windows-based servers with SAP GUI installed. The clients connect to the SAP system in the SAP application tier via a Cisco Catalyst 4500 switch. The Storage Navigator is installed on a Windows based management. The SAP Solution Manager was installed on a Linux server. The SAP application tier and database tier are installed on a Linux server. The SAP Solution manager and the SAP ERP application server are connected to the Hitachi Virtual Storage system via two Brocade 5300 switches. All servers were connected to a Hitachi Virtual Storage Platform via two Brocade 5300 switches

Table 1 describes the hardware used to test this solution in the Hitachi Data Systems lab.

**Table 1. Tested Deployment Hardware**

<i>Hardware</i>	<i>Quantity</i>	<i>Configuration</i>	<i>Role</i>
Hitachi Virtual Storage Platform storage system	1	Firmware version 70-00-53-00/00 4 Fibre Channel ports 1 pair of front-end directors 2 pairs of back-end directors 28 x 300GB 15K RPM SAS Drives 40 x 300GB 10K RPM SAS Drives 38 GB Cache	Primary storage
Brocade 5300 SAN Fibre Channel switch	2	Rev03 00.20.48.00 10 8Gb Fibre Channel ports used	N/A
Dell Power Edge R905 server	1	4 quad core AMD Opteron 1.9GHz processors 64GB memory Emulex LPe12002-M8 8Gb 2-port PCIe Fibre Channel Adapter, fwrev : 1.10A5 (U3D1.10A5), sli-3	SAP Solution Manager server
Dell Power Edge R905 server	1	4 quad core AMD Opteron 1.9GHz processors 64GB memory Emulex LPe12002-M8 8Gb 2-port PCIe Fibre Channel Adapter, fwrev : 1.10A5 (U3D1.10A5), sli-3	SAP ERP server
Dell PowerEdge 750	2	Intel Pentium 3.0GHz processor 1GB memory	SAP clients
Dell PowerEdge 750	1	Intel Pentium 3.0GHz processor 1GB memory	Management server with access to Hitachi Storage Navigator software

## Tested Deployment Software

Table 2 lists the software used to test this solution in the Hitachi Data Systems lab.

**Table 2. Tested Deployment Software**

<i>Software</i>	<i>Version</i>
Hitachi Dynamic Link Manager	6.0.1.0.804
Hitachi Storage Navigator	N/A
SAP ERP	6.0 Special Release 3 (SR3)
SAP Solution Manager	7.0
SAP GUI	7.1
Oracle Database 10g Enterprise Edition	10.2.0.4

## Solution Implementation

A successful deployment of this solution involves the following high-level steps:

1. Configure the SAN.
2. Configure servers and the operating environment.
3. Configure storage.
4. Install, configure and verify SAP software.

Your checklist might vary based on your environment. For more information about each of these high-level tasks, see the following resources:

- Hitachi Virtual Storage Platform Storage Navigator User's Guide
- Hitachi Dynamic Link Manager Users Guide for Linux
- [SAP Notes](#)

Best practice is to deploy SAP Solution Manager and SAP ERP software on separate servers.

## Configure the SAN

Figure 2 illustrates zones and the redundant paths that are managed by Hitachi Dynamic Link Manager Advanced software in the deployment tested for this solution. The multiple I/O paths provide high-availability from the hosts to storage via a Fibre Channel switches.

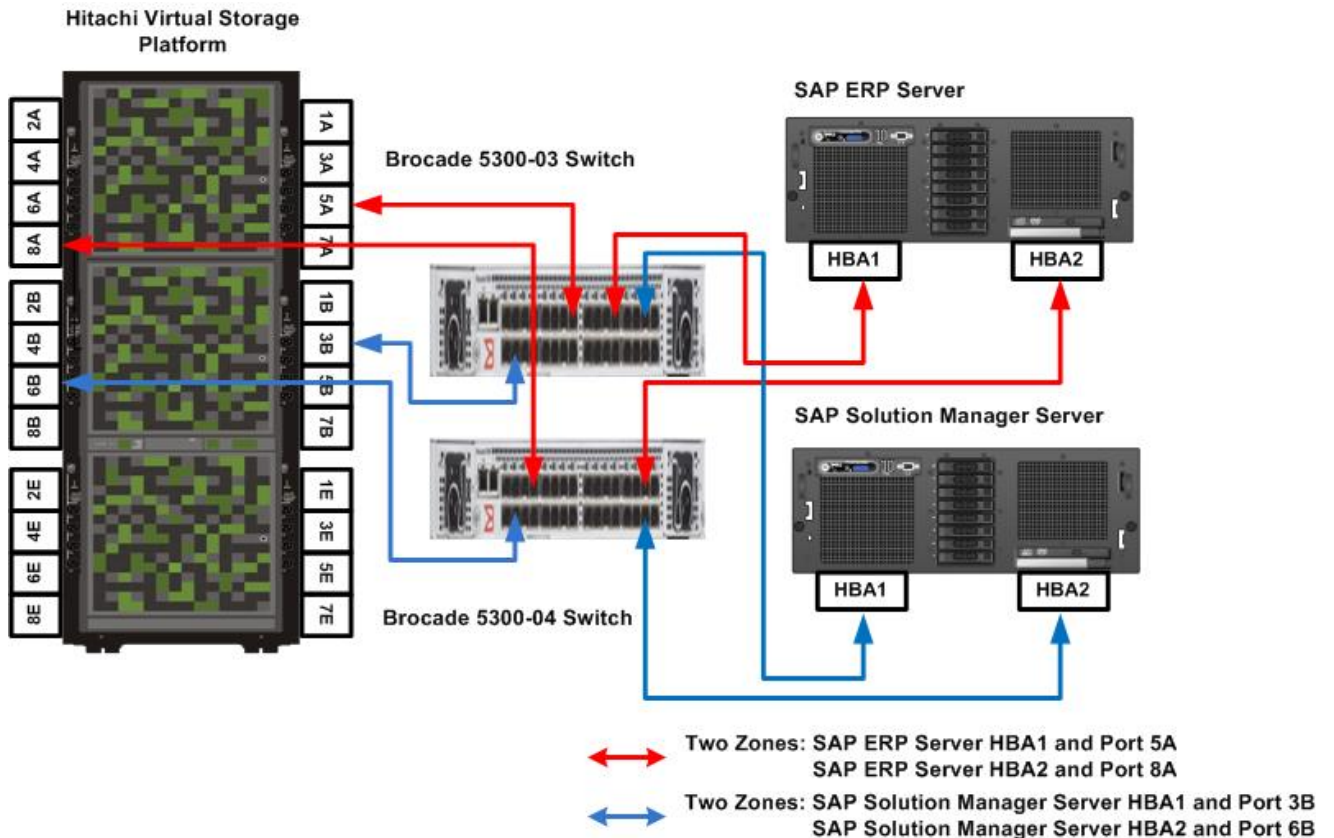


Figure 2

When deploying an SAP system, the storage used by SAP Solution Manager server and SAP ERP server must be isolated from each other in a SAN. As shown in Figure 2, four separate Fibre Channel zones are created, one from each HBA of the SAP Solution Manager server and SAP ERP server. The Fibre Channel zones for the SAP Solution Manager and the SAP ERP server do not share the Fibre Channel port of the storage systems.

World wide names (WWNs) of the HBAs residing on a physical server can be assigned to a host group and then a logical device (LDEV) can be associated with each host group. This means that a group of LDEVs can be isolated to be used by only the assigned physical servers.

The Hitachi Virtual Storage Platform supports active-active multipath connectivity. To obtain maximum availability, design and implement your host-storage connections so that at least two unique paths exist from the host to the storage system. Hitachi Data Systems recommends the use of dual SAN fabrics, multiple HBAs and host-based multipathing software when deploying SAP systems.

Attach the SAP ERP server, and SAP Solution Manager Server to redundant HBAs via a SAN consisting of a redundant fabric. The Hitachi Virtual Storage Platform, which was used in testing this solution, has 8Gb Fibre Channel ports. In the tested deployment, the SAP ERP and SAP Solution Manager servers were connected via two HBAs each to two Brocade 5300 switches. The Brocade 5300 switches, in turn, are connected to four different ports of the Hitachi Virtual Storage Platform.

The choice of load-balancing algorithm depends on the specific environment and access patterns of the application. In most cases, the extended round robin algorithm provides the best overall performance. In some environments, such as an environment shared with other applications, one of the other algorithms might give the best overall performance.

## Configure Servers and the Operating Environment

When deploying SAP Solution Manager and SAP ERP software for this solution, Hitachi Data Systems used two Dell PowerEdge R905 servers with four quad core CPUs and 64GB memory. Each server had two Emulex LPe12002-M8 8Gb two-port HBAs and two 1GB network interface cards.

Two SAP client servers and one management server for access to Hitachi Storage Navigator software were deployed using Dell PowerEdge 750 servers with two CPUs and 2GB memory.

Hitachi Data Systems installed Red Hat Enterprise Linux version 5.2 (x86 64 bit) on the SAP Solution Manager and SAP ERP servers. Microsoft Windows 2008 was installed on SAP clients and the management server to access Hitachi Storage Navigator.

For more information about supported operating systems, see the following SAP Notes, which are available to licensed customers from SAP's [web site](#):

- SAP Note 171356, *SAP Software on Linux: Essential*
- SAP Note 1048303, *Red Hat Enterprise Linux 5.x: Installation and Upgrade*
- SAP Note 1172419, *Linux: Supported Java Versions on the x86\_64 Platform*

### *Complete SAP Pre-installation Tasks*

You must complete several pre-installation tasks before installing either SAP Solution Manager or SAP ERP software or configuring your storage. For more information about pre-installation tasks, see the following SAP Notes, which are available to licensed customers from SAP's [web site](#):

- SAP Note 855498, *Installation Pre-requisite Checker, SAP Guides*
- SAP Note 1145779, *SAP Solution Manager 7 Installation*
- SAP Note 1052298, *SAP NetWeaver 7.0(2004s) SR3 Installation on UNIX/Oracle*
- SAP Note 830576, *Parameter Recommendations for Oracle 10g*
- SAP Note 1289199, *Information about Oracle Parameters*
- SAP Note 828268, *Oracle Database 10g: New Functions*
- SAP Note 871735, *Current Patch Set for Oracle 10.2.0*

Hitachi Data Systems recommends using a logical volume manager (LVM) to create volume groups and logical volumes. For this solution, the logical volumes, volume groups and file systems are carefully laid out based on the application needs. Both the SAP Solution Manager Server and the SAP ERP server use similar volume group, logical volume and file system configuration.

## Configure Storage

The following sections describe how to configure Hitachi Virtual Storage Platform for this solution. Use Hitachi Storage Navigator software to create and manage the storage objects.

### *Hitachi Dynamic Tiering Overview*

Hitachi Dynamic Tiering creates a hierarchy of multiple media types (SSD, SAS, SATA) with different cost and performance characteristics within a storage pool. It also enables relocating frequently accessed data to high-speed media and infrequently accessed data to low-speed media.

A tier is defined by the performance characteristics of the media type. Media with short response times are positioned as higher tiers, and media with longer response times are positioned as lower tiers.

A Hitachi Dynamic Tiering pool is simply a Hitachi Dynamic Provisioning pool with the multi-tiering function enabled.

The Hitachi Dynamic Provisioning pool is divided into pages. Virtual volumes (V-VOLs) are created from the Dynamic Provisioning pool. The virtual volume is a logical device is also referred to as an LDEV. When data is written to a virtual volume in a Dynamic Tiering pool, pages are allocated to it from the top tier. When pages are accessed in a Hitachi Dynamic Tiering environment, the I/O activity is monitored and the most frequently referenced pages are placed in the highest tier and the least frequently referenced pages are placed in the lowest tier. The lowest tier also contains much of the spare capacity. The pages are relocated to appropriate tiers within the pool for performance optimization purposes. By default, Hitachi Dynamic Tiering software automatically initiates page relocation periodically. However, the storage administrator can change this to manually initiate the relocation between tiers.

### *Sample Hitachi Dynamic Tiering Configuration*

Table 4 lists storage configuration details of the Dynamic Tiering pools created for the SAP ERP server and the SAP Solution Manager server in the Hitachi Data Systems lab. This configuration has separate Dynamic Tiering pools for the SAP Solution Manager and SAP ERP servers. The parity groups are of RAID-5 (3D+1P) type. Multi-tiering is enabled on both Dynamic Provisioning pools, which are referred to as Hitachi Dynamic Tiering pools. Each Dynamic Tiering pool has two tiers and five pool volumes. Four pool volumes are carved from SAS 10k RPM, RAID-5 (3D+1P) parity groups and one pool volume is carved from a SAS 15k RPM, RAID-5 (3D+1P) parity group. Each Dynamic Tiering pool has 14 V-VOLs. These V-VOLs are presented to the servers as host LUNs.

**Table 4. Hitachi Dynamic Tiering Pool Configuration**

<i>Dynamic Tiering Pool</i>	<i>Number of Drives</i>	<i>Drive Capacity</i>	<i>Pool Capacity (TB)</i>	<i>Number of Pool Volumes</i>	<i>Number of V-VOLs</i>	<i>Role</i>
ERPPOOL	16	300GB SAS 10K RPM	3.6	4 1	14	SAP ERP server
	4	146GB SAS 15K RPM				
SMPOOL	16	300GB SAS 10K RPM	3.6	4 1	14	SAP Solution Manager server
	4	146GB SAS 15K RPM				

All V-VOLs created in the ERPPOOL Dynamic Tiering pool were mapped to storage ports 5A and 8A and all V-VOLs created in the SMPOOL Dynamic Tiering pool were mapped to storage ports 5A and 8A.

Table 5 lists configuration details for the SAP ERP server.

**Table 5. V-VOL Configuration in ERPPOOL for SAP ERP Server**

<i>V-VOL</i>	<i>Host LUN</i>	<i>Size (GB)</i>
00:00:07	07	50
00:00:08	08	50
00:00:0A	0A	50
00:00:0B	0B	50
00:00:0C	0C	50
00:00:0D	0D	200
00:00:0E	0E	200
00:00:0F	0F	200
00:00:016	16	200
00:00:017	17	200
00:00:018	18	200
00:00:019	19	200
00:00:1A	1A	200
00:00:1B	1B	200

Table 6 lists configuration details for the SAP Solution Manager server.

**Table 6. V-VOL Storage Configuration for SAP Solution Manager Server**

<i>V-VOL</i>	<i>Host LUN</i>	<i>Size (GB)</i>
00:00:07	07	50
00:00:08	08	50
00:00:0A	0A	50
00:00:0B	0B	50
00:00:0C	0C	50
00:00:0D	0D	200
00:00:0E	0E	200
00:00:0F	0F	200
00:00:016	16	200
00:00:017	17	200
00:00:018	18	200
00:00:019	19	200
00:00:1A	1A	200
00:00:1B	1B	200

Table 6 lists details of disk groups, logical volumes and file systems for the SAP ERP server and the SAP Solution Manager server. Both servers have the same layout.

**Table 6. SAP Solution Manager Server and SAP ERP Server Disk Group, Logical Volume and File System Layout**

<i>Host LUN</i>	<i>Disk Group</i>	<i>Logical Volume</i>	<i>File System Mount Point</i>	<i>Usage</i>
07	appbin_vg	appbin_sap_lv	/sapmnt, /usr/SAP/<SID>, and /usr/SAP/trans	File system for central repository for SAP Solution Manager, SAP binaries, and central location for moving data to another system when required
08	appbin_v	appbin_ora_lv	/oracle	File system for Oracle binaries for installation of Oracle, Oracle client, home directory for Oracle user, and to store data temporarily for online data re-organization respectively
0A	appl ogA_vg	appl ogA_lv	/oracle/<SID>/origlogA and /oracle/<SID>/mirrlogB	File system for online redo logs and duplex logs of the database
0B	appl ogB_vg	appl ogB_lv	/oracle/<SID>/origlogB and /oracle/<SID>/mirrlogA	File system for online redo logs and duplex logs of the database
0C	apparch_vg	apparch_lv	/oracle/<SID>/saparch	File system for archived logs
0D	appsapdata_vg	appdata_sapdata_lv	/oracle/<SID>/sapdata1	File system for SAP Solution Manager data
0E	appsapdata_vg	appdata_sapdata_lv	/oracle/<SID>/sapdata2	File system for SAP Solution Manager data
0F	appsapdata_vg	appdata_sapdata_lv	/oracle/<SID>/sapdata3	File system for SAP Solution Manager data
16	appsapdata_vg	appdata_sapdata_lv	/oracle/<SID>/sapdata4	File system for SAP Solution Manager data
17	apporadata_vg	appdata_oradata_lv	/oracle/<SID>/oradata	File system for Oracle data dictionary, temporary, UNDO, and users table space data

## Create Dynamic Provisioning Pools

This procedure assumes that zone creation within the Fibre Channel fabric and the RAID groups creation steps for your environment are complete. Figure 3 shows the **Create Pools** window in Hitachi Storage Navigator software. To create Dynamic Provisioning pools with Hitachi Dynamic Tiering, follow these steps:

1. Choose **Actions > Pool > Create Pools**.

The **Create Pools** window displays.

The screenshot shows the 'Create Pools' wizard in Hitachi Storage Navigator. The window title is 'Create Pools' and it shows a progress bar with '1. Create Pools' selected and '2. Confirm' next to it. Below the progress bar, there is a brief description: 'This wizard lets you create pools for Dynamic Provisioning and COW Snapshot. Enter the information. Click Finish to confirm the creation, or click Next if you want to create LDEVs (virtual volumes) first.' The main configuration area includes: 'Pool Type' set to 'Dynamic Provisioning'; 'Multi-Tier Pool' with 'Enable' selected; 'Pool Volume Selection' section with 'Drive Type/RPM' and 'RAID Level' both set to 'Any', and a 'Select Pool VOLs' button; 'Pool Name' field with 'Prefix' and 'Initial Number' sub-fields, and a note '(Max. 32 characters total including max. 9-digit number)'; 'Options' section with 'Initial Pool ID' set to '0', 'Subscription Limit' set to a blank field, and 'User-Defined Threshold' set to '70'; 'Tier Management' section with 'Auto' selected; and a 'Cycle Time' set to '24' and 'Monitoring Period' from '00:00' to '23:59'.

2. From the **Pool Type** drop-down menu, choose **Dynamic Provisioning**.
3. In the **Multi-Tier Pool** field, ensure that the **Enable** radio button is selected.  
This enables Hitachi Dynamic Tiering for the Dynamic Provisioning pool.
4. Click the **Select Pool VOLs** button.

The **Select Pool VOLs** window displays.

5. Highlight the required pool volumes in the **Available Pool Volumes** pane and click the **OK** button.

You can select drives of different types, but you must select pool volumes of the same RAID type.

The **Create Pools** window displays with the **Total Selected Pool Volume** and **Total Selected Capacity** fields populated.

6. Assign a prefix for the pool name in the **Prefix** field.
7. (Optional) Assign an initial number for the pool name in the **Initial Number** field.
8. Expand the **Options** pane.
9. Assign a pool ID in the **Pool ID** field.
10. Assign a subscription limit in the **Subscription Limit** field.

This is the percentage of oversubscription for this pool that you allow in your environment.

11. Choose a value from the **User-Defined Threshold** drop-down menu and click the **Add** button.

The **User Defined Threshold** value determines when a pool capacity alert is triggered.

The **Selected Pools** pane is populated.

12. Ensure that the **Tier Management** option's **Auto** radio button is selected.

If you choose the **Manual** option, you must specify the start and end monitoring settings and start and stop relocation setting from the CLI interface instead of the GUI. For more information, see the Hitachi CCI User and Reference Guide.

Use the following command to set the monitoring period:

```
raidcom monitor pool -pool_id 0 -operation stop/start  
raidcom reallocate pool -pool_id 0 -operation start/stop
```

If you use the Auto option, you can set the start and end time for monitoring using the GUI.

13. Choose a value from the **Cycle Time** drop-down menu.

The default value is 24 for **Auto Tier Management** setting.

14. Assign a start time and an end time in the **Monitoring Period** fields.

The default settings are 00:00 to 23:59.

15. Click the **Finish** button.

The **Create Pools** window displays.

16. Click the **Apply** button.

## *Create LDEVs Within a Dynamic Provisioning Pool*

This procedure assumes that Dynamic HDT Pool creation in your environment is complete. To create one or more LDEVs within a Dynamic Provisioning pool using Hitachi Storage Navigator GUI, follow these steps:

1. Choose **Actions > Logical Device > Create LDEVs**.

The **Create LDEVs** window displays.

2. From the **Provisioning Type** drop-down menu, select **Dynamic Provisioning**.
3. From the **Emulation Type** drop down menu, select **OPEN-V**.
4. (Optional) Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

These options allow you to filter the available pools.

5. Click the **Select Pool** button.

The **Select Pool** window displays.

6. Highlight a pool in the **Available Pools** pane and click **OK**.

The **Create LDEVs** window displays with the **Selected Pool Name** and the **Selected Pool Capacity** fields populated.

7. Enter a capacity amount in the **LDEV Capacity** field and choose a unit of measure from the drop-down menu.
8. Enter the number of LDEVs of that size to be created in the **Number of LDEVs** field.
9. In the **LDEV Name** pane, assign a prefix in the **Prefix** field and assign an initial number in the **Initial Number** field.
10. Expand the **Options** pane.
11. Review the value in the **LDKC** field.

Modify the LDKC value if the default of 00 is not appropriate. This is most often the case if the storage will be configured with more than one LDKC.

12. Choose a value from the **CU** drop-down menu.
13. Choose a value from the **DEV** drop-down menu.
14. (Optional) Choose a value from **Interval** drop-down menu.

Leave this value at the default of 0 for sequential numbering of LDEVs. If you want a different numbering sequence, choose a different value.

15. Review the default values in the **Initial SSID** field, the **CLPR** field and **Processor Blade** field.

In most situations, use the default values. Change them only if your environment requires different values.

16. Click the **Add** button.

The **Selected LDEVs** pane is populated.

17. Click the **Finish** button.

The **Create LDEVs** window displays.

18. Click the **Apply** button.

### *Create Host Groups*

Hosts (WWNs) can be assigned to host groups and then the desired LDEVs can be associated with each host group. Each host in the host group has access to the LDEVs associated with that group. To create host groups using Storage Navigator software, follow these steps:

1. Choose **Actions > Ports/Host Groups > Create Host Groups**.

The **Create Host Groups** window displays.

2. Assign a name in the **Host Group Name** field.

3. From the **Host Mode** drop-down menu, choose **00[Standard]**.

4. In the **Available Hosts** pane, if the required host is not in the list click **Add New Host** button

The **Add New Host** window displays.

5. Enter the HBA WWN of the hosts HBA port.

6. Enter the name of the host.

7. Click the **OK** button.

8. In the **Available Hosts** pane, highlight one or more hosts.

9. In the **Available Ports** pane, highlight one or more ports.

10. Click the **Add** button.

The **Selected Host Groups** pane is populated.

11. Click the **Finish** button.

The **Create Host Groups** window displays.

12. Click the **Apply** button.

## Map LDEVs

To map LDEVs to host groups using Hitachi Storage Navigator software, follow these steps:

1. Choose **Actions > Logical Device > Add LUN Paths**.

The **Add LUN Paths** window displays.

2. In the **Available LDEVs** pane, highlight one or more LDEVs.
3. Click the **Add** button.

The **Selected LDEVs** pane is populated.

4. Click the **Next** button.

The **Add LUN Paths** window displays.

5. In the **Available Host Groups** pane, highlight one or more host groups.
6. Click **Add**.

The **Selected Host Groups** pane is populated.

7. Click **Next**.

The **Add LUN Paths** window displays.

8. Click the **Finish** button.

The **Add LUN Paths** window displays.

9. Click the **Apply** button.

## Install and Configure Hitachi Dynamic Link Manager Advanced Software

Follow these steps to install Hitachi Dynamic Link Manager Advanced software and to create LVM objects and file systems on the disks on the ERP and Solutions Manager servers:

1. Execute the following command on the server on which you want to install Hitachi Dynamic Link Manager Advanced software:

```
. /i nstal l hdl m
```

When the installation is complete, Hitachi Dynamic Link Manager Advanced displays a message indicating that the software was successfully installed.

Hitachi Dynamic Link Manager Advanced software names your disk device using the following format:

```
sddl m<di sk-devi ce>
```

<di sk-devi ce> is a combination of letters from a to z identifying your disk device.

2. On each server, using LVM, create a physical volume on each disk device, using the following syntax:

```
pvccreate /dev/sddl m<di sk-devi ce>
```

3. On each server, using LVM, create volume groups, using the following syntax:

```
vgcreate -s 64 <volume group name> <physical volume name or names>
```

4. On each server, create logical volumes, using the following syntax:

```
lvcreate -l <size of logical volume> -n <volume name> <volume group name>
```

For more information about the logical volumes and volume groups needed for this solution, see Table 6.

5. On each server, create a file system for each logical volume, using the following syntax:

```
mke2fs -t ext2 /dev/<volume group name>/<logical volume name>
```

6. Mount each file system, using the following syntax:

```
mount -t ext2 -o async /dev/<volume group name>/<logical volume name>  
/<mount point>
```

For more information about the file systems and mount points needed for this solution, see Table 6.

## Install, Configure and Verify SAP Software

Install SAP Solution Manager and SAP ERP software on dedicated servers. SAP GUI software is installed on SAP clients.

It is important to verify that the newly installed and configured hardware and software are operating properly before you make your deployment live. The verification process involves logging in to each server and executing a few transactions.

For more information, see SAP Notes, which are available to licensed customers from SAP's [web site](#).

## Appendix — Tested Deployment Oracle Parameter File

The following listing shows the Oracle ERP server database initialization parameter settings for this deployment. These settings are dependent on the server configuration and other database requirements.

ini tPRD.ora

```
*. _B_TREE_BITMAP_PLANS=FALSE
*. _IN_MEMORY_UNDO=FALSE
*. _INDEX_JOIN_ENABLED=FALSE
*. _OPTIM_PEEK_USER_BINDS=FALSE
*. _OPTIMIZER_MJC_ENABLED=FALSE
*. _SORT_ELIMINATION_COST_RATIO=10
*. _TABLE_LOOKUP_PREFETCH_SIZE=0
*. background_dump_dest='/oracle/PRD/saptrace/background'
*. compatible=' 10. 2. 0'
*. control_file_record_keep_time=30
*. control_files='/oracle/PRD/origlogA/cntrl/cntrl PRD. dbf' , '/oracle/PRD/origlogB/cntrl/cntrl PRD. dbf' , '/oracle/PRD/sapdata1/cntrl/cntrl PRD. dbf'
*. core_dump_dest='/oracle/PRD/saptrace/background'
*. db_block_size=8192
*. db_cache_size=9133306675
*. db_files=254
*. db_name=' PRD'
*. dml_locks=4000
*. event=' 10191 trace name context forever, level 1'
*. FILESYSTEMIO_OPTIONS=' setall '
*. job_queue_processes=1
*. log_archive_dest='/oracle/PRD/saparch/PRDarch'
*. log_buffer=1048576
*. log_checkpoint_interval =0
*. log_checkpoints_to_alert=true
*. open_cursors=800
*. parallel_execution_message_size=16384
*. pga_aggregate_target=12177742233
```

- \*. processes=80
- \*. query\_rewrite\_enabled=' false'
- \*. recyclebin=' off'
- \*. remote\_login\_passwordfile=' exclusive'
- \*. remote\_os\_authent=true
- \*. replication\_dependency\_tracking=' false'
- \*. sessions=96
- \*. sga\_max\_size=18266613350
- \*. shared\_pool\_reserved\_size=913330667
- \*. shared\_pool\_size=9133306675
- \*. sort\_area\_retained\_size=0
- \*. sort\_area\_size=2097152
- \*. star\_transformation\_enabled=' true'
- \*. undo\_management=' AUTO'
- \*. undo\_retention=43200
- \*. undo\_tablespace=' PSAPUNDO'
- \*. user\_dump\_dest=' /oracle/PRD/saptrace/usertrace'

 **Hitachi Data Systems Corporation**

---

*Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries. All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.*

*Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation*

© Hitachi Data Systems Corporation 2010. All Rights Reserved. AS-065-01 November 2010

**Corporate Headquarters**

750 Central Expressway,  
Santa Clara, California 95050-2627 USA  
[www.hds.com](http://www.hds.com)

**Regional Contact Information**

**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)