

Designing Storage Tiers

Application Optimized Storage™ Solutions from Hitachi Data Systems

A White Paper

By Hugh Ujhazy

June 2005

Executive Summary

Whether within a Fortune 100 enterprise or a small- to medium-sized business, IT applications, and the storage environments that they depend on, are now critical to success. The processes and decisions that drive organizational growth and profitability, and minimize corporate risk, rely on the close alignment of IT infrastructure and business goals. The new challenge for IT is to create an infrastructure that supports the ongoing alignment of storage system capabilities and business application needs.

Application Optimized Storage™ solutions from Hitachi Data Systems provide the answer to this new business challenge. With Application Optimized Storage solutions, organizations are able to strategically align business applications and the storage infrastructure to reduce costs, boost performance, improve availability, and enhance functionality. A key component of an Application Optimized Storage solution is a virtualized, multitiered infrastructure, based on storage from Hitachi Data Systems and other vendors.

This paper describes a process for matching applications to designated tiers of storage. Each storage tier identified maps to a desired level of service quality required by the enterprise. Using this process, the storage manager can plan an optimal arrangement of storage platforms to meet the availability, performance, and scalability service-level requirements of business applications.

In addition to a tiered storage infrastructure, organizations can leverage the comprehensive suite of Application Optimized Storage solutions from Hitachi Data Systems. This suite of services, software, and hardware solutions can manage and optimize the storage environment, address the storage needs of applications, and further the close alignment of business and IT goals.

Contents

The Challenge	1
Mandate 1: No Time for Downtime	1
Mandate 2: Squeezing Every Last Dollar	3
Mandate 3: Size without Scale	4
Planning for Tiered Storage	5
Define Storage SLAs	5
Define Application SLA Needs	6
Define Tiered Storage Model	8
Application Optimized Storage Solutions.....	8

Designing Storage Tiers

Application Optimized Storage™ Solutions from Hitachi Data Systems

A White Paper

By Hugh Ujhazy

The Challenge

With IT budgets under scrutiny and continual demand for application service-level improvements, CIOs are feeling the pressure to do more with less. The business value of every dollar invested in IT must now be clearly demonstrated, as the focus of attention moves to ensuring scarce budget resources are deployed in alignment with the strategic priorities of the organization.

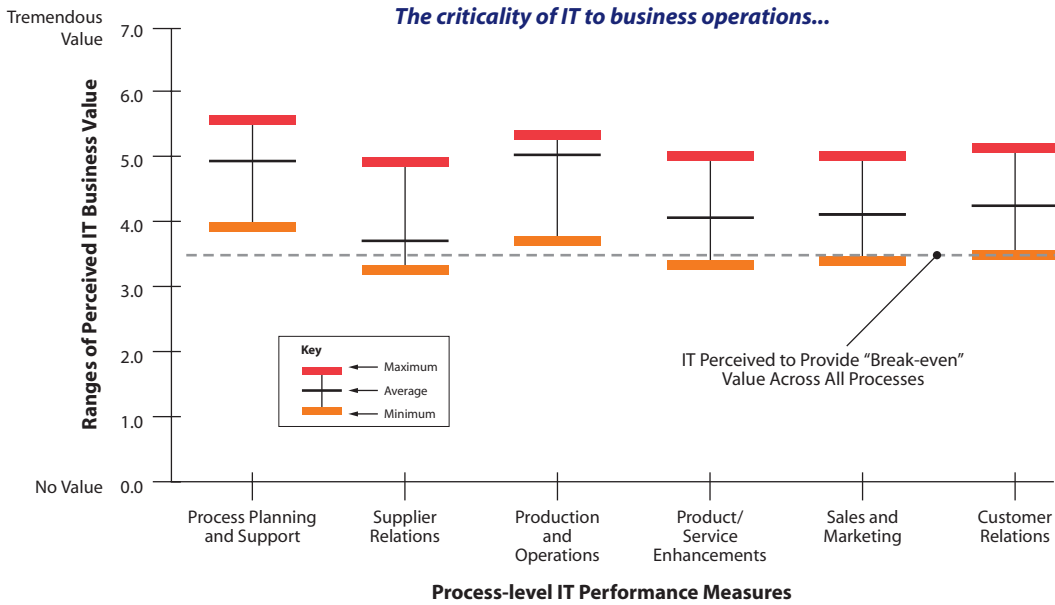
Driven by increasing business demand for IT-enabled solutions, technology infrastructures have grown rapidly and now command a large share of the corporate IT budget. While the majority of companies have successfully centralized their IT infrastructures, executives still face pressure to develop new strategies that address the unrelenting business demand for new services, higher availability, and larger transaction volumes. This has led to the establishment of three principle mandates.

Mandate 1: No Time for Downtime

With core business functions relying heavily on IT for support, the infrastructure that delivers application services is under tremendous pressure to improve systems availability and resilience, as shown in Figure 1. Although the corporate mandate for reliable service delivery is not new, two acute factors recently emerged to highlight system vulnerabilities:


- :: A flood of new security threats demonstrated the inadequacy of current risk management strategies.
- :: Rapid increases in the scale and scope of infrastructure services revealed the shortcomings of non-standard, ad hoc operational processes.

Figure 1. Improving Service, Availability, and Resilience



n = Most senior executives in 300 large companies

...evident in the high cost of downtime...



...heightens the need for impeccable infrastructure service integrity

No Room for Error

"Today, so much of our business is dependent on IT infrastructure. Our critical business applications have to be supported around the clock, and we have to make sure the infrastructure they run on is in excellent shape every second of the day. There is usually no room for error."

Source: Tallon, Paul, Kenneth L., Kenneth L. Krammer, and Vijay Gurbaxani, "Fact or Fiction: The Reality Behind Executives' Perceptions of IT Business Value"; "How Safe is the Business?" META Group, 2002; IEC research. (Figure redrawn from original.)


It is critical to keep business processes moving. The cost of letting operations falter or become hampered in any way as a result of downtime can be devastating to businesses.

The bottom line is that organizations rely heavily on their IT infrastructures to run their businesses, and IT must build, maintain, and optimize the infrastructure to support all business-critical application requirements.

Mandate 2: Squeezing Every Last Dollar

Despite a return of business optimism, many companies continue to pursue aggressive cost-containment initiatives as a means to improve profitability. While a focus on cost cutting over the last few years has yielded substantial unit cost reductions from the infrastructure, research suggests that unchecked business demand for increased service quality, volume, and functionality is driving up total infrastructure spending.

Figure 2. Identifying the Next Frontier of Sustainable Cost Savings



Note: Numbers indicate average range of savings reported from each initiative.
*Assumes no increase in scope, volume, or quality of infrastructure services.

Time to Partner

“We have gotten our unit costs as low as they can go every year, but the increased demand continues to drive costs up. To be able to provide year-over-year cost reductions, we need the users to make intelligent decisions about volume and quality of services consumed.”

SVP of IT
Technology Company

Source: IEC Leadership Survey 2003; Network World Fusion; Computer World; Morgan Stanley CIO Survey Series; Business Wire; Information Week; Optimize; Cahners In-Stat; Gartner; CERT; SNIA; Compass Group; IEC research (Figure redrawn from original.)

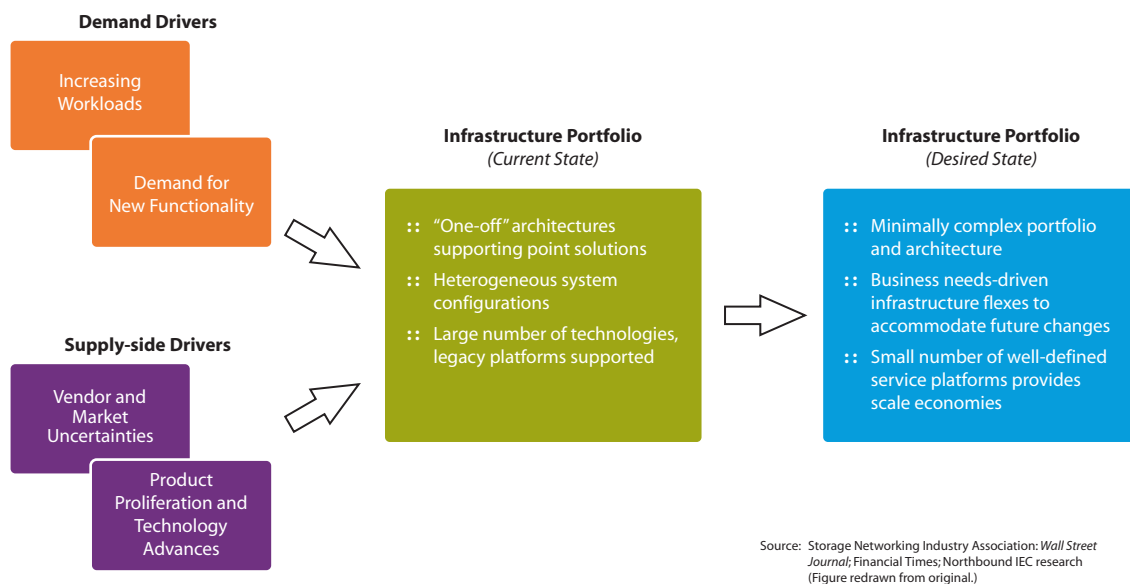
As demand for IT services outpaces unit cost reductions achieved through supply-side efforts, managing demand for infrastructure services emerges as the next source of cost savings.

To accrue further cost savings, the application delivery infrastructure must have the ability to partner with business customers to manage demand growth, as shown in Figure 2. At the same time, new, unexploited savings opportunities, based on the performance, availability, growth, and protection lifecycle requirements of applications, must be identified.

Mandate 3: Size without Scale

Rapid infrastructure expansion in the late 1990s left many organizations with complex technology portfolios containing numerous legacy platforms connected by a network of middleware and point-to-point solutions. The heterogeneity of these portfolios has imposed significant maintenance and licensing costs on these businesses, and it is crippling each infrastructure’s ability to rapidly provision new services.

Figure 3. Building a Flexible, Business-relevant Technology Portfolio



Infrastructure must create a service portfolio that can adapt to changing business and technology marketplace conditions while managing existing systems complexity.

With rejuvenated world economies, businesses are once again looking to expand, and IT departments are under pressure to streamline their existing technology assets, yet ensure service-level agreements (SLAs) are met. In order for future technology provisioning to be responsive, decisions must be informed by the needs of the business and by long-term strategic goals. (See Figure 3.)

Planning for Tiered Storage

A critical foundation for addressing these three mandates is the design and deployment of a tiered storage infrastructure. By understanding the specific needs of each business application and how these map to a storage tier's capabilities, organizations can sharply reduce their storage infrastructure total cost of ownership. There are several steps involved in the planning and implementation of a tiered storage infrastructure.

Define Storage SLAs

The first step is to inventory and understand the classes of data owned by the enterprise. Such an inventory can be performed at a gross level and can reflect the tiers proposed by IDC in their 2004 data classification study as shown in Figure 4.

Figure 4. IDC Data Classification Study—Recommended Data Classifications

<p>Mission-critical Data</p> <ul style="list-style-type: none">:: Most valuable to an enterprise, high access:: High performance, high availability, near zero downtime, highest cost <p>Business-critical Data</p> <ul style="list-style-type: none">:: Important to the enterprise, average cost:: Reasonable performance, good availability, less than eight-hour recovery <p>Accessible Online Data</p> <ul style="list-style-type: none">:: Cost sensitive, low access, often compliance of fixed content:: Online performance, high availability, less than eight-hour recovery <p>Nearline Data</p> <ul style="list-style-type: none">:: Cost sensitive, low access, large volumes:: Less than one-hour access time, automated retrieval <p>Offline Data</p> <ul style="list-style-type: none">:: Archived data, backup or compliance related:: Very cost sensitive, limited access, ~72-hour seek time
--

Source: IDC Classifications 2004

Taking an inventory of the types of data in your organization is the first step understanding the storage service levels that are needed to serve that data.

The five classifications derived by IDC focus on the availability and recoverability of the data within the enterprise. Each data type can be assigned a service level, as shown in Table 1. And additional sub-classifications can be created, based on application demand for snapshots, improved business continuity, and differing forms of backups.

Table 1. Storage Tier Metrics

Tier Name	Mission Critical	Business Critical	Accessible Online	Nearline	Offline
Availability	99.999%	99.999%	99.99%	99.90%	Offline
Performance	5	4	3	2	1
Maximum Backup Time (hours)	0.02	0.02	3	0.02	N/A
Maximum Local Recovery Time (hours)	1	4	6	4	1
Local Recovery Point (hours)	24	24	24	24	N/A
Remote Recovery Time (hours)	1	72	72	72	72
Remote Recover Point (days)	1	7	7	30	N/A
Offering Cost/GB (relative)	100%	80%	60%	40%	10%
Point-in-time Snapshots	Yes	Yes	No	No	No

The five tiers of data detailed by IDC map to a variety of availability and performance metrics. Availability demands may decrease or increase across different tiers, along with the urgency to recover to a particular point in time and the timeliness of that recoverability. These metrics make up the SLA for each tier. Each enterprise will define SLAs appropriate to the needs of their own applications.

A cost metric can also provide IT managers with both a method of evaluating vendor performance and a method of controlling storage costs within the enterprise. For example, if a tier is defined as Accessible Online, storage managers can determine that solutions in this category must not exceed 60 percent of the cost of storage in the Mission Critical category.

Define Application SLA Needs

The second step in the process is to match the needs of each application to storage tier SLAs. Each application—whether mainframe or open systems—is assigned to one or more of the following environments:

- :: **Production.** This is the environment where users interact with the application and where live data is stored and updated. This environment drives the business side of the application.
- :: **Test.** Close to the production environment in platform, size, and data model, the test environment allows the business to evaluate the impact of changes before they are introduced to production.
- :: **Development.** This is an environment for raw code development and for customizing packaged applications. Availability of the development environment may be on par with production, as interruptions to coding can delay the implementation of business functionality and much-needed maintenance.
- :: **Stage.** This environment provides a holding area for conducting end-user training against replicas of the production environment, or it may be used to hold data prior to loading to decision support environments.

The specific requirements of each organization will vary, and if other application environments exist they can be added to the list as needed. Once this information is known a simple study can be undertaken to match the application environment needs with storage tier SLAs, as shown in Table 2.

Table 2. Matching Application Service-level Needs to Storage-tier Capabilities

Application Name	Platform	Production		Test	
		Tier Name	Capacity (TB)	Tier Name	Capacity (TB)
Customer Management	Mainframe	Mission Critical	2.5	Accessible Online	1
Company E-mail	Open	Mission Critical	4.5	Accessible Online	2
Enterprise Resource Planning (SAP)	Open	Mission Critical	8	Business Critical	2
Customer Order Management	Mainframe	Mission Critical	2	Business Critical	1
Company Web Environment	Open	Mission Critical	0.75	Accessible Online	0.25
Financial Consolidation	Mainframe	Business Critical	0.5		
Decision Support/Data Warehouse	Open	Business Critical	9	Accessible Online	2

The SLA associated with a given application will invariably depend on the environment in which it is running. The production e-mail environment, for example, is likely to be mission critical—loss of access to e-mail severely impacts the ability of the enterprise to conduct business. However, loss of the test or training e-mail environment will have a minimal impact on the enterprise for the first 24 hours. The test e-mail environment can be categorized as accessible online.

In this straightforward model, the application environment is defined according to the nature of the application and the platform on which it runs. This mapping is helpful when making decisions about storage tiers with application owners, since it removes the infrastructure components from the equation.

Once applications have been assigned one or more SLAs, the total for each tier is assessed. Data volatility—the rate at which data owned by the application changes—and application data growth patterns—increases in users, transactions, and results of regulatory demands—also factor into the final equation.

Define Tiered Storage Model

Once steps one and two have been completed, a tiered storage infrastructure model is developed. Performing an inventory of the existing enterprise storage environment and mapping the results to application SLAs, will provide an understanding of current tier candidates. This approach may be expanded to group hardware storage assets by their performance characteristics, creating pools of storage. For example, modular storage platforms using SATA might be grouped as a “Modular—Low Performance” tier rather than a specific instance of hardware.

For example, an enterprise may currently deploy a mixture of enterprise and modular disk components with a variety of interfaces—mainframe and open. Assessing each component and its ability to meet SLA requirements by tier helps storage managers understand how the systems can be mapped to a Hitachi TagmaStore™ Universal Storage Platform managed environment.

The SLAs defined earlier in the process can now be mapped to candidate storage platforms. The Hitachi Lightning 9900™ V Series enterprise storage systems and EMC Symmetrix systems owned by an enterprise may be deemed acceptable for SLAs that are Business Critical. Similarly, the Hitachi Thunder 9500™ V Series modular storage systems might be appropriate to the Accessible Online tier.

The application-to-SLA mapping performed in step two indicates to storage managers how much physical storage is needed at each tier level. Adding metrics such as data volatility and growth rates allows future storage needs for each tier to be forecast. And, an analysis of the current storage platform inventory will help in creating a consolidation and data migration plan. Initially, all existing storage platforms can be virtualized; then, over time, Hitachi Volume Migration software (formerly CruiseControl) can move data transparently to its appropriate tier, without disrupting applications.

Application Optimized Storage Solutions

The steps outlined in this paper provide a clear methodology for designing a tiered-storage infrastructure. However, the design and implementation is only one component, albeit a critical one, in helping organizations create a tiered-storage infrastructure that enables closer alignment of business and IT. Application Optimized Storage solutions from Hitachi Data Systems help organizations address the critical need for aligning business and IT by providing:

- :: A virtualized, multitier infrastructure that optimizes data storage cost across the entire data lifecycle, and is managed by a single family of software tools
- :: Universal data services for managing data backup, replication, migration, and protection
- :: A common platform supporting diverse application solutions, including messaging, databases, and file services
- :: Application-centric, quality-of-service (QoS) data delivery, based on performance, availability, cost, and protection requirements
- :: Consulting and technology services to assess current and future requirements, successfully deploy appropriate solutions, and maximize an organization’s return on investment

With the introduction of the Universal Storage Platform, Application Optimized Storage solutions have reached previously unattainable levels of performance, consolidation, and virtualization. Heterogeneous, externally attached resources can now be brought together in a single, common pool of storage. This allows support for application-centric QoS, storage area management, and universal data management.

The Universal Storage Platform also enables storage managers to logically partition storage resources. This ensures maximum application QoS, aids the movement of data across storage tiers when matching storage SLAs and application requirements, and provides enterprise-wide business continuity support. The Universal Storage Platform allows all storage resources to be consistently managed with a single family of software.

With Application Optimized Storage solutions at the center of an IT strategy, organizations large and small are able to directly address their need for tiered storage. A tiered storage infrastructure greatly enhances the organization's business value and efficiency, and supports consolidation, automation, and consistent storage and data management to reduce the total cost of storage ownership.

For further information regarding Application Optimized Storage solutions and the Universal Storage Platform, please visit www.hds.com

 **Hitachi Data Systems Corporation****Corporate Headquarters**

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd.

Lightning 9900, Thunder 9500, TagmaStore, and Application Optimized Storage are trademarks of Hitachi Data Systems Corporation.

All other company names are, or may be, trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express 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 sells and licenses its products subject to certain terms and conditions, including limited warranties. To see a copy of these terms and conditions prior to purchase or license, please go to http://www.hds.com/products_services/support/warranty.html or call your local sales representative to obtain a printed copy. If you purchase or license the product, you are deemed to have accepted these terms and conditions.

©2005, Hitachi Data Systems Corporation. All Rights Reserved.

WHP-189-01 June 2005