



Tiered Storage Design Guide

Best Practices for Cost-effective Designs

By John Harker

September 2010

Table of Contents

Executive Summary	3
Tiered Storage Design	4
It Starts with the Applications	4
Aligning Storage Tiers with Applications	4
Storage Options	7
RAID Types	9
Drive Types	12
SAS Drives	13
SATA Drives	13
Summary	15
Appendix A – Related Software from Hitachi Data Systems	16
Hitachi Basic Operating Software V and External Storage	16
Hitachi Dynamic Provisioning Software	16
Hitachi Dynamic Tiering Software	17
Hitachi Tuning Manager Software	18
Hitachi Tiered Storage Manager Software	19

Executive Summary

When you design storage systems to encompass multiple service levels and price points, you can deliver superior storage service levels for less cost. This paper introduces and explains a design methodology for developing tiered storage architectures. It then discusses the options and best practices for configuring storage tiers using available software facilities, disk drive types and RAID group levels to meet performance, reliability and cost goals.

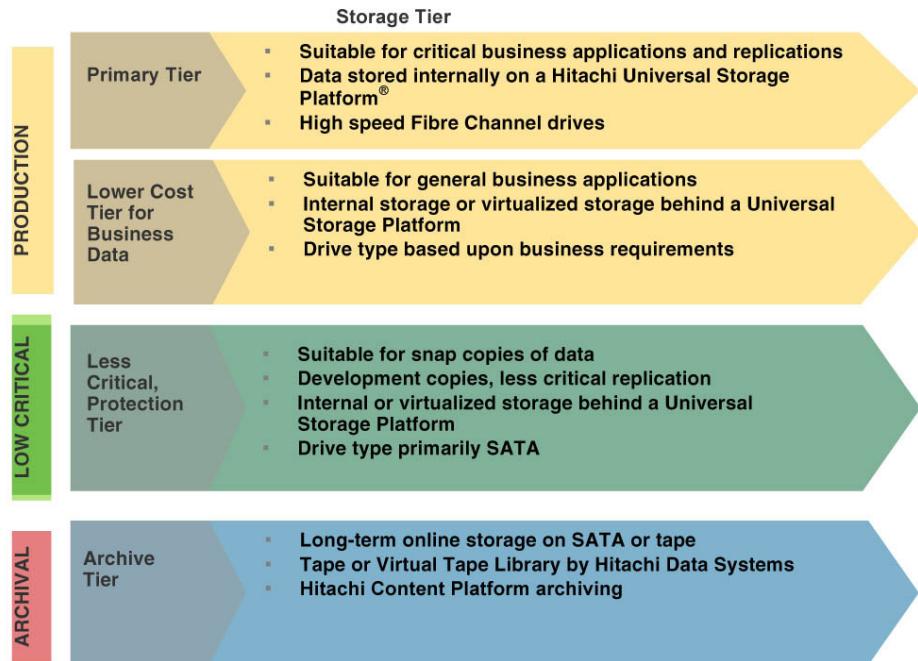
Tiered Storage Design

It Starts with the Applications

The design of a virtualized, tiered storage system starts with the applications. It is the business needs and applications that drive the storage requirements, which in turn guide tier configuration. Most applications can benefit from a mix of storage service levels, using high performance where it is important and less expensive storage where it is not.

But operationally it is not efficient to configure unique tiers for each application. Individually configuring a unique scheme for each application leads to extra work, cost and provisioning delays. Instead, the recommended practice is to develop a catalog of tiers with pre-defined characteristics and then allocate storage to applications as needed from the catalog. Figure 1 outlines a four-tier model; your individual requirement may call for more or less.

Figure 1. Recommended Practice: Develop a Catalog of Tiers



Aligning Storage Tiers with Applications

When designing storage tiers, start by looking at your application requirements in a variety of areas. The goal is to align each application to the right tier of storage with correct storage characteristics.

What are "storage characteristics?" As a starting point, individual consideration needs to be given to:

- **Availability.** How important is it that the data always be available? A retail storefront on the web might need guaranteed 100 percent uptime, while an archival system might only be needed during business hours.
- **Performance.** What is the I/O response time needed? How much bandwidth?
- **Cost.** Drive types and RAID configuration options offer a wide range of price points for different service levels.
- **Protection and Recoverability.** What are the RTO and RPO requirements? Is a simple nightly backup OK? Is replication necessary for backup window timing or business continuity? How many replication copies? Is distance replication needed?
- **Retention and Compliance.** How long does the data need to be kept on this tier? Note that if you back up and archive to a lower tier, you don't need elaborate retention controls on higher tiers.

It is useful to develop a specific map of the individual characteristics of each attribute for every tier as shown in Figure 2. These will be the storage service characteristics an application gets when assigned storage from that tier. When defining your storage tier characteristics, see how your applications map to those tiers. Your design should be able to accommodate the storage service level requirements of all applications.

One way is to develop this in an iterative fashion for the starting set of applications. Start by estimating values for a starting set of numbers or characteristics for each tier. Then examine each application in turn and optimize those numbers for that application. Then do the same for the next application to further adjust the numbers so they are optimized between this and the prior application. Continue the process until arriving back at the first application, which starts the cycle over again until converging on a design that meets the requirements. This is shown in the black box overlay of "Desired Attributes by Tier" in Figure 2. During this process you may discover that you need more or fewer tiers.

Figure 2. Estimated Values for a Starting Set of Numbers or Characteristics for Each Tier

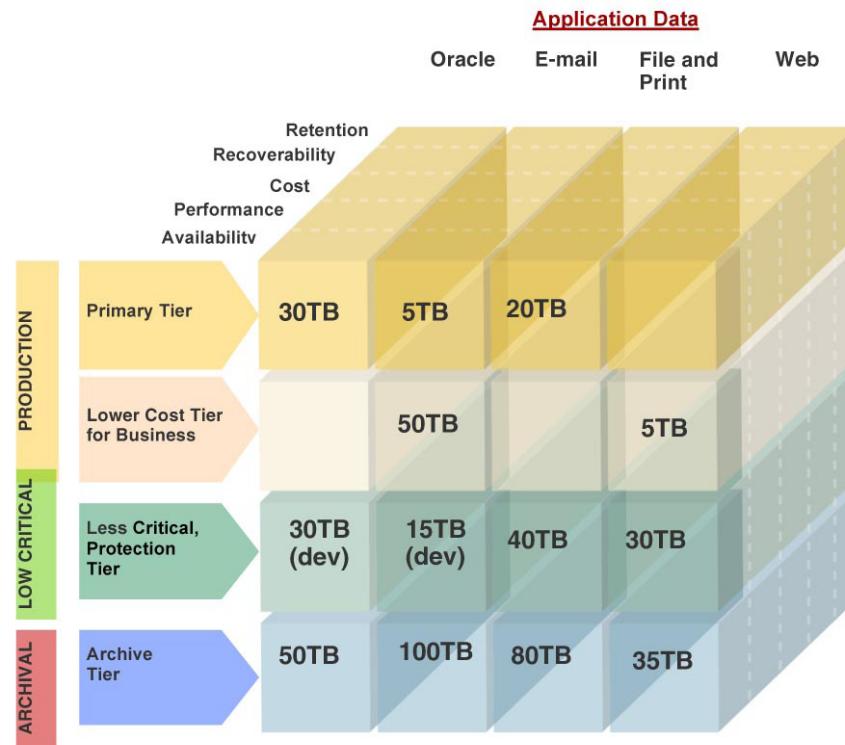
Data Tier	Desired Attributes by Tier					
	Availability	Performance	Cost (relative)	Recoverability	Retention	
PRODUCTION	Primary Tier	99.999%	5	\$\$\$\$\$	A	6 months
	Lower Cost Tier for Business Data	99.99%	4	\$\$\$\$	A	6 months
LOW CRITICAL	Less Critical, Protection Tier	99.9%	3	\$\$	B	1 year
	Archive Tier	99.9%	1	\$	C	7–100 years

It is worth noting that in Figure 2, for purposes of illustrating the layered design effort, tier characteristics such as "Performance", "Recoverability" and "Retention" are simplified. In actual design, each of these areas in turn faces a number of design options that are not fully captured with the linear values specified.

The last step is to map in storage for each application from the appropriate tiers as needed. Figure 3 provides an example of how this might look for a variety of applications.

For the company's mission critical Oracle business applications, primary tier storage is used. However, the backup mirrors and lower priority development copies are run on the more cost-efficient "Less Critical" tier. Old records in the Oracle databases are regularly archived to the "Archive" tier.

Figure 3. Storage Mapped for Each Application from the Appropriate Tiers



- The Microsoft Exchange email application needs to keep its log files on primary storage for performance, but the bulk of the storage for the mailboxes themselves can be mapped to the less expensive but still performing "Lower Cost" Tier for business data. A small amount of storage space is also mapped in from "Less Critical" for development purposes. With stringent retention policies and an expanding amount of emails with large attachments, a large amount of "Archive" Tier storage is needed.
- The NAS Head File and Print functions need some "Primary Tier" storage for several critical image processing applications. However, the bulk is file sharing used for shared directories within the company and print spooling and can use inexpensive "Low Critical" tier.
- Additionally, the company's web server uses the "Lower Cost" Tier for business data for the core set of often accessed pages. The bulk of what is online is infrequently accessed and can be kept on "Less Critical" storage.

Storage Options

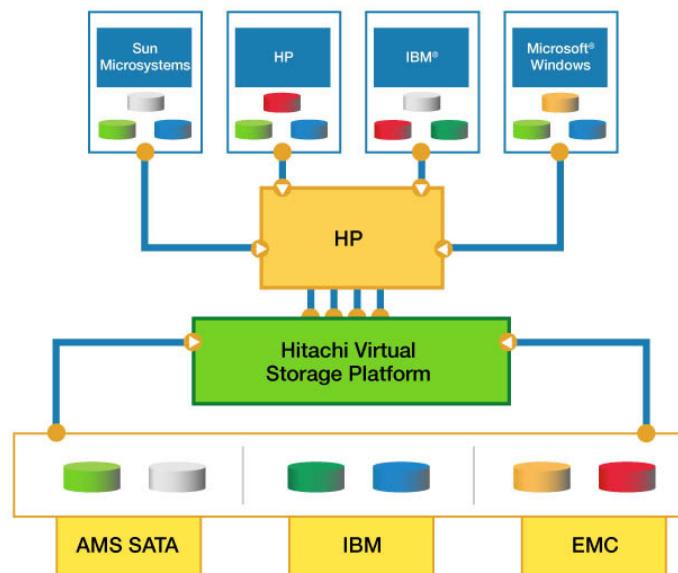
Now that we have designed our tiers from a requirements standpoint, how do you configure a system to match? There are a variety of ways to configure tiered storage architectures. You can

dedicate specific storage systems for each tier, or you can use different types of storage within a storage system for an "in-the-box" tiered storage system. The Hitachi best practice is to use the virtualization capabilities of the Hitachi Virtual Storage Platform (VSP) and the Hitachi Universal Storage Platform(r) (USP) family to eliminate the inflexible nature of dedicated tiered storage silos and seamlessly combine both. This allows for the best overall solution possible.

For example, for the highest tier you could start with a VSP configured with Fibre Channel drives and a high performance RAID configuration. Here the highest levels of performance and availability for mission critical applications are required. As a second tier you could add the USP with Fibre Channel drives, which are configured at a RAID level that is more cost-effective and still highly reliable but with a little less performance. As a third tier you could add even more cost-effective internal or external SATA drives.

The Hitachi storage virtualization architecture is differentiated by the way in which Hitachi storage virtualization maps its existing set of proven storage controller-based services, such as replication and migration, across all participating heterogeneous storage systems. See Figure 4.

Figure 4. Hitachi Storage Virtualization Architecture



Most environments will contain a mix of workload and applications within a single system. Thus, it is important to design a system where I/O load or activities (such as formatting, shredding, replicating or parity rebuilding) on a tier do not disrupt or degrade other tiers. In particular, it is important that a lower tier load does not affect higher tier performance.

In looking at performance requirements for designing or selecting a tier it is also important to look at the applications' use of that tier's storage in terms of read and write characteristics. How much is sequential read? How much random read? How much sequential write? How much random write? The mix of read and write characteristics is a major determining factor in the storage performance it will see. As an illustration of why, consider an I/O's data path inside the VSP:

- Random read hits are serviced at electronic speed by cache and do not reach the back end.
- Random read misses go through cache unaltered and go straight to the appropriate back-end disk drive.
 - This is the only type of I/O operation where the host always sees the performance of the back-end disk drive.
- Random writes
 - Host sees random writes complete at electronic speed; it only sees delay if too many pending writes build up.
 - Each host random write is transformed going through cache into a multiple I/O pattern that depends on RAID type.
- Sequential I/O reads and writes
 - Host sequential I/O is at electronic speed. Cache acts like a holding tank and the storage system back end puts [removes] "back-end buckets" of data into [out of] the tank to keep the tank at an appropriate level.

High performance drive type and RAID configurations will have a much bigger impact on a random read and write intensive application, versus one that does mostly sequential reads and writes. And proper cache sizing, I/O queue depth settings and cache partitioning can be equally important as drive type and configuration in tiered storage designs.

There are also a number of tiered storage design tools useful for configuring and managing storage tiers. These are software products, interfacing to and supported by your storage system. They provide enhanced levels of functionality that are used to optimize and tune all drive and RAID combinations.

As with all storage systems, when sizing it is necessary to keep in mind not only the amount of storage needed, but also application I/O characteristics and performance requirements. The final drive count may wind up being driven more by performance requirements than required capacity. Proper drive choices as well as the right configuration will provide the highest performance and will create a more cost-effective tiered storage environment.

RAID Types

An industry and Hitachi best practice is to always use some form of RAID configuration beyond RAID-0 as a first line of defense in the data protection war. In a tiered storage environment, you can take advantage of different (multiple) RAID configurations to provide varying levels of service and capacity optimization according to individual tier requirements.

There is a price for using a RAID configuration both in storage capacity used and performance. This price varies according to the type of RAID used. For example with RAID-1, whenever a host writes data to the storage system, at least two disks need to be updated. The amount of extra disk drive I/O activity needed to handle write activity is the key factor in determining the lowest cost solution combination of disk drive RPM, disk drive capacity and RAID type.

Particularly for random I/O, back-end disk I/Os are one of the finite performance factors. A drive head can only go so fast. For example, take a RAID-5 LUN and a RAID-1+0 LUN, issue four I/Os and observe how many back-end disk I/Os this generates (see Table 1).

TABLE 1. RAID PERFORMANCE COMPARISON

	Total Number of Disk I/Os
RAID-5	
Random Read	4
Random Write	16
Sequential Read	4
Sequential Write	5
RAID-1+0	
Random Read	4
Random Write	8
Sequential Read	4
Sequential Write	8

Most enterprise storage systems are configured with a mix of three types of RAID levels. The following sections discuss what they are and how they work.

RAID-1 and RAID-1+0 (Mirroring)

RAID-1 mirrors the entire disk. RAID-1 configurations have higher performance in most applications because each disk has a copy of the data. Thus, it is possible for the system to be responding to two I/O requests at the same time. In general, a mirrored configuration will perform as much as 30 percent better than a parity-protected (RAID-5) configuration.

It is a Hitachi best practice to use RAID-1+0 (sometimes called RAID-10) instead of RAID-1. RAID-1+0 is a variation on RAID-1 that adds increased performance through added striping.

Mirrored configurations also provide higher performance if there is a disk failure since another complete copy of the data is immediately available. They also offer the shortest rebuild time for the failed drive.

- Since RAID-1 and RAID-1+0 require doubling the number of disk drives to store the data, people tend to think of RAID-1 as the most expensive type of RAID.
- However, RAID-1 and RAID-1+0 offer the lowest "RAID penalty" of only having two disk I/Os per random write, compared to four I/Os for RAID-5, and six for RAID-6.
- And there is no penalty and a potential improvement for read operations.

For this reason, when the workload is sufficiently active and has a lot of random (as opposed to sequential) writes, RAID-1 or RAID-1+0 will often be the cheapest RAID type. This is because it has the least disk drive I/O operations per random write. Other RAID configuration will have to be used at low densities to meet the random write requirements.

RAID-5 (Parity Protected)

RAID-5 configurations are parity protected. In the event of a physical drive failure, the missing data is rebuilt by reading the remaining drives in the RAID group and performing parity calculations.

RAID-5 achieves redundancy with less disk space overhead, but at the expense of having a higher "RAID penalty" for random writes. There is also a larger performance degradation upon a drive failure (while the failed drive is rebuilt from the parity data).

- For sequential reads and writes, RAID-5 is very good.
- It is very space efficient (smallest space for parity), and sequential reads and writes are efficient, since they operate on whole stripes.
- For low access density (light activity), RAID-5 is very good.
 - The 4x RAID-5 write penalty is (nearly) invisible to the host, because it is nonsynchronous.
- For workloads with higher access density and more random writes, RAID-5 can be throughput-limited due to all the extra parity group I/O operations to handle the RAID-5 write penalty.

RAID-6 (Double Parity)

RAID-6 adds on to the single error correcting capabilities of RAID-5 with two types of parity correction that can continue running and rebuild the data even if two drives fail. It has generally the same read and write characteristics as RAID-5. It provides the highest level of data availability, albeit with more of an impact on write performance than RAID-5 due to increased overhead from the double error-correcting capability.

The extra protection of RAID-6 is particularly well suited for SATA drives, which have both lower reliability numbers and, due in part to their size, longer rebuild times. Because of this, RAID-6 is considered a best practice for SATA configurations.

A Note on Parity Groups

To ensure data integrity a RAID configuration of drives involves both storage for data and storage for extra information. This extra information is referred to as parity. For each of the RAID types there are different ways of laying this out. For example, for the USP the following are available:

RAID-1+0: 2D+2P (2+2)

RAID-5: 3D+1P (3+1), 7D+1P (7+1)

RAID-6: 6D+2P

The VSP and the USP family also offer the ability to create "concatenated parity groups," which are parity groups where the striping extends across multiple coupled physical parity groups. This extended striping brings more drives (spindles) to bear and improves performance.

RAID-1: 2 x (2+2), often called 4+4

RAID-5: 2 x (7+1) and 4 x (7+1)

There are a lot of ramifications behind selecting the correct parity group configuration. At a high level, smaller parity groups offer better performance at the cost of consuming additional disks. Larger parity group configurations sacrifice performance and rebuild times to reduce the number of physical disks needed. But this is only a rough guideline. Among other things the power of the platform matters.

For example, when 7D+1P was introduced many expected there would be an increased parity generation overhead. However, it turns out that because of the optimal efficiency in the Hitachi

enterprise storage architecture, the 7D+1P is a very efficient RAID group size. This compensated for the theoretical impact and gave a net result of equality between 2 x 3D+1P and 1 x 7D+1P.

For each fixed drive type, by using different RAID and parity group types, a number of configurations are possible, each offering different performance and reliability levels.

Now let's look at the drive types.

Drive Types

An increasing number of enterprise storage systems now offer options for hard disk drive types and sizes varying by performance, capacity and reliability. Like Fiber Channel drives, SAS drives are available in various sizes and rotational speeds. High capacity SATA drives are available for lesser tier applications such as low speed replication and archiving. And solid state drives (SSD) add a new high performance Tier 0 option.

In considering the performance impact of the various physical drive options, it is important to understand that higher speed drives perform better than lower in terms of seek time and rotational latency. Drives of the same rotational speed, latency and seek time will have roughly the same performance, regardless of capacity.

Table 2 illustrates the industry standard specified capacity (base-10), the actual (computer base-2) raw size and the typical usable (formatted) size of each type of disk. Also shown is the rule of thumb average maximum random IOPS rate for each type of disk when using most of the surface. The number and size of the LDEVs created per array group determine how much of its disks' surfaces are in active use.

TABLE 2. HARD DISK DRIVE (HDD) TYPES AND SPECIFICATIONS

HDD Type	RPM	Form Factor	Port Speed	Advertised Size (GB)	Usable Size (GB)	Nominal Random Read IOPS
2TB SATA	7200	LFF	3Gb/sec	2000	1832	80
600GB SAS	10K	SFF	6Gb/sec	600	545	130
400GB SSD	n/a	LFF	6Gb/sec	400	364	5000
300GB SAS	10K	SFF	6Gb/sec	300	268	130
200GB SSD	n/a	SFF	6Gb/sec	200	191	5000
146GB SAS	15K	SFF	6Gb/sec	146	137	180

The type and quantity of drives used in a VSP and the RAID levels chosen for those drives will vary according to analysis of the user workload mix, cost, application performance targets and usable protected capacity requirements.

Note that as the percentage of active data on the drive increases, the effective I/O rate per volume of data decreases. This effective I/O rate is referred to as access density, or IOPS/GB (I/Os per second per gigabyte). It decreases because the more the disk surface is active, the farther the disk

heads must seek. This creates higher average seek times. For example (using a 15K RPM disk), when only using the outer 25 percent of the surface the average read seek time may typically be about 1.8ms (providing about 267 IOPS), With 100 percent of the surface average seek time will be about 3.8ms (about 182 IOPS). For a SATA disk at 7,200 RPM, these values are about 4.3ms (119 IOPS) and 8.5ms (79 IOPS).

So it is important to consider not only the I/O characteristics you are optimizing for, but also the access density. The lowest cost configuration in terms of RPM, capacity and RAID type depends strongly on the tier's access density and read and write characteristics.

Another important design factor in selecting the size of the drive is recoverability. How long will it take to rebuild the physical drive from a spare if the physical drive fails? The answer is, for a given RAID type the bigger it is the longer it will take, which is an exposure. This is one of the factors influencing the RAID configuration selected. Options such as RAID-1+0 offer rapid recovery while RAID-6 offers an extra level of security. So consider disk recovery time in system design.

The following sections describe the performance characteristics of each drive type and give some use case examples and configuration recommendations.

SAS Drives

SAS drives are increasingly the primary drives used now in both enterprise and midrange storage systems. These are the workhorse drives, having both high performance and good capacities. They fall in between the random performance levels of SSD and SATA (much closer to the SATA end, though). All three drive types will have about the same sequential read performance per matching RAID levels. SAS drives are the same as Fibre Channel drives but with a SAS interface. SAS drives come in 15K RPM, 10K RPM, and 7.2K RPM in sizes currently ranging from 146GB to 600GB. The SAS interface speeds on all Hitachi supported drives is 6Gb/sec.

SAS drives are very high technology, having dual processors, dual host ports, large caches and a wealth of internal error detection or correction mechanisms. They will perform an automatic write check every ten minutes or so on a hidden track (per surface) to verify that the write heads are functioning properly. [Note: SATA drives do not do this, hence the read verify after write phase done on Hitachi storage systems.] The dual host ports allow two concurrent interactions with the host (storage system controller). One port may be receiving or responding to a new I/O command while the other is transferring data (only one such transfer per drive at a time).

The first SAS drives were of the usual large form factor size (3.5 in.). Now they are shifting to the smaller form factor (2.5 in.). This reduces cost and heat, as well as reducing seek times (smaller platters). In general, a 3.5 in. 15K RPM SAS drive and a 2.5 in. 10K RPM SAS drive have nearly identical performance. The 15K RPM variant of the 2.5 in. drives will disappear within the next two years.

SATA Drives

SATA drives offer very high capacities (now at 2TB per drive but with 4TB coming soon) at an "economic" level of performance. These are best suited for archival duty. These are not suitable for high levels of random workloads, such as online transaction processing (OLTP), with even modest (20 percent or so) sustained levels of write.

The decision on how to use SATA should be carefully considered since there are two options. You can use SATA drives internal to your enterprise storage system or use externally attached SATA storage. Which is best, often depends how much SATA storage you are configuring.

Due to the large capacities, most SATA drives are used in RAID-6 configurations. This use avoids the likelihood of a dual drive failure during the potentially long rebuild times of a failed drive onto a spare drive. However, the use of RAID-6 carries a very high RAID write penalty factor of six internal array drive operations per host write request. With SATA, there are three more such operations (read-verify) per write, or nine in all per host write request. You can trade off usable capacity for a large increase in write performance by the use of RAID-1+0 rather than RAID-6. RAID-1+0 only carries a write penalty of two drive operations per host write request. So the usable capacity is reduced from 80 percent (RAID-6 6D+2P) to 50 percent (RAID-1+0 4D+4D). But in many cases the cost per usable gigabyte is not as important as the cost per IOPS.

On Hitachi storage systems using SAS technology, the SATA disk canister has a proprietary logic board that creates two SAS ports from the single SATA port on each SATA disk. This interface also performs a read-after-write check on the last sector from the range of blocks (whether data or parity/mirror) just written to the disk. After an application block is written to the disk, the last sector of it is read from the surface and compared to the buffer in the canister logic board. While this reduces the write rate by some amount depending on the RAID level, this read from the disk does guarantee that the data was accurately written. SATA drives do not validate proper functioning of the write heads as do SAS disks. This is what drives the need for a read verify after write operation on SATA disks.

Other considerations on the use of SATA versus SAS disks when using RAID-6 include:

- No-load format times are about 18 times longer than with SAS disks.
- No-load corrective copy times are about 4.4 times longer than with SAS disks.
- No-load disk copy time is about 10 times longer than with SAS disks.
- There should be one spare SATA disk for every 16 SATA disks in use.

The Hitachi Data Systems recommendation on suitable uses for SATA volumes is to use them for near-line storage with a low frequency of access but online availability, such as:

- Email archives, backups, tape replacement
- Second or lower tier NAS file and print
- Second and third stage replication copies
- Second or lower tier primary or secondary replication pairs
- Historical data with limited access (medical records, etc.)
- Archived documents for regulatory purposes
- Sound or video files

It is generally a Hitachi best practice to configure a cache partition for the SATA disk drives, or for externally attached SATA to turn off USP caching. This restricts cache contention and prevents a potential performance impact on higher tiers.

On the Virtual Storage Platform and the Universal Storage Platform family the decision to use up valuable internal chassis disk slots with SATA disks, rather than the much higher performing SSD or

SAS disks, should be carefully considered. In many cases, the use of SATA disks within virtualized storage on an external midrange product might make better sense. Generally this works out to be that the more SATA storage there is, the more external systems are favored. For smaller amounts of SATA, it is more likely to be more economic to configure it within the primary storage system. However, there will be individual cases where the use of some internal slots for SATA drives will solve a specific customer business solution in a straightforward manner. It is not expected that a VSP would normally be configured with a high percentage of internal SATA disks.

Summary

Making the proper design decisions for tiered storage designs requires an understanding of what can be done and what the tradeoffs are. This paper has introduced and explained a design methodology for developing tiered storage architectures. It also portrayed some of the options and relative costs for configuring the system at the disk drive and RAID group levels.

With a broad selection of software, physical media and configuration options, including SATA, available in today's enterprise storage systems, you have the flexibility to design your environment for your organization's specific needs. Financially, the larger the system is, the greater the benefits that will be seen from internal and external SATA drives. For storage systems that are or will become large, the cost benefit can be dramatic.

Hitachi Data Systems offers a number of software products that add important design options to tiered storage systems. More information about this software and how it supports and enhances tiered storage environments can be found in Appendix A.

In order to get the maximum value from your storage system, it is vital to invest the time required to design the best possible solution to meet your unique requirements, whether they include capacity, reliability, performance and cost. This paper has just touched on some of these areas. We encourage you to work closely with your storage partner, who can offer advanced configuration and sizing tools and can assist in delivering the highest possible quality solution.

Appendix A – Related Software from Hitachi Data Systems

Hitachi Basic Operating Software V and External Storage

One of the unique abilities of the Hitachi Virtual Storage Platform and the Hitachi Universal Storage Platform family is to be able to connect and virtualize externally attached storage systems, creating a single pool of tiered storage. Enabled with Hitachi Basic Operating System V software (BOS V), this ability brings the capabilities of the VSP or USP to bear on a mult-tiered storage area network (SAN) of heterogeneous storage systems. It allows single point access from and to other old and new storage systems in the SAN.

BOS V allows the operation of multiple storage systems connected to a VSP or USP as if they were all one system with a common management toolset and software capabilities. Using BOS V, the shared storage pool composed of both internal and external storage volumes can be used with all Hitachi storage system-based software for data migration and replication, as well as any host-based application.

This software also provides the system designer with the option to reuse existing storage, flexibility and scalability for storage expansion. One of the design factors involving the use of SATA drives is how much growth is expected. In general, external storage becomes much less expensive than internal if there is going to be a lot of it.

Hitachi Dynamic Provisioning Software

Hitachi Dynamic Provisioning (HDP) software is an important option to consider when developing your tiered architecture. Hitachi Dynamic Provisioning is bundled with BOS V on the VSP and sold separately as a layered product on the USP. It is also bundled with the Hitachi Adaptable Modular Storage 2000 family.

HDP is thin provisioning software that permits the complete virtualization of a volume provisioned to an application. Actual storage capacity is assigned on an on-demand basis when data is written by the application. Actual storage capacity is collected in one or more HDP physical storage pools. The pools are created from multiple LDEVs from multiple array groups of any RAID level. Each HDP pool supports virtual host-accessible volumes referred to as HDP volumes (short form: DPVOLs). From an application's standpoint, an HDP volume looks no different from any other storage volume. As applications write data to HDP volumes, Hitachi Dynamic Provisioning software assigns actual capacity from the HDP pool to the HDP volumes. This just-in-time method means physical storage allocations remain available until they are actually needed.

Hitachi Dynamic Provisioning simplifies storage management by decoupling the provisioning of capacity to an application from the management of physical resources in the storage system. For example, physical storage is nondisruptively added as needed to the storage system, placed in centralized pools, and made available to new and existing HDP volumes.

HDP also simplifies performance optimization by transparently spreading many hosts' individual I/O patterns across many physical disks. You can create pools with the RAID level desired and use as many disks as needed to support the desired target IOPS rates for the DPVOLs connected to that pool.

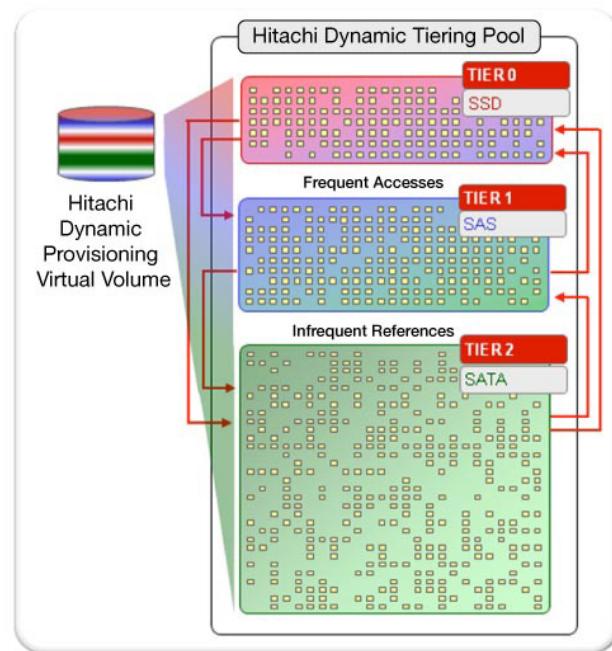
For example, you might create four database pools, with 16 LDEVs taken from 8 RAID-1+0 array groups (two LDEVs each), and use these for the database files that contain user data. These database pools are configured to support high levels of random small block OLTP workloads.

You might then create a log pool of eight LDEVs from four RAID-5 array groups (two LDEVs each) and use this pool for DPVOLs assigned to database logs or temp space. This log pool is configured to accept high levels of large block sequential write workloads.

Hitachi Dynamic Tiering Software

Hitachi Dynamic Tiering (HDT) software provides fine grain tiering within virtual volumes (See Appendix A — Figure 1). It automates the complexities and overhead of implementing data lifecycle management and optimizing tiered storage. HDT eliminates the need for time consuming manual data classification and movement of data to optimize usage of tiered storage.

Appendix A — Figure 1. Fine Grain Tiering with Hitachi Dynamic Tiering Software

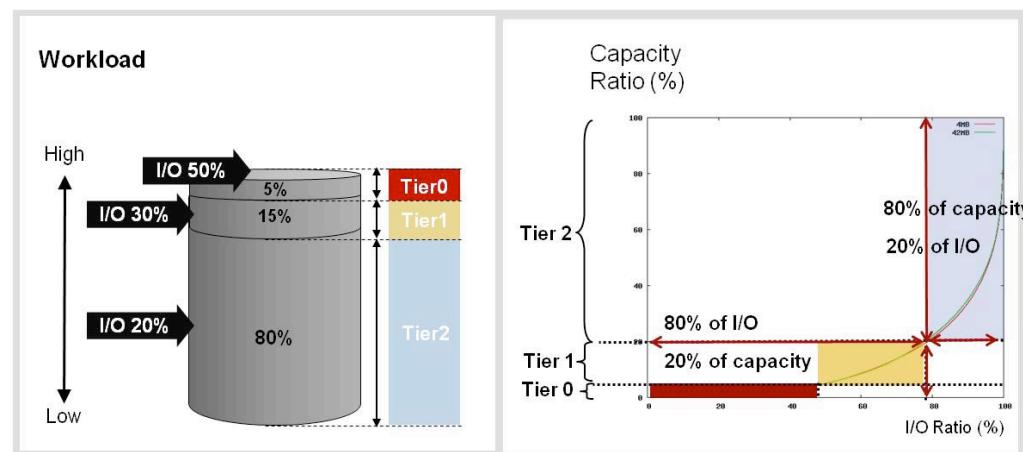


No elaborate decision criteria are needed; data is automatically moved according to simple rules. One, two or three tiers of storage can be defined and used within a single virtual volume using any of the storage media types available for the Hitachi Virtual Storage Platform. Tier creation is automatic, based on user configuration policies, including media type and speed, RAID level and sustained I/O level requirements. On a periodic basis using ongoing I/O activity analysis as a heat index, the data in the volume is moved to the most appropriate internal tier. The most active data moves to the highest tier and less active data moves to the lowest tier. The data is moved on a fine grain page-based level so system overhead is minimized and only the right data is moved. During

the process, the system automatically maximizes use of storage, keeping the higher tiers fully utilized.

Typically, over 80 percent of data access activities involve less than 20 percent of your storage (see Appendix A — Figure 2). This pattern can be applied with Hitachi Dynamic Tiering to both improve performance and reduce storage costs. Dynamic Tiering uses Hitachi Dynamic Provisioning technology and thus inherits its advantages of simplified provisioning, capital savings and self-optimizing performance as well.

Appendix A — Figure 2. Typical Pattern of 80% Capacity Serving only 20% of I/O

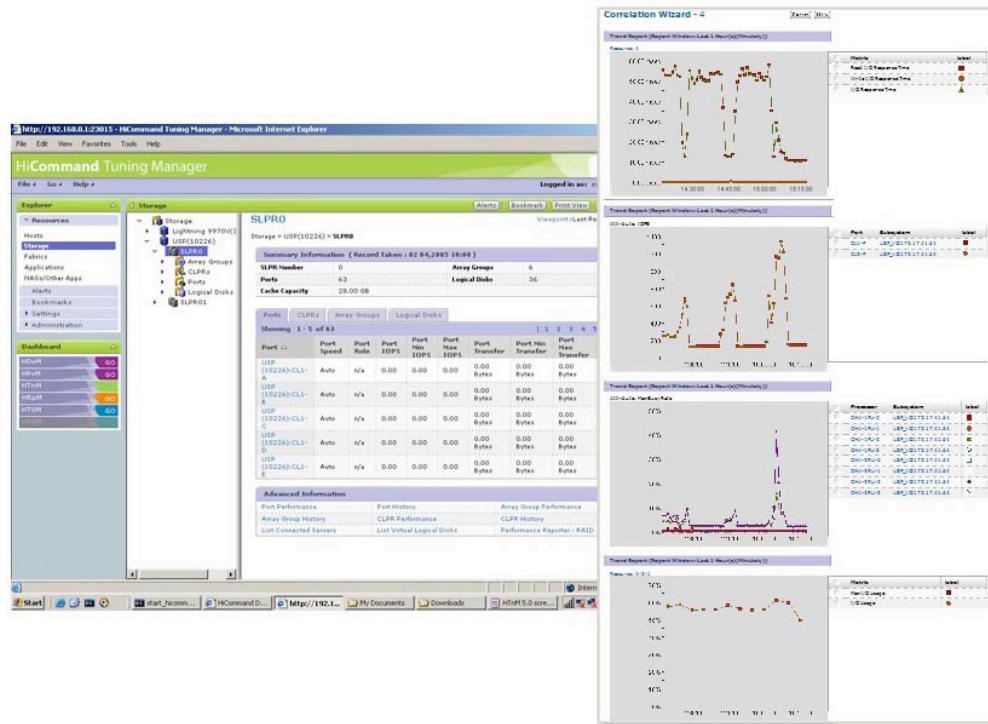


Hitachi Dynamic Tiering also takes the guesswork out of deploying expensive solid-state drives (SSDs) as a Tier 0 data accelerator. Because it moves only the most active data to the highest tier, SSD usage is automatically optimized, and even a small percentage of SSDs can dramatically improve overall performance. Also, SSD technology requires less power than disks and, when combined with a higher percentage of energy-efficient SATA drives, can reduce drive count and improve overall system power consumption.

Hitachi Tuning Manager Software

Tuning Manager is an optional Hitachi Command Suite product .It can deliver real-time and historical performance and capacity information useful in defining and monitoring individual storage tier QoS to administrators and other products such as Tiered Storage Manager. Tuning Manager can also feed real-time performance data to Tiered Storage Manager for improved tier definition and management (see Appendix A — Figure 3).

Appendix A — Figure 3. Hitachi Tuning Manager User Interface

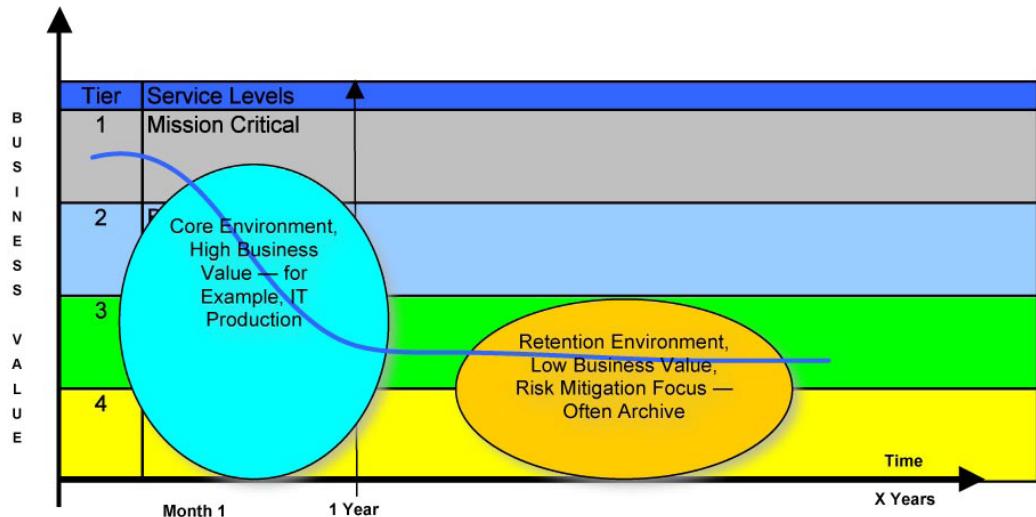


Hitachi Tiered Storage Manager Software

Tiered Storage Manager is an optional Hitachi Command Suite product that provides easy tier definition and dynamic data movement capabilities for all storage in the USP pool - both internal and external. It enables administrators to easily and interactively identify available storage to match application quality-of-service requirements and then define storage tiers based on this information. As such, it is an excellent tier management tool. It also provides the ability to migrate or schedule migrations of online data within or between tiers without disrupting application access.

It is important to be able to nondisruptively move a running application's data between tiers in a tiered storage environment. Information value changes over time, driving content movement requirements across service levels to realize an optimal cost balance (see Appendix A — Figure 4).

Appendix A — Figure 4. Data Lifecycle of a "Transaction"



Tiered Storage Manager is also an important option for management flexibility. The ability to nondisruptively migrate the data to a different tier, drive or RAID type is important for rapid response to new requirements, rapid deployments or data migrations. Storage can be allocated to an application, observed and then nondisruptively changed if needed.

Tiered Storage Manager is a useful tool for data migrations off old equipment, which can be virtualized behind a USP. The data can be nondisruptively migrated to another system or to storage internal to the USP.



Corporate Headquarters

750 Central Expressway
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

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 in this document or website 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. WP-310-B DG September 2010