

HITACHI
Inspire the Next



Increasing SAP Infrastructure Value with Hitachi Storage

Solution Guide

By James Stephens

October 2009

Summary

This solution guide introduces storage system deployment concepts that can help increase the total value of ownership for your storage investment. Implement these concepts to leverage storage system investments to provide additional data management functionality.

Hitachi Data Systems achieves this goal by expanding SAP's three-tier architecture to a fourth tier for storage. SAP's three-tier architecture spreads the work load across the tiers, enabling extended growth and flexibility. This white paper explains how to take best advantage of the features of Hitachi storage, and introduces the concept of a conceptual fourth tier for storage that handles all storage-related functions.

This solution guide explains SAP architecture, challenges that businesses face when implementing SAP, and the value of Hitachi Data Systems solutions in addressing those challenges. It is written for business teams that are responsible for designing and implementing their organizations' storage infrastructures for SAP. It assumes familiarity with basic storage terminologies and technologies.

For best results use Acrobat Reader 8.0.

Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The author recognizes and sincerely thanks the following contributors and reviewers of this document:

- Malcolm Brickwood
- Federick Brillantes
- Rick Butcher
- Singh Garewal
- Charles Hickman
- Lisa Pampuch
- Dhiren Patel
- Thomas Stutesman
- Rick Szarama

Table of Contents

SAP's Three-tier Approach	6
Client Tier.....	7
Application Tier.....	7
Database Tier.....	7
Landscapes	7
Test.....	8
Development.....	8
Training.....	8
Quality Assurance.....	9
Production.....	9
Other Landscapes.....	9
Tier Four Storage Services	9
Consolidation.....	10
Performance.....	10
Data Replication.....	12
Data Migration.....	13
Storage Management.....	14
Archiving.....	14
Putting it All Together	14
Conclusion	15

Increasing SAP Infrastructure Value with Hitachi Storage

Solution Guide

While many conversations about storage involve the total cost of ownership (TCO), it's also important to understand the total value of ownership (TVO) associated with a storage solution. A complex application such as SAP presents many opportunities to leverage your storage infrastructure beyond simply storing data.

While TCO is an investment metric that businesses want to lower, TVO is a metric that businesses want to increase. All investments have two critical factors: cost and return. That is also true of investments in infrastructure. Hitachi Data System's focus is not just on how to lower the TCO but also on how to increase the TVO. This guide introduces storage system deployment concepts that can help increase your storage investment TVO. Implement these concepts to leverage storage system investments to provide additional data management functionality.

Hitachi Data Systems achieves this goal by expanding SAP's three-tier architecture to a fourth tier for storage. SAP's three-tier architecture spreads the work load across the tiers, enabling extended growth and flexibility. To take best advantage of the features of Hitachi storage, perform storage-related functions in a fourth SAP tier dedicated to storage.

This solution guide explains the following concepts:

- SAP architecture
- Challenges that businesses face when implementing SAP
- Value of Hitachi Data Systems solutions in addressing those challenges

This paper is written for business teams that are responsible for designing and implementing their organizations' storage infrastructures for SAP. It assumes familiarity with basic storage terminologies and technologies.

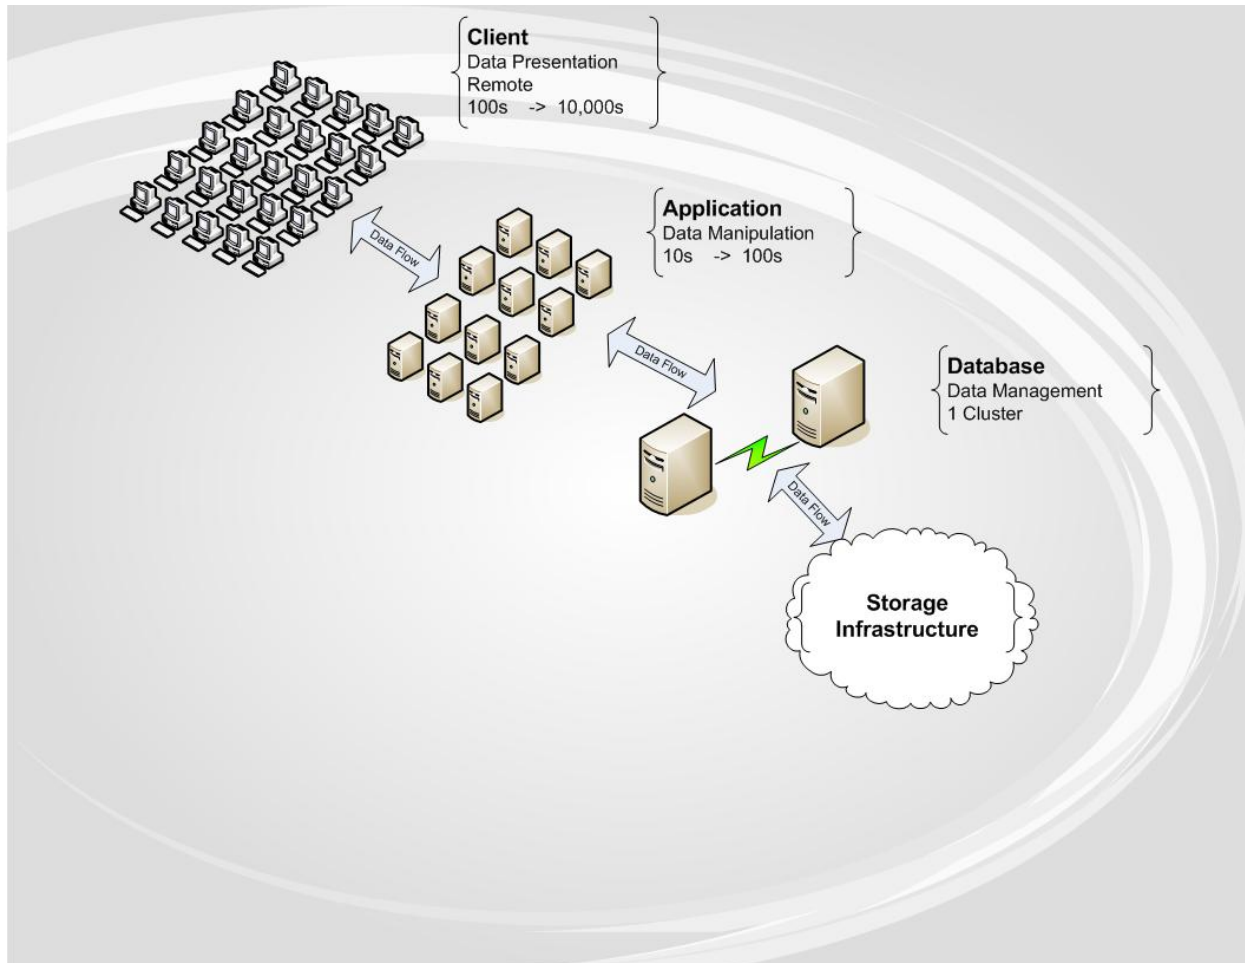
SAP's Three-tier Approach

SAP software was initially developed as a mainframe solution. When the application was ported to the open systems platform, SAP implemented the following three-tier architecture:

- Client tier
- Application tier
- Database tier

Figure 1 illustrates SAP's three-tier architecture.

Figure 1. SAP Three-tier Architecture



Note that SAP uses the term *tier* in a very different way than storage companies use the term. SAP's three-tier architecture design enables deployment of SAP applications in large centralized global implementations that are accessible from anywhere at any time. This architecture allows SAP applications to be scaled to meet the most complex, dispersed and voluminous requirements.

Each tier performs a specific function in the application stack. Although all functions can be performed within a single tier, doing so limits the flexibility and scalability of the application.

Client Tier

The client tier is the front end of the SAP application that is presented to the user. You can think of the client tier as the end user's desktop. The client tier is made up of client servers or devices that manage end-user facing functions. An SAP deployment might have thousands of clients using a variety of devices, including hand-held scanners, label readers, laptops, thin clients and more. Any device that supports end-user I/O is a potential client.

Simple data entry and display functions are performed on data that is provided by the application tier. All data to and from the client tier comes through the application tier.

The client tier has no direct impact on the development of the storage architecture and infrastructure because it has no direct interaction with that infrastructure.

Application Tier

The application tier is made up of application servers that manage data manipulation and calculations. It feeds data to the client tier and receives the data from that tier. The application tier is where all data manipulation occurs. The application servers in this tier perform specialized functions to optimize the application stack. For example, a server might be dedicated to performing a specific financial function. The SAP application balances those requests among the application servers. If a specific server fails, the load is spread across the remaining servers. An SAP implementation can have hundreds of application servers; these are candidates for virtualization.

Data used by the application tier is not stored locally but routed through the database tier. Because no application server maintains a static copy of SAP data, these servers do not need and do not have any direct interaction with the storage infrastructure. Consequently, the structure of the application tier has no effect on the storage infrastructure architecture or implementation.

Database Tier

The database servers that make up the database tier are responsible for all data management functions. All data that is stored within the storage infrastructure is managed by the servers in the database tier. This tier normally consists of a clustered pair of servers. The database tier is generally the source of any performance issues and scalability constraints in an SAP deployment. If these servers are tasked with functions such as host-based mirroring, they can become I/O constrained and can have a negative effect on the entire deployment.


It is important to note that the average database in an SAP deployment contains tens of thousands of objects. Many of these objects have little or no data at the initial installation but grow over time. However, they require a space allocation at the initial installation and in many cases the allocation is the result of estimates of future growth.

Landscapes

Every SAP deployment has multiple environments that are referred to as *landscapes*. A landscape is a complete SAP deployment. When organizations deploy SAP applications, they actually deploy them multiple times and these all run in parallel. When customers are asked about the size of their SAP deployments, they often respond with the size of their production (PRD) landscapes. The actual storage required for SAP deployments is many times that number, depending on the number and type of landscapes.

An SAP implementation has several landscapes, including test, development, training and more. However, given the mission-critical nature of SAP implementations, most CIOs focus on protecting their production landscapes. Figure 2 illustrates some of the landscapes in a typical SAP deployment.

Figure 2. SAP Landscapes



Test

Any SAP implementation begins with a test landscape, which is generally small. In fact, all three tiers might be implemented on a single server in a test landscape. Although fully functional, the data demands on both scale and functionality are minimal. During the life cycle of any SAP implementation, when new business processes and strategies are added to the implementation, the test landscape is where they are introduced.

Development

Although SAP is a commercial, off-the-shelf package, most organizations require extensive development efforts to manage an implementation. The development landscape is a fully functional environment but with a database that is a subset of the production data.

Training

End user training is a key part of any SAP implementation. Implementing SAP applications is a major change for a business and the critical nature of the application increases the demand for end user training. The training landscape is where new SAP users are introduced to SAP software.

Quality Assurance

Any change that is considered for a production SAP environment must be tested thoroughly before it is introduced to the production landscape. The QA landscape is where these tests are performed. The data in the QA environment is not a subset of the production data but a full set of that data. This protects against the possibility that a change that works in the limited test landscape might fail in the larger production landscape. A key issue is that performance tests are often done in the QA environment. If this is the case, the QA landscape must have identical server and storage configurations as production. The data in the QA environment is refreshed from production on a regular basis.

Production

The core of any SAP implementation is the production landscape. In the SAP world, this landscape is often called the PRD. All other landscapes are of secondary importance and any SAP customer expends a great deal of effort to ensure that its production landscape is available to users when it is needed. Not all organizations require 24-hours-a-day, seven-days-a-week availability for their production landscapes. In some cases, the business process is limited to one shift or a limited work week.

At the beginning of an SAP deployment, the functionality being implemented might be limited. For example, a company might decide to start by putting its human resources functions in an SAP environment before turning over other business functions. Although these limited deployments exist even after a full deployment, the majority of SAP implementations require 24/7 availability. Downtime for any production landscape must be very limited. Functions like system upgrades, hardware refreshes and system maintenance that require the production landscape to be shut down require extensive planning.

Other Landscapes

In addition to the landscapes that support and protect the production landscape, SAP deployments can also include several other landscapes, like data mining and disaster recovery. These can be as large as the production landscape and can also be of great importance to the business.

A variety of processes make use of production data. To ensure that nothing can impact the performance or availability of the production landscape, the data is replicated and a separate SAP instance provides for the required functionality.

Any organization with a mature SAP implementation needs a disaster recovery (DR) plan. These plans vary greatly from one organization to the next. Generally a DR strategy includes a remote SAP landscape.

Tier Four Storage Services

Hitachi Data Systems solutions for SAP extend the SAP model to a conceptual fourth tier. The strength of the SAP model is that the work load is spread across the tiers that enable extended growth and flexibility. To take best advantage of the features of Hitachi storage, perform storage-related functions in the fourth SAP tier.

When SAP is deployed on independent disks, database servers perform all of the storage management functions. Although this is possible in a small, non-critical implementation, it is not an optimal solution. In addition, because the database tier has limited scalability, it can become the source of performance issues. A more successful approach is to implement a conceptual fourth tier. This tier performs only those functions related to storage management.

This section addresses at a high level the value proposition of the Hitachi Data Systems fourth tier approach and the functions that are performed in that tier. The focus is on SAP environments and their potential problems and requirements.

Consolidation

In an SAP environment, the primary concern is ensuring the quality of service for the production landscape (PRD). To achieve this, production data is isolated from all other data. This is usually done by deploying a separate storage system that contains only the production data. Stringent change control procedures and isolation for the production landscape guard against some unforeseen event in the non-production environments affecting production processes.

Earlier in the evolution of storage technology, the challenge was for a single storage system to scale large enough to manage a complete PRD landscape. The scalability of the current storage technologies is such that multiple SAP landscapes can easily be maintained in a single storage system. For many SAP customers, based on storage management and floor space requirements, a single storage system is the preferred solution. In addition, data must be able to be moved across landscapes and a common storage system greatly simplifies that process,

The issue that prevents consolidation of multiple storage systems is not capacity or performance but rather concern about quality of service for the PRD landscape, but Hitachi storage systems address this concern with storage partitions and multi-protocol connectivity.

Storage Partitions

The ability to partition Hitachi Universal Storage Platform® V and Universal Storage Platform VM storage systems resolves the need to consolidate the landscapes and at the same time guarantees the quality of the production landscape. Hitachi enterprise storage systems can be divided with partitions for legacy applications and multiple landscapes, for example PRD, QA and miscellaneous landscapes.

The partitioning capabilities of the Hitachi Universal Storage Platform V and Universal Storage Platform VM ensure that the QoS for the PRD landscape is not affected by any event in other environments. This feature enables you to take advantage of the benefits of a consolidated environment.

Although the partitioning capabilities of the Hitachi Adaptable Modular Storage 2000 family are not as complete, the ability to partition cache and isolate data within 2000 family storage systems can provide a similar value for customers with only modular Hitachi products.

Multi-Protocol Connectivity

A second key to consolidation in the SAP space is the ability to service multiple I/O connectivity requirements from a single storage system. Often legacy data is mainframe based and co-exists with the SAP environment for an extended period of time. In addition, businesses often have other connectivity requirements that can be addressed by a consolidated storage system.

Performance

Hitachi Data Systems experience is that approximately 60 percent of all performance issues in SAP deployments can be traced to the storage infrastructure. When a storage-related performance issue arises, it's tempting to dive into the storage infrastructure to isolate and resolve the problem, however, that is a tactical approach that does not address the strategic error that is the root of the problem. A successful strategy is to design an environment that consistently delivers high performance and can grow as the SAP implementation evolves.

Data Levels

The diverse types of data profiles in an SAP environment almost always dictate that a one-size-fits-all approach is not appropriate for selecting storage technology. Implementing the highest performance and most costly storage might work but it is never the most cost-effective method. A better strategy is to create a solution with a systematic two-step approach:

1. Define data levels.
2. Create storage tiers to match the data levels.

The first step is to define data levels determined by the characteristics of the data. Data levels group data by I/O characteristics, retention needs and more. After you create data levels, define storage tiers that match the requirements of the data levels.

When defining data levels, the primary focus is the production environment. Your first concern is taking care that the production environment meets performance requirements. Within the production landscape, it's important to distinguish between extremely high I/O requirements for data such as log files and temporary database files and the high I/O requirements for database files.

Table 1 provides data level definitions and purposes. These definitions are a rule of thumb and can vary according to each customer's requirements.

Table 1. Data Levels

<i>Data Level</i>	1	2	3	4	5
<i>System Type</i>	Production, business continuity (BC), quality assurance High I/O files	Production, business continuity, quality assurance	Development, test	Backup, training	Archive, file servers
<i>Performance Requirement</i>	High	High	Medium	Low	Low
<i>Availability Requirement</i>	High	Medium	Medium	Medium	Medium
<i>Scalability Requirement</i>	High	Medium	Medium	Medium	High
<i>Retention Requirement</i>	Medium	Medium	Medium	High	High
<i>Storage Type</i>	SAN	SAN	SAN	NAS	NAS or SAN
<i>Data Activity Level</i>	Active	Active	Active	Active	Inactive

Storage Tiers

Storage tiers are developed by combining various storage attributes. Following is a typical list of attributes that are combined to define storage tiers:

- Storage system type
- Drive type
- Drive capacity
- RAID technology
- Internal storage
- Externally attached storage

Note that in some (if not most) instances, a tier strategy might be exclusively internal with no externally attached storage. In addition, some implementations are modular based. These determinations are made on implementation-by-implementation basis.

Thin Provisioning

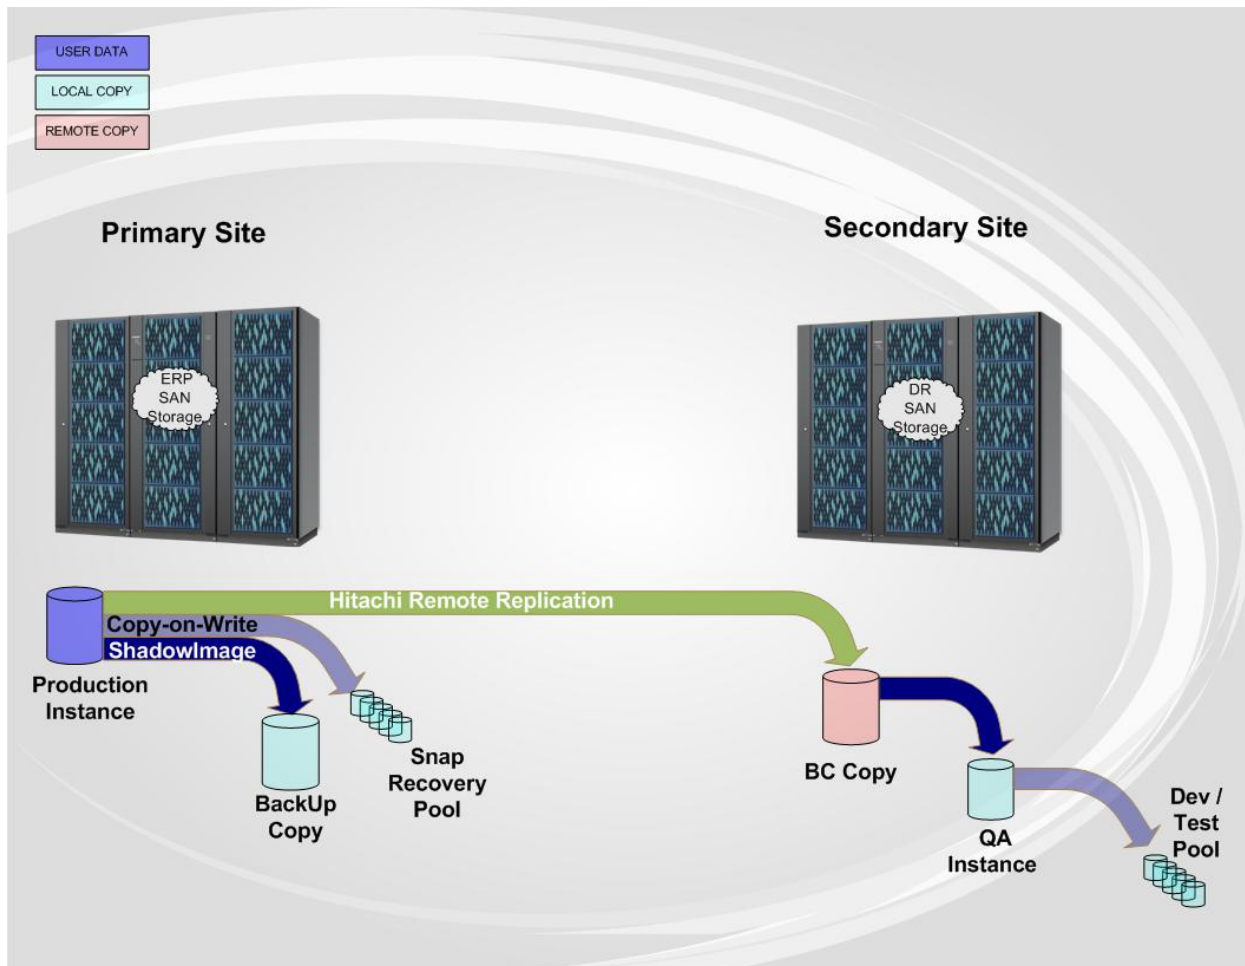
The generally accepted best practice for SAP is to allow the storage system to stripe the data across as many spindles as are available. Thin provisioning is the latest advancement of that technique and can significantly improve performance. Data is allocated across the pool as it is written, which spreads the I/O load across all of the spindles in the pool and delivers improved performance.

Data Replication

The ability to replicate SAP data at the storage level is the most valuable advanced technology in the storage infrastructure. The primary value is that it helps ensure the quality of service of the production environment.

Figure 3 illustrates replication solutions for an SAP deployment.

Figure 3. Replication Solutions



This replication solution offers the following features:

- Primary site:
 - Local Hitachi ShadowImage® Replication software copy for backup prevents the backup process from affecting production performance.
 - Hitachi Copy-on-Write Snapshot software volumes can provide rapid recovery from data corruption.
 - Remote copy enables business continuance.
- Secondary site:
 - ShadowImage copy for QA allows QA servers to be dual purposed as business continuance servers.
 - Copy-on-Write volumes are used in development environments.

To isolate the production landscape from secondary processes such as backup, create a full copy of production landscape data and perform the secondary function on the replicated data. Although data can be replicated both locally and remotely at the server level, storage system-based replication is commonly accepted as the preferred approach.

Backup Copies

The first use of data replication is to enable backups that do not affect the production landscape. The backup process frequently causes access contention issues with production landscape processes. To alleviate this, a full copy of the production landscape is created that the backup server accesses.

Landscape Refresh

The ability to move data from the production landscape to other landscapes is important in SAP deployments. Following are examples of landscape refresh scenarios:

- **QA** — The QA environment requires a refresh from the PRD environment on a regular basis. Utilizing the replication technologies of the storage system is the fastest and most effective method of copying all of the PRD data into the QA environment
- **Data mining and data warehousing** — Frequently, a multitude of other processes require a copy of production landscape data. It is not unusual for an SAP deployment to use seven or eight copies of the production landscape data.

In non-PRD environments where multiple copies of data are needed for functions such as testing, Hitachi Copy-on-Write Snapshot software can significantly reduce costs and improve functionality.

Business Continuance

Despite the mission-critical nature of SAP deployments, many organizations fail to address business continuance early in the deployment process. Frequently, the effort to simply get the SAP deployment into production consumes so much of an organization's capabilities that business continuance is relegated to a future effort. A common mistake is to put in place a process that does not meet recovery point objectives (RPOs) and recovery time objectives (RTOs). Storage system-based replication is a key technology in most SAP business continuance plans.

Data Migration

The fact that most SAP implementations require 24/7 availability implies that any disruption, planned or unplanned, is a major concern. Traditionally, hardware or software upgrades require significant planning efforts. Critical upgrades can be delayed significantly if a sufficient downtime window is unavailable. Replication and migration solutions from Hitachi dramatically reduce the planning and downtime required to perform these upgrades. When combined with the ability to integrate external storage into the infrastructure, most migrations can be done with no effect on application availability.

Storage Management

A key concern in SAP deployments is the growth of the storage infrastructure with no increase in the storage management staff. A majority of the repetitive efforts of the storage team are focused on storage allocation. Hitachi Dynamic Provisioning software provides critical technology to simplify the storage allocation process

Archiving

Businesses that use SAP face ever-increasing reliance on unstructured content, growing database sizes with resulting performance impacts and data retention requirements for regulatory and governance compliance. The Hitachi Content Archive Platform is the ideal storage foundation for addressing these challenges. It offers these key benefits:

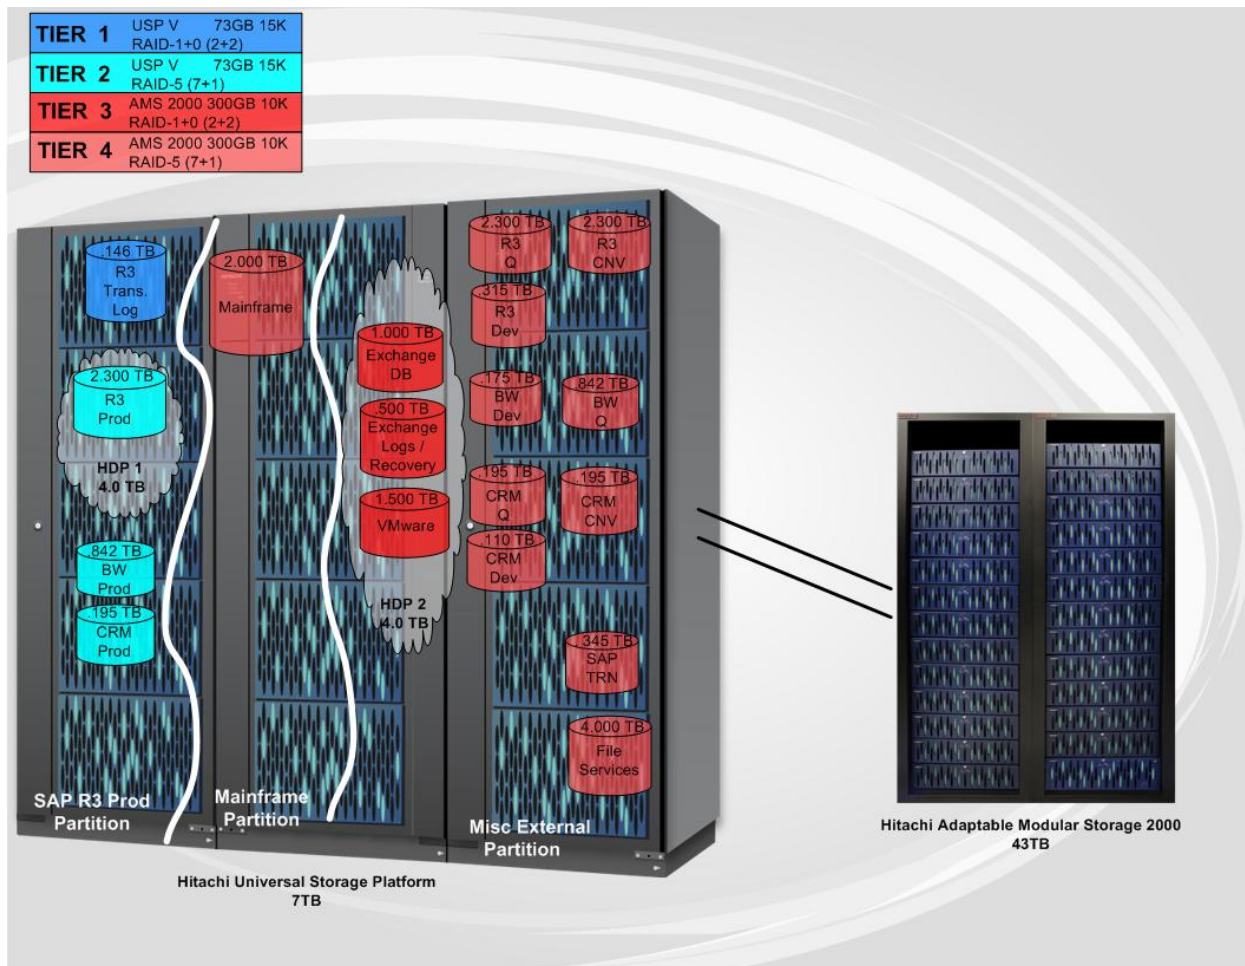
- Faster and less resource-intensive system upgrades
- Integration of SAP and non-SAP documents into SAP processes, allowing faster access and enabling improved customer service
- Improved compliance and litigation readiness
- Reduced TCO by archiving of aged content out of the database
- Improved enablement of customer and vendor self service
- Reduced TCO because it is easier to manage and scale than optical and alternative disk-based offerings such as network attached storage (NAS) or content addressed storage (CAS)

Putting it All Together

The technologies that Hitachi Data Systems brings to any SAP deployment can deliver significant benefits beyond simple data storage. This is true regardless of where in the cycle of SAP deployment you might be. Even if your business is at the beginning of the SAP deployment process where the immediate need might be a Hitachi Adaptable Modular Storage 2000 family system for development, that modular system can be part of the long-term deployment plan that can grow with no disruption or hardware obsolescence.

Figure 4 illustrates a sample optimized SAP architecture. Note that although the storage solution involves a Hitachi Universal Storage Platform system, the majority of the storage capacity is housed on an attached Hitachi Adaptable Modular Storage 2000 system.

Figure 4. Sample Optimized SAP Architecture



This architecture has the following features:

- Virtual partitions ensure QoS for critical environments
- Data levels and storage tiers are defined by data requirements
- Hitachi Data Provisioning software pools provide performance and management simplification
- External storage ensures cost management

Conclusion

SAP deployments are complex and have significant business impact whenever they are deployed. Understanding SAP's three-tier architecture and Hitachi storage as a conceptual fourth tier is a key step in the storage deployment process. The concepts introduced in this paper can help businesses maximize the TVO for their storage investments. Implement these concepts to leverage your storage system investments to provide additional data management functionality.

For more information about Hitachi storage systems and software, see the Hitachi Data Systems [Web site](#), your sales representative or your channel partner.



Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

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 or Web site 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. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

© Hitachi Data Systems Corporation 2009. All Rights Reserved.