Enterprise Storage

Hitachi Freedom Storage™ Thunder 9200™ Architecture and Performance Configuration Guidelines

by
Gil Rangel
Bruce Bush
Open Systems Performance Consultants
Storage Performance Management Group
Hitachi Data Systems Solutions Engineering
San Diego, CA

Contents

Preface

The intent of this document is to familiarize both the sales and technical support staff, as well as customers and value-added resellers, of Hitachi Data Systems with the performance capabilities of the Hitachi Freedom Storage™ Thunder 9200™ architecture. Please keep in mind that this information is provided as a guideline. Individual implementations will vary due to differences in OS architecture, application design, SAN implementation, and other factors. All of the initial testing was done using Sun® servers performing I/O through the raw disk interface. Information will be added to reflect additional server configurations, database workload simulations, Thunder 9200 hardware changes, and microcode updates.

The people who will benefit most from this document are those responsible for advising or implementing the hardware and data layout phase for an application or database system—those who require the best performance from the Thunder 9200 storage system. A working knowledge of storage arrays and configuring our existing midrange products is assumed.

Hitachi Freedom Storage™ Thunder 9200™ Architecture and Performance Configuration Guidelines

by
Gil Rangel
Bruce Bush

Thunder 9200 Hardware Architecture

Internals

The Hitachi Freedom Storage™ Thunder 9200™ consists of one or two controllers, with up to four fibre ports (two per controller), and 10 disk drives in the Command Module (also called RK unit). Additional disk bays, called Expansion Modules (also called RKA units) and consisting of 10 disks, can be added up to a maximum of 100 disk drives per two controllers. Keep in mind that there are performance or capacity configurations. Under some conditions, a fully configured Thunder 9200 may not offer the best performance. The hard disk drives (HDDs) in the Thunder 9200 are high performance and can saturate the four back-end fibre loops with less than 50 HDDs at high I/O loads. For random or sequential workloads, the saturation point may be reached at 20 to 40 HDDs, depending on the read/write ratio, cache available, and I/O load. For most workloads, it would be more advisable to add Command Module/Expansion Module combinations than to add more Expansion Modules to systems already configured with 40 HDDs.

A PowerPC 750, 300MHz processor manages each controller. At the front and back ends there are two Tachyon® TS chips per controller. The bus between the data controller and cache performs at 805MB/sec. The buses between the controllers perform at 235MB/sec each. These features improve performance at varying percentages, depending on the workload. Figure 1 illustrates the internal bus and processors structures.

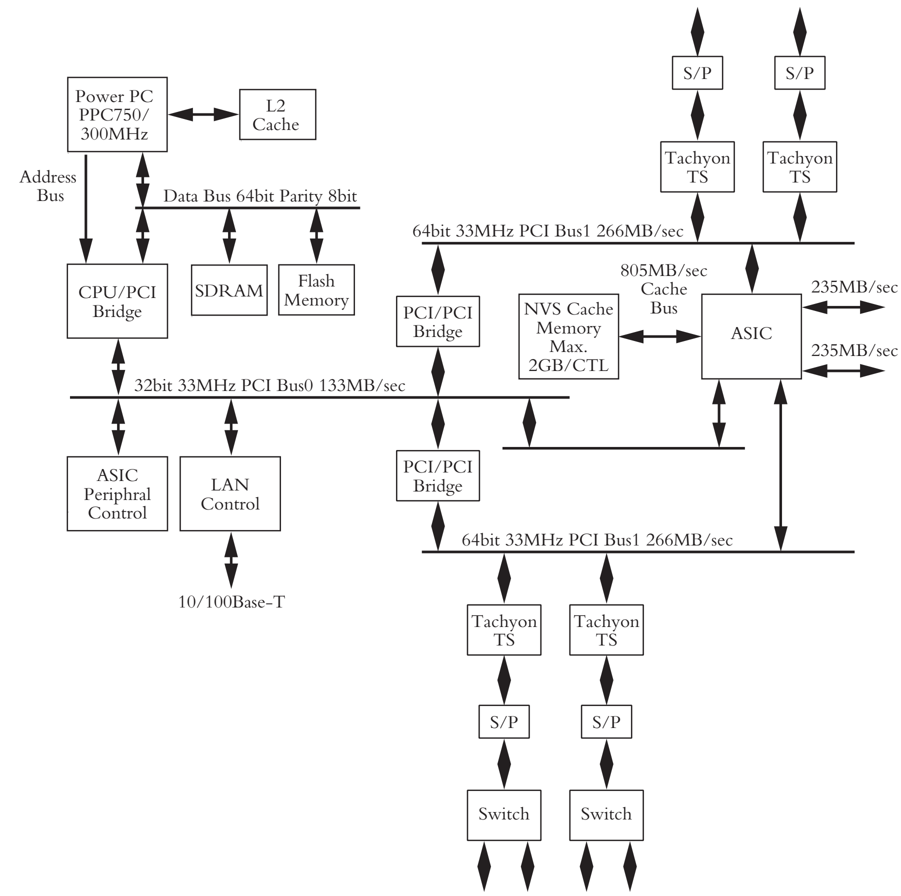

Figure 1: Internal Controller Function Diagram.

Back-end Fibre Channel Architecture

At the back end, there are four active fibre loops to x-number of dual-ported fibre HDDs. Figure 2 illustrates how these loops are used for the primary paths and alternate paths in case of failure. Since each controller has two active paths, HDDs within a RAID group will use alternating fibre loops. This design allows for a fairly balanced workload if the RAID groups are configured correctly. The details of how the physical HDDs are connected to the loops are very important in creating the RAID groups and assigning LUNs to a controller. This will be discussed in further detail later in this document.

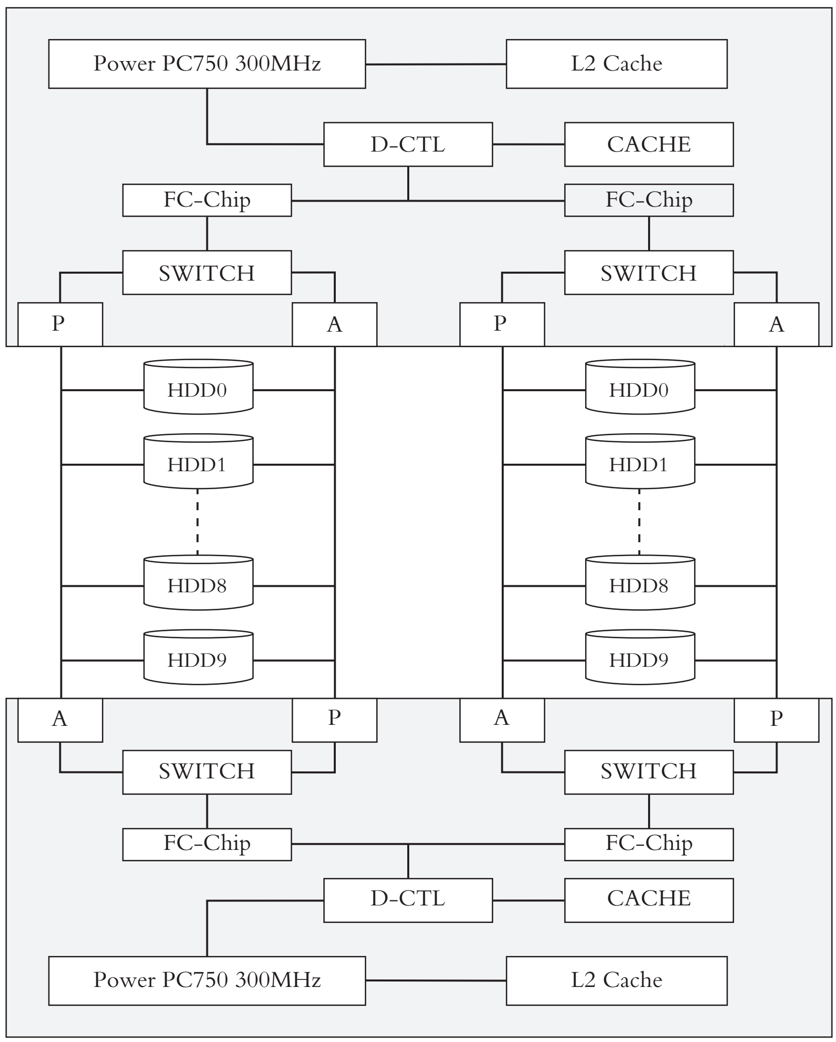

Figure 2: Fibre Loop Diagram.

Figure 3 illustrates the fibre paths for LUNs owned by controller 0. This example is a 9D+1P parity group, with one LUN owned by controller 0. The Fibre Channel loops alternate at the physical HDD level in an even/odd mapping.

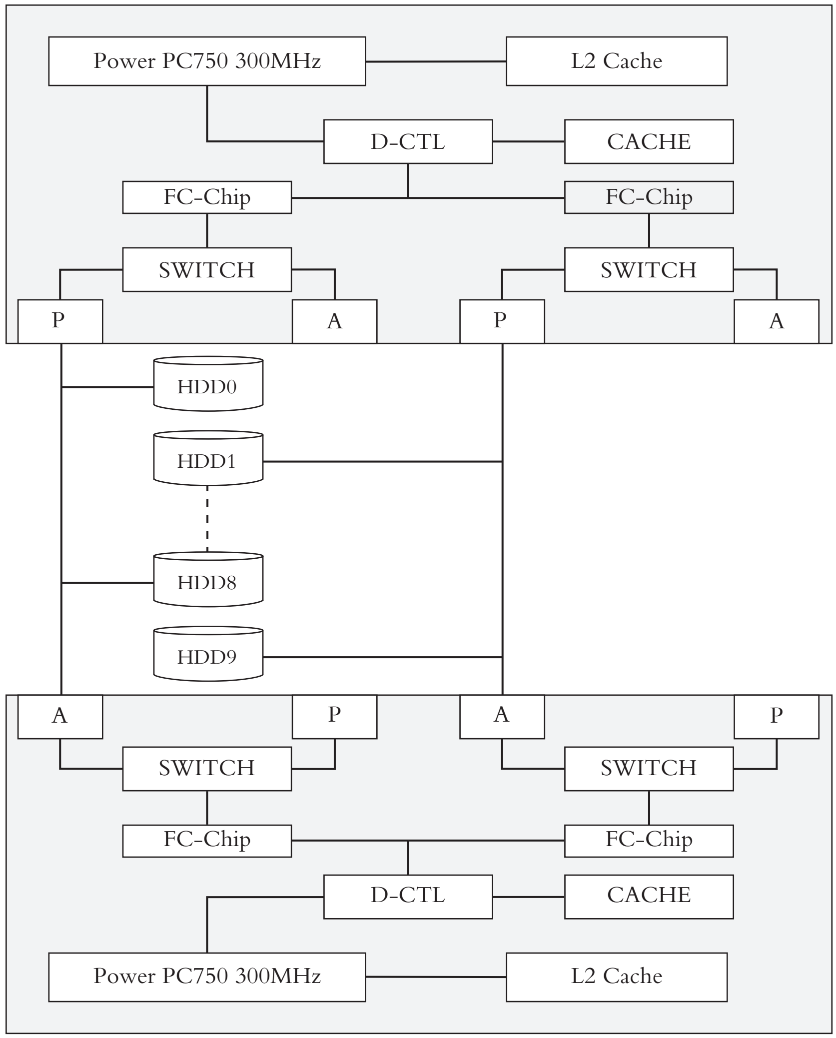

Figure 3: LUNs Assigned to Controller 0.

Figure 4 illustrates the fibre paths for LUNs owned by controller 1. This example depicts a 9D+1P parity group, with one LUN owned by controller 1. The Fibre Channel loops alternate at the physical HDD level in an even/odd mapping.

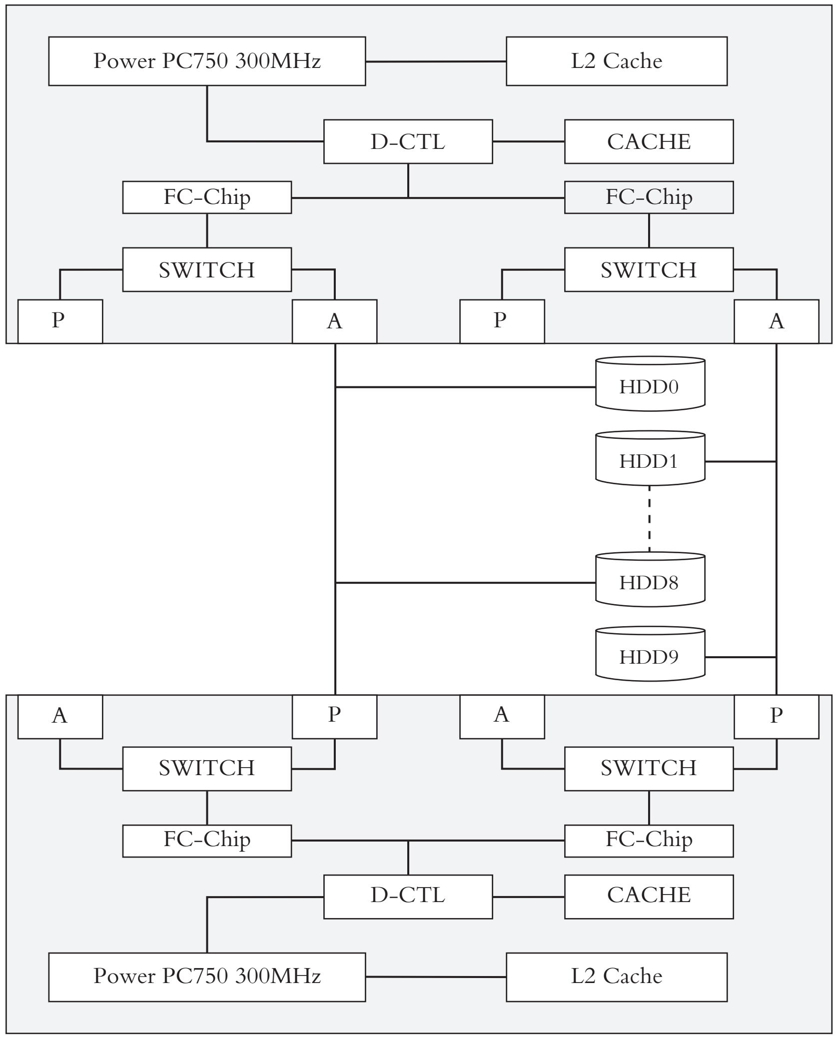

Figure 4: LUNs Assigned to Controller 1.

Configuration Guidelines and General Rules

Understand Your Customer's Environment

Before recommending a Thunder 9200 configuration, it is important to know how the customer will use the product. Although we attempt to simulate a variety of environments, it would be impossible to know in advance each customer's workloads. The items of most importance are as follows.

A. What is the customer trying to accomplish? The more data you have, the easier it is to architect a solution.

B. What are the current workload statistics? If none, then try to characterize the workload based on several models. (Various models will be included later in this document.)

C. What server platform(s)—CPUs, memory, HBAs, OS levels (including patches)—will be attached?

D. What database is being used, if any?

E. What is the network infrastructure and average number of users?

F. What are the high-availability requirements—clusters, DMP, etc.?

G. What are the capacity requirements?

Disk Drive Arrangement

The back-end communications between the controllers and the disks are not designed to dynamically load-balance. Thus, it is important for maximum performance to plan the Parity Group configuration to have the same number of physical disks on each back-end loop pair. An unbalanced drive arrangement may cause performance degradation. Figure 5 illustrates six parity groups, 2D+1P, that are not evenly distributed across the fibre back-end paths. In this configuration, the physical disk 9 from units 0 and 1 can be used as spares.

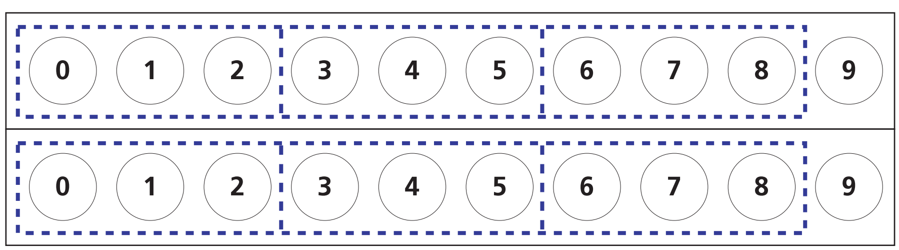

Figure 5: Unbalanced RAID Group Configuration for RAID5, with six parity groups—2D + 1P.

Figure 6 illustrates six parity groups, 2D+1P, that are evenly distributed across the fibre back-end paths. In this configuration, physical disk 9 from unit 0, and physical disk 0 from unit 1 can be used as spares.

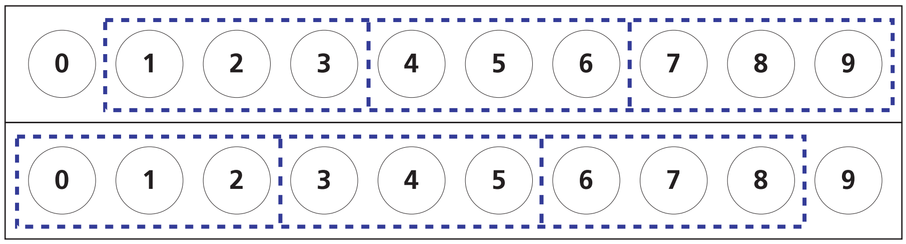

Figure 6: Balanced RAID Group Configuration for RAID 5, with six parity groups—2D + 1P.

RAID Levels

The Thunder 9200 supports RAID Levels 0, 1, 0+1, and 5. The factors in determining which RAID level to use are cost, reliability, and performance. Table 1 shows the benefits and disadvantages of each RAID type.

	RAID0	**RAID1**	**RAID0+1**	**RAID5**
Description	Data Striping NO DATA PROTECTION	Disk Mirroring	Data Striping and Mirroring	Data Striping with distributed parity
Minimum # Disks	2	2	4	3
Maximum # Disks	16	16	16	16
Benefit	High Performance for most workload models	Data Protection through redundancy	High Performance with data redundancy	The best balance of cost, reliability, and performance
Disadvantages	NO DATA PROTECTION	Higher cost per number of physical disk	Higher cost per number of physical disk	Performance penalty for high percentage of Random Writes

Table 1: RAID Level Comparisons.

RAID Groups and Parity Groups

RAID groups can contain one or more parity groups. When configured in this manner, the potential exists for a LUN assignment to span parity groups, which will cause delays in the parity group and drive selection process. In addition, LUNs that span parity groups will cause parity to be generated for each parity group during writes. This will reduce the overall performance; therefore, for maximum performance, each RAID group should be configured as a single parity group.

LUN Capacity per RAID Group

The Thunder 9200's flexibility allows multiple LUNs per RAID group, but the goal should be to keep this to a minimum. The number of LUNs will depend on the RAID type and capacity requirements. For example, you can create two 4D+1P RAID groups, with four LUNs each, as opposed to one 9D+1P RAID group with eight LUNs. The overall capacity for the two 4D+1P groups would be ~32GB less, but the 9D+1P group may reduce the overall performance due to undesirable mechanical seek delays between LUNs in the same parity group. In addition, multiple LUNs in the same parity group should be assigned to the same controller. Additional testing with different RAID types and larger disk configurations still needs to be performed.

Stripe Size within a Parity Group

The Thunder 9200 supports stripe widths 16KB, 32KB, and 64KB. Most of the caching algorithms have been optimized for a 64KB stripe. Altering this will negatively affect performance. Additional testing will be forthcoming to illustrate the performance degradation.

LUN Assignment and Controller Ownership

The Thunder 9200's controllers manage LUNs by ownership. This creates a reference for the controller's cache to manage hit/miss determination and for duplexed writes.

Selecting the Proper Disk Drive Form Factor

In all cases, small-capacity, high-RPM drives will provide better performance in terms of full random access. When cached data locality is low, multiple, small-capacity, high-RPM drives should be used. In some cases, a higher rotation speed drive can provide better sequential-access performance due to more sophisticated drive technology, but this effect is minimized by our intelligent caching algorithms for sequential reads and writes.

Performance Considerations for the Thunder 9200 Architecture

Cache Data Allocation Single Controller

When using the Thunder 9200, all read/write data goes through cache. In a single controller configuration there is no mirroring of write data. This is a default that cannot be modified. Figure 7 illustrates the cache management for a single controller configuration.

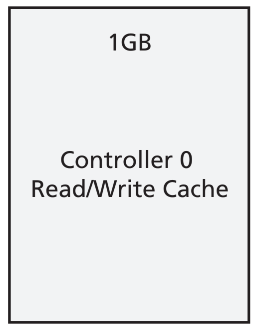

Figure 7: Single Controller with 1GB Total Cache Configuration.

Cache Data Allocation Dual Controller/Dual Active

When the Thunder 9200 is configured with dual active controllers, active/active, all Read/Write data goes through cache. All write data written to the primary controller is mirrored to the other controller for data integrity. Figure 8 illustrates the cache management for a dual active controller configuration.

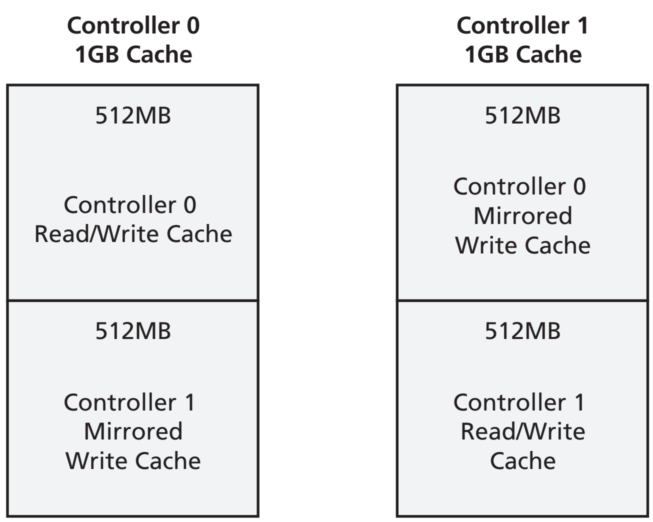

Figure 8: Dual Active Controller Cache Management.

Cache Data Allocation Dual Controller—Active/Hot Standby

When the Thunder 9200 is configured with dual controllers, active/hot standby, all Read/Write data goes through cache. All write data is mirrored to the standby controller for data integrity. Figure 9 illustrates the cache management for an active/hot standby controller configuration.

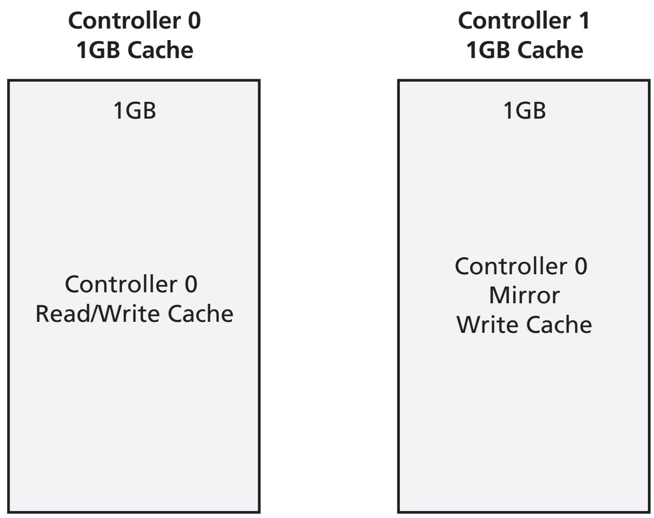

Figure 9: Dual Controllers Active/Hot-Standby Cache Management.

Hitachi FlashAccess 9200™

FlashAccess 9200, formally known as Hitachi TurboLUN™, is basically one or more small LUNs staged in cache the first time they are accessed. All access thereafter will be a cache hit. Systems with a large cache can benefit from this in many different environments, but in a small cache system, such as the Thunder 9200, there may be performance implications. Since the FlashAccess LUN(s) will use cache from both controllers, a configuration with many non-FlashAccess LUNs may see a reduction in overall performance due to a reduced cache hit rate. Additional testing will be done to determine the largest size FlashAccess LUN that can be used without negatively affecting performance of the entire system.

Data Share Mode

When the Thunder 9200 is configured for environments that use a "shared-everything" concept, such as Oracle® Parallel Server or VERITAS® VxVM/DMP, there is a possibility of a LUN being accessed through a path that does not "own" the LUN. When LUN ownership boundaries are crossed, the Thunder 9200 must copy cache information from the owning controller to the accessing controller, update the cache information, make a write to disk, and then copy the cache information back to the owning controller. This effect, known as ping-ponging, degrades performance significantly. Data Share mode should always be enabled to avoid unexpected LUN ownership changes due to access through the non-owning controller path. This becomes more important when configuring the Thunder 9200 in a SAN environment. Figure 10 illustrates normal I/O requests to LUNs 0 through 3 by the owning controller.

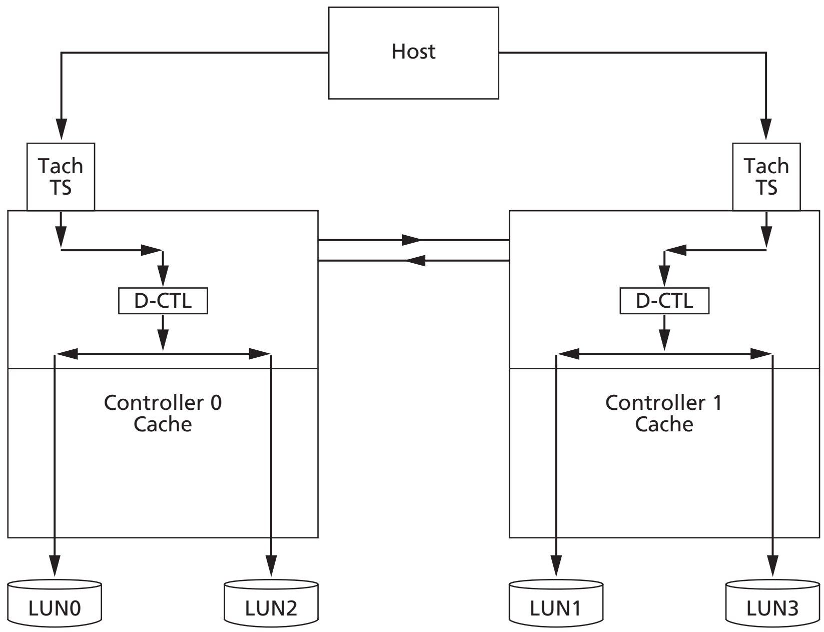

Figure 10: Normal I/O Access.

Data Share mode allows access to the non-owned LUNs through the controller-to-controller communication paths. In the following example, LUNs 0 and 2 are owned by controller 0, LUNs 1 and 3 are owned by controller 1, and the host is requesting data from LUN 0 through controller 1. Figure 11 illustrates the I/O path when Data Share mode is ON.

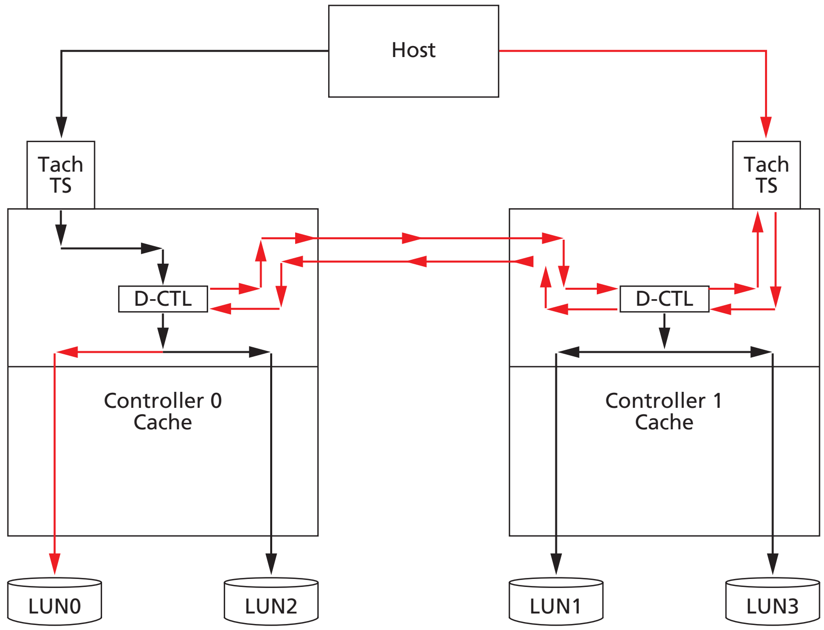

Figure 11: Data Share ON I/O Access.

Multipath Setting

The purpose of the Multipath setting is to determine how Sequential-detect will function.

Multipath (Array Unit)—When determining if access to a drive is sequential, access information for both controllers is examined. This is the default and should be used for configurations with multiple hosts attached to both controllers in an ACTIVE/ACTIVE mode.

Multipath (Controller)—When determining if access to a drive is sequential, examine only the owning controller. This should be used for high availability (HA) environments that are configured with ACTIVE/ACTIVE controllers, ACTIVE/PASSIVE at the host. This is used in combination with TRESPASS mode.

Tagged Command Queuing (TCQ)

Always use TCQ. TCQ primarily benefits Random Read performance. Turning off TCQ will cause an undesirable I/O wait time at the host for all workload types.

The Thunder 9200 will support 256 commands per Fibre Channel port. Excessive queue depth per LUN may cause error recovery routines to fail. For best performance, use the formula: ***queue depth = 256 / <number of LUNs per port>.***

Performance

Workload Characteristics

The following is a list of workload models. These are intended to be guidelines only, as each customer's environment will differ. There is a high probability that some customers will have a combination of these workloads. Therefore, it's important to understand the characteristics of each in determining the best possible Thunder 9200 configuration.

- On-line Transaction Processing (OLTP)
  OLTP is characterized by a large number of small I/O requests to a number of different databases. The typical OLTP workload can have I/O block sizes ranging from 2 through 32KB, with 67 to 85 percent Random Read.
- Decision Support Systems (DSS) and On-line Analytical Processing (OLAP)
  DSS is characterized by a number of different workloads. The first is moving data from the OLTP databases into a larger database or data warehouse (DW). Data marts (DM) are smaller, department-specific databases that are extracted from the DW and add external items. The data for the first part of DSS is typically large block, 512KB to 8MB, Sequential Read of OLTP databases, and Sequential Write to the DW. Once in the DW or DM, the data is typically accessed in large blocks, 512KB to 8MB Sequential Read.
- Microsoft® Windows NT®—File Systems
  Windows NT with a File System uses a 64KB block regardless of what the application requests, or whether the I/O is Random or Sequential.
- Windows® 2000—File Systems
  Unlike Windows NT, Windows 2000 will transfer I/O blocks to the disk at the application's requested size.
- Web Servers
  The content of most Web servers is accessed at 8KB or 16KB, Random Read.
- Video Servers
  Video servers are characterized by streaming large amounts of data at large blocks. The typical VS workload will have I/O block sizes 1MB and larger Sequential Read. Updates to these servers, internally or externally, will be 1MB or larger Sequential Write.
- Satellite Data Collection and Manipulation
  Data blocks of 1 to 32MB are downloaded and written sequentially to disk. The stored data is then read sequentially by multiple hosts. In this environment, speed is king. The faster the data can be downloaded the better.
- UNIX® File Systems (UFS) and Email Servers (Exchange is not included)
  It's very difficult to characterize a UFS or Email Server, primarily because of two factors: The varying numbers of users in the read/write ratio and the operating system's buffering mechanism.

Server Configuration

The server used for this testing was a Sun E3000 running Solaris® 2.6 with four 250MHz UltraSparc™ processors, 1GB RAM, and four LP8000 PCI Fibre Channel adapters. Additional testing with different servers will be performed and reported at a later date.

Thunder 9200 Configuration

The Thunder 9200 was configured with dual controllers, dual active, with 512MB cache per controller. There were numerous RAID group configurations, with one LUN per RAID group for each test. Please note that performance can vary significantly with cache size. Additional testing with varying cache sizes will be performed and reported at a later date.

RAID Comparison Performance Charts

The following charts need to be carefully examined to determine the best possible configuration for the various workload models covered earlier in this section. Keep in mind that these are guidelines, which, when used in combination with the understanding of your customers workload, performance, cost, and capacity requirements, will help to determine which RAID type and number of physical disk per parity group to use. In addition, RAID0 was included for comparison only; it is not recommended because it provides no data protection.

Random Read Performance

In examining Figure 12, it's clear that for two Fibre Channels, 20 physical disk—RAID5 or RAID0+1—perform relatively the same for high Random Read workloads, such as OLTP. The determining factors as to which to use at this point become capacity and cost requirements.

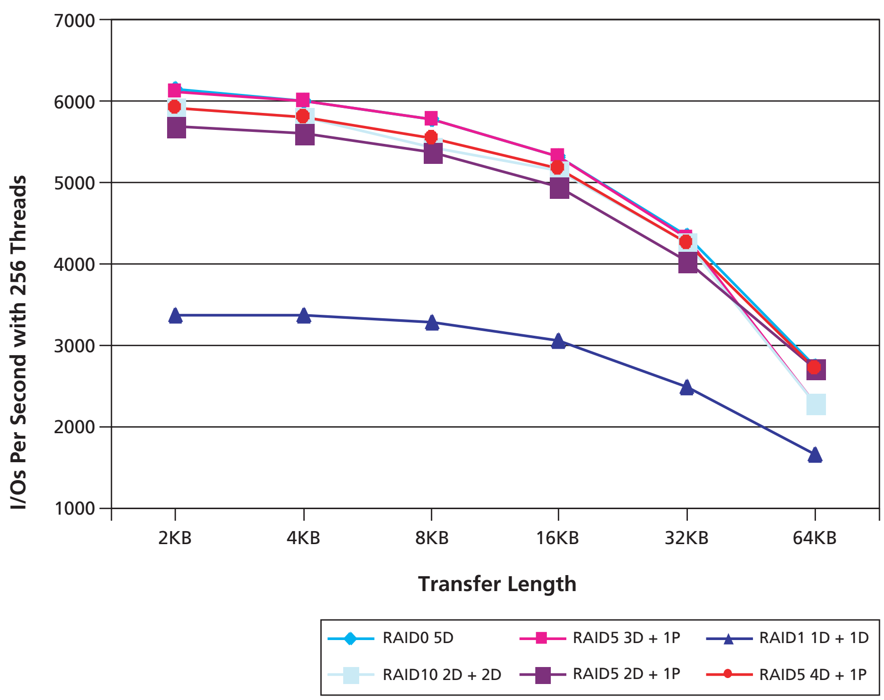

Figure 12: 100 Percent of Random Read Performance for RAID0, 1, 10, and 5, with 20 HDDs and two FCs.

In examining Figure 13, it's clear that for four Fibre Channels and 40 physical disk, RAID5 performs better than RAID1 or RAID0+1 for highly Random Read workloads, such as OLTP. Also in examining Figures 12 and 13, a clear performance advantage is apparent if you were to have two Command Module/Expansion Module combinations with 20 HDDs each, instead of one Command Module/ Expansion Module with 40 HDDs.

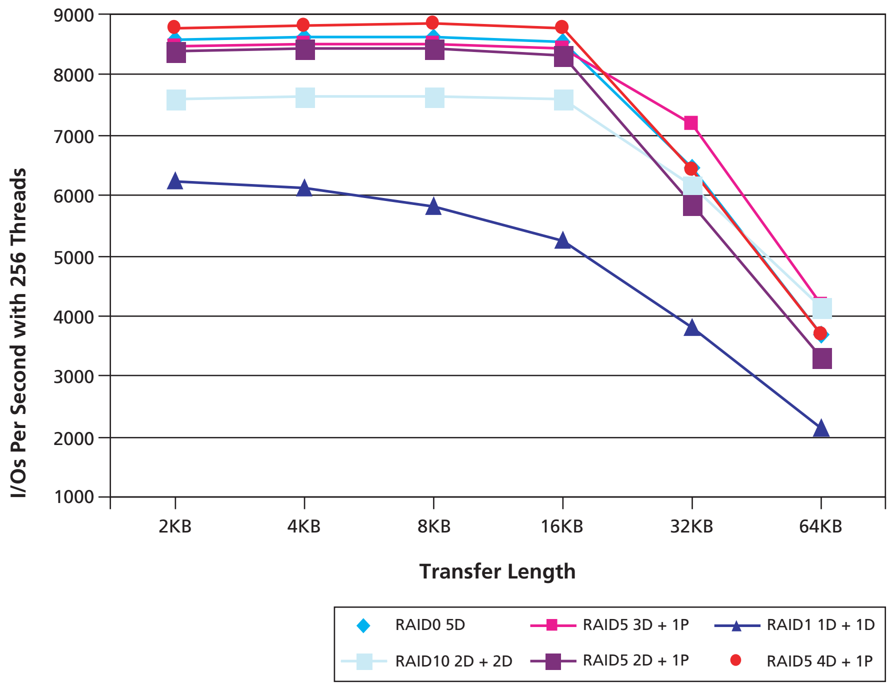

Figure 13: 100 Percent of Random Read Performance for RAID0, 1, 10, and 5, with 40 HDDs and four FCs.

Random Write Performance

In examining Figure 14, it's clear that RAID1 and RAID0+1 offer better performance than RAID5 for highly Random Write workloads. High Random Write is more common in general UNIX and some Windows NT file system workloads.

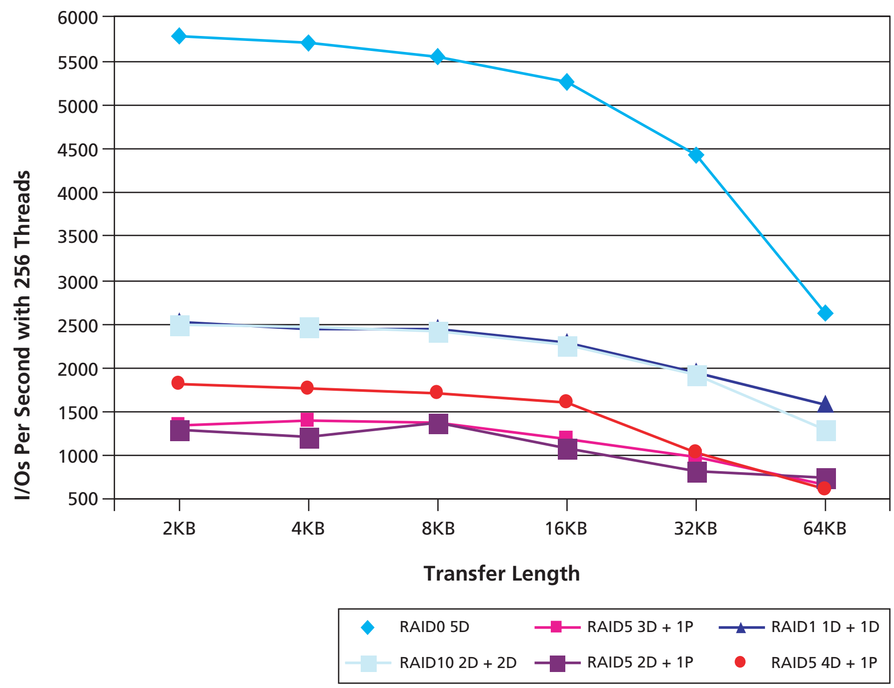

Figure 14: 100 Percent of Random Write Performance for RAID0, 1, 10, and 5, with 20 HDDs and two FCs.

In examining Figure 15, it's still apparent that RAID1 and RAID0+1 offer better performance than RAID5 for highly Random Write workloads. The probability of having this type of workload is fairly common in general UNIX and Windows NT file systems. In addition, in examining Figures 14 and 15, you can determine that it would be wiser to have two Command Module/Expansion Module combinations with 20 HDDs each versus one Command Module/Expansion Module with 40 HDDs for better performance.

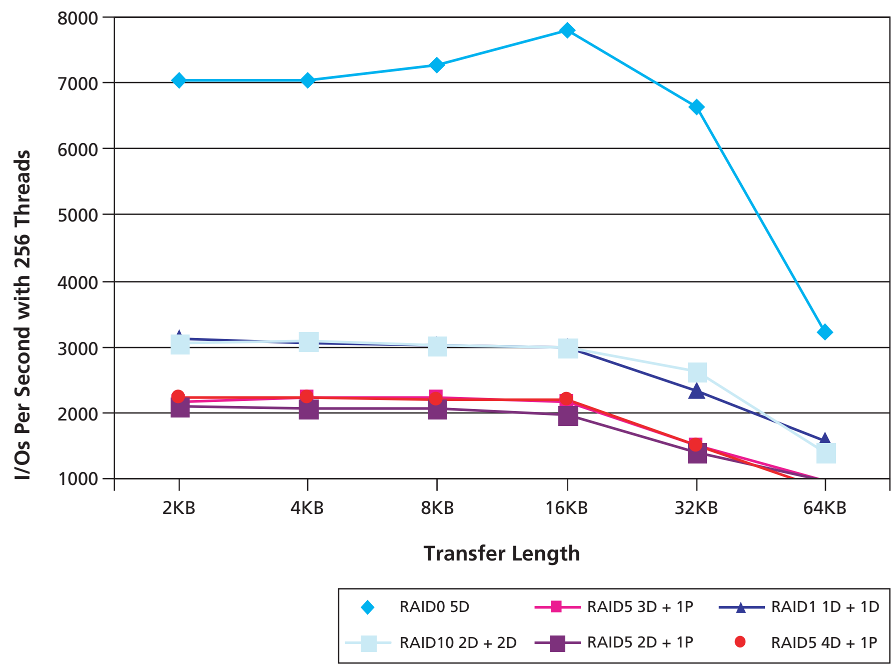

Figure 15: 100 Percent of Random Write Performance for RAID0, 1, 10, and 5, with 40 HDDs and four FCs.

Sequential Read Performance

In examining Figure 16, it can be determined that the various RAID5 configurations offer better performance than RAID1 or RAID0+1 for block sizes less than 1MB. At 1MB and 2MB blocks, performance for RAID0+1 and RAID5 is almost identical. At block sizes larger than 2MB, RAID5 is again superior to RAID1 or RAID0+1. The design for this type of workload is probably one of the more difficult to configure. With the various RAID5 configurations, some will be more expensive. The determining factors as to which to use will most likely be cost, capacity requirements, and the application or RDBMS block size most commonly used.

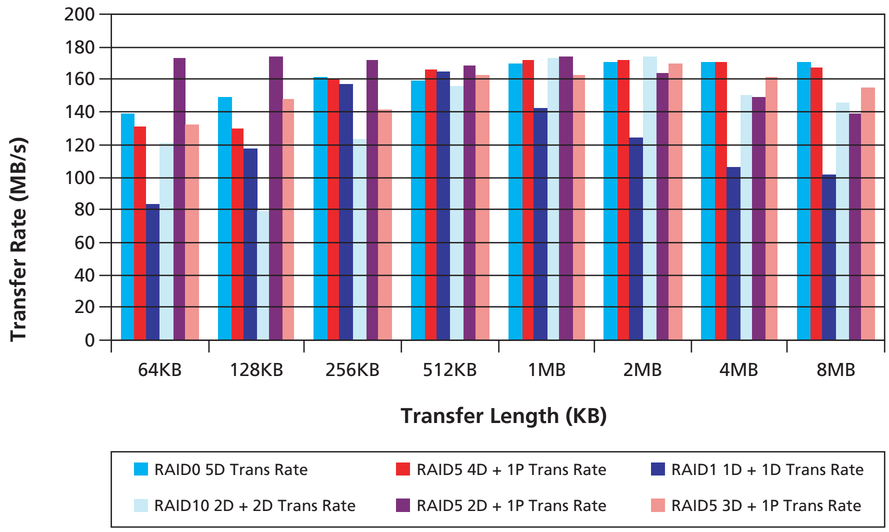

Figure 16: 100 Percent Sequential Read Performance for RAID0, 1, 10, and 5, with 20 HDDs and two FCs.

In examining Figure 17, it can be seen that most RAID5 variations offer better performance than RAID1 or RAID0+1 for the majority of block sizes. With the various RAID5 configurations, some will be more expensive. The determining factors as to which to use will most likely be cost, capacity requirements, and the application or RDBMS block size. In addition, at 40 HDDs and maximum I/O load the Thunder 9200 reaches the saturation point for the back-end fibre loops. Adding Expansion Modules to this configuration might decrease the overall performance.

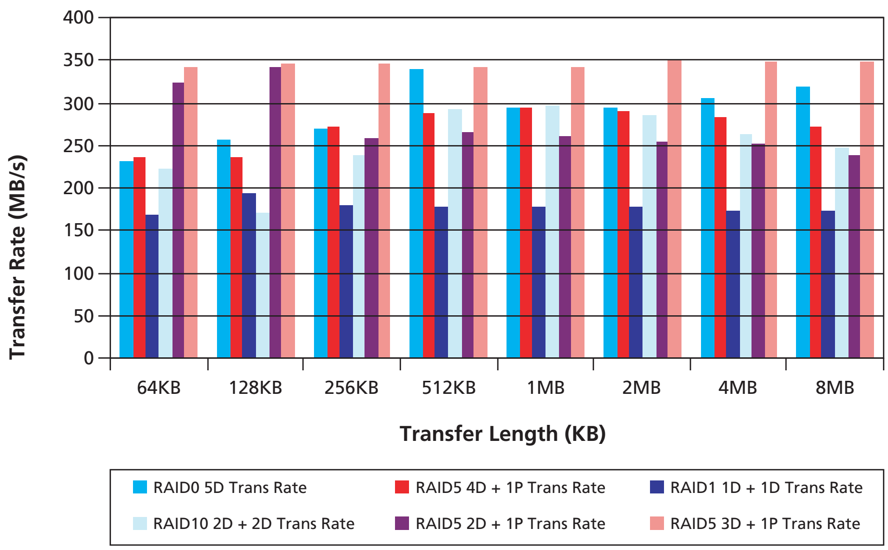

Figure 17: 100 Percent Sequential Read Performance for RAID0, 1, 10, and 5, with 40 HDDs and four FCs.

Sequential Write Performance

In examining Figure 18, it is apparent that most RAID5 variations offer better performance than RAID1 or RAID0+1 for the majority of block sizes. The design for this type of workload is probably one of the more difficult to configure. With the various RAID5 configurations, some will be more expensive. The determining factors as to which to use will most likely be cost, capacity requirements, and application block sizes.

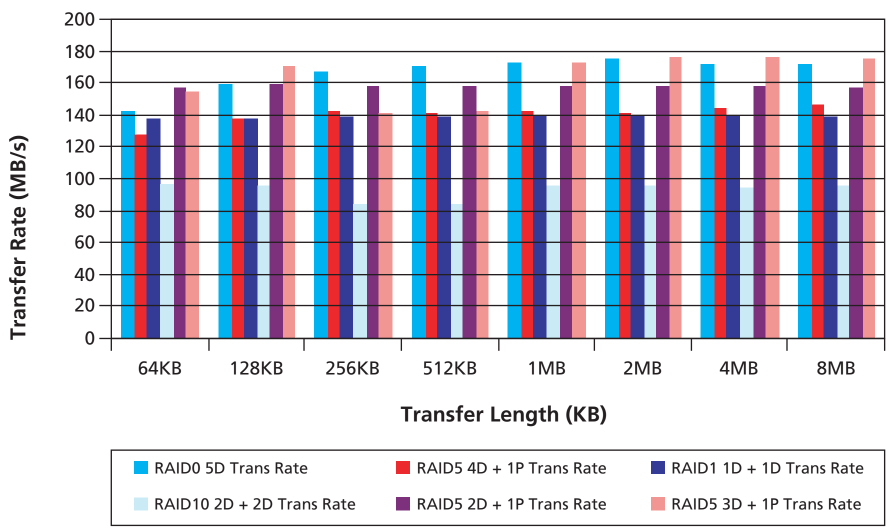

Figure 18: 100 Percent Sequential Write Performance for RAID0, 1, 10, and 5, with 20 HDDs and two FCs.

In examining Figure 19, it can be observed that most RAID5 variations offer better performance than RAID1 or RAID0+1 for the majority of block sizes. With the various RAID5 configurations, some will be more expensive. The determining factors as to which to use will most likely be cost, capacity requirements, and application or satellite download block sizes. In addition, at 40 HDDs and maximum I/O load the Thunder 9200 reaches the saturation point for the back-end fibre loops. Adding Expansion Modules to this configuration might decrease the overall performance. A configuration with two Command Module/Expansion Module combinations would offer better overall performance for this type of workload.

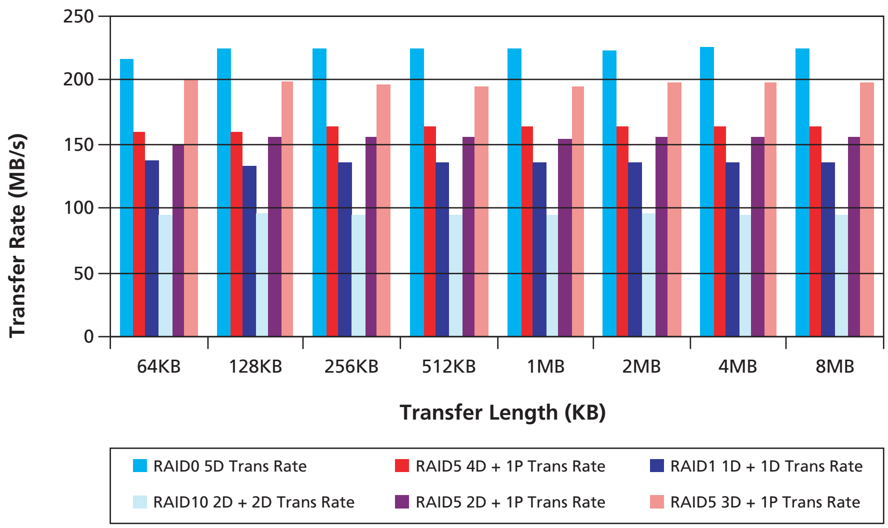

Figure 19: 100 Percent Sequential Write Performance for RAID0, 1, 10, and 5, with 40 HDDs and four FCs.

High Availability Environment Settings

Sun Solaris and VERITAS VxVM Parameters

Section	Parameter	Value for v. 2.6	Value for v. 3.0.1
System	System Startup Attribute	Dual Active	
Host Interface	Data Share Mode	ON	
Host Port	LU Mapping	Recommended	
	Host Connection Mode	None	Multipath (Controller)
	Host Connection Mode 1	Standard Mode	Trespass Mode
	Host Connection Mode 2	VxVM DMP Mode Enabled	
	LIP Reset Mode	Target Authentication	
	Reset/LIP Mode (Process)	Reset Within the Port	
	Reset LIP Mode (Signal)	Reset Within the Port	
Name	Target Reset (Bus Device Reset) Mode	Reset Within the Port	
	Vendor ID	N/A to DF500F	HITACHI
	Product ID	N/A to DF500F	DF400
	Serial Number		
	Controller Identifier	Enabled	
	Controller ID	DF500-00 [1]	

[1] Should be unique between DF units in the same site.

Table 2: Sun Solaris and VERITAS VxVM Parameters

Sun Solaris and Hitachi Dynamic Link Manager™ Parameters

Section	Parameter	Value
System	System Startup Attribute	Dual Active
	Data Share Mode	ON
Host Interface	LU Mapping	Recommended
Host Port	Host Connection Mode	None
	Host Connection Mode 1	Standard Mode
	Host Connection Mode 2	Report Inq Page 83H
	LIP Reset Mode	Target Authentication
	Reset/LIP Mode (Process)	Reset Within the Port
	Reset LIP Mode (Signal)	Reset Within the Port
	Target Reset (Bus Device Reset) Mode	Reset Within the Port
Name	Serial Number	XXXX [1]

[1] Should be the same value between CTL0 and 1, and unique between DF units in the same site.

Table 3: Sun Solaris and Hitachi Dynamic Link Manager Parameters.

HP-UX® Alternate Link Parameters

Section	Parameter	Value
System	System Startup Attribute	Dual Active
	Data Share Mode	ON
Host Interface	LU Mapping	Recommended
Host Port	Host Connection Mode	None
	Host Connection Mode 1	Standard Mode
	Host Connection Mode 2	HP® Mode
	LIP Reset Mode	Target Authentication
	Reset/LIP Mode (Process)	Reset Within the Port
	Reset LIP Mode (Signal)	Reset Within the Port
	Target Reset (Bus Device Reset) Mode	Reset Within the Port

Table 4: HP-UX and Alternate Link Parameters.

IBM® AIX®—Hitachi Dynamic Link Manager Parameters

Section	Parameter	Value
System	System Startup Attribute	Dual Active
	Data Share Mode	ON
Host Interface	LU Mapping	Recommended
Host Port	Host Connection Mode	None
	Host Connection Mode 1	Standard Mode
	Host Connection Mode 2	Report Inq Page 83H
	LIP Reset Mode	Target Authentication
	Reset/LIP Mode (Process)	Reset Within the Port
	Reset LIP Mode (Signal)	Reset Within the Port
	Target Reset (Bus Device Reset) Mode	Reset Within the Port
Name	Serial Number	XXXX [1]

[1] Should be the same value between CTL0 and 1, and unique between DF units in the same site.

Table 5: AIX—Hitachi Dynamic Link Manager Parameters.

Windows NT—Hitachi Dynamic Link Manger Parameters

Section	Parameter	Value
System	System Startup Attribute	Dual Active
	Data Share Mode	ON
Host Interface	LU Mapping	Recommended
Host Port	Host Connection Mode	None
	Host Connection Mode 1	Standard Mode
	Host Connection Mode 2	Report Inq Page 83H
	LIP Reset Mode	Target Authentication
	Reset/LIP Mode (Process)	Reset Within the Port
	Reset LIP Mode (Signal)	Reset Within the Port
	Target Reset (Bus Device Reset) Mode	Reset Within the Port
Name	Serial Number	XXXX [1]

[1] Should be the same value between CTL0 and 1, and unique between DF units in the same site.

Table 6: Windows NT—Hitachi Dynamic Link Manager Parameters.

Windows NT—Microsoft Cluster Server Parameters

Section	Parameter	Value
System	System Startup Attribute	Dual Active
	Data Share Mode	ON
Host Interface	LU Mapping	Mandatory
Host Port	Host Connection Mode	None
	Host Connection Mode 1	Wolfpack Mode
	Host Connection Mode 2	None
	LIP Reset Mode	LIP Port All Rest Mode
	Reset/LIP Mode (Process)	Spread Reset to Other Port(s)
	Reset LIP Mode (Signal)	Spread Reset to Other Port(s)
	Target Reset (Bus Device Reset) Mode	Reset Within the Port

Table 7: Windows NT—Microsoft Cluster Server Parameters.

Sequent® NUMA-Q®—MP Driver Parameters

Section	Parameter	Value for v. 2.6	Value for v. 3.0.1
System	System Startup Attribute	Dual Active	
	Data Share Mode	ON	
Host Interface	LU Mapping	Recommended	
Host Port	Host Connection Mode	None	Multipath (Controller)
	Host Connection Mode 1	Standard Mode	Trespass Mode
	Host Connection Mode 2	VxVM DMP Mode Enabled	
	LIP Reset Mode	Target Authentication	
	Reset/LIP Mode (Process)	Reset Within the Port	
	Reset LIP Mode (Signal)	Reset Within the Port	
	Target Reset (Bus Device Reset) Mode	Reset Within the Port	
Name	Vendor ID	N/A to DF500F	HITACHI
	Product ID	N/A to DF500F	DF400
	Serial Number		
	Controller Identifier	Enabled	
	Controller ID	DF500-00 [1]	

[1] Should be unique between DF units in the same site.

Table 8: Sequent NUMA-Q—MP Driver Parameters.

Summary

The Hitachi Freedom Storage Thunder 9200 offers great flexibility, and it should perform superbly in any environment, provided a little time and effort go into designing a solution specific to each customer's unique workload, capacity, reliability, and cost requirements.

Hitachi Data Systems
www.hds.com
storage@hds.com

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com
channelone@hds.com

Asia Headquarters
39-09 Tower One
Lippo Centre
89 Queensway
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters
11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325-3300
info@hds.com.au

Canada Headquarters
380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdscanada.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-618000
info@hds.co.uk

Latin America Headquarters
750 Central Expressway, MS 3468
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
infolatin@hds.com

U.S. Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
ussalesinfo@hds.com

PERI-101
May 2001