Configuration Rules for Mission Critical Storage

A step-by-step process for selecting and configuring application storage

Storage Technical Marketing
Network Storage
Sun Microsystems

Revision 2.1, 25 May 2000

ALPHA VERSION

Please
Recycle

Revision History

Version	Author	Date	Comments
1.0	Miroslav Klivansky	07 June 1999	First public version
1.1	Miroslav Klivansky	13 July 1999	Fixed diagrams, naming conventions
1.2, 1.3	Miroslav Klivansky	1 October 1999	Legal edits
1.4	Jean-François Marie	12 October 1999	Added T300
1.5	Steve Rogers	2 December 1999	Updated T300
1.6	Miroslav Klivansky	7 February 2000	Updated Methodology; Formatting
1.7	Jenny Stedman	3 April 2000	Added in Rvalues caluculations
2.0	Miroslav Klivansky	8 May 2000	Major update and clean-up
2.1	Mark Elkins, Sadi Marroquin, Jim Heuser	24 May 2000	Updated Components

Introduction

This document outlines some general guidelines for configuring storage platforms to match application requirements. These guidelines are designed to help ensure high levels of availability first, good performance second, and low cost last. Although these guidelines are adequate for most needs, determining the optimal configuration under stringent or unusual requirements require more precise planning.

This document is primarily concerned with the highly available and mission-critical environment, where the goal is to have data always available to the application. Configurations for which a rare failure event may compromise availability are clearly marked. The intended audience is system architects, capacity planners, and technical decision makers who need to suggest storage solutions to address business challenges. Such decision makers include system engineers, installation engineers, personnel from Professional Services, and other technical persons responsible for planning or verifying storage configurations.

Applicability

These guidelines are written to be as generally applicable as possible without sacrificing clarify and technical rigor. As such, they are not universally appropriate. To determine whether these guidelines are for you, consider the conditions listed in Table 1.

Table 1. Applicability of Guidelines

Use These Guidelines	When you . . . • Desire a starting point for designing a storage solution • Have insufficient time, information, or resources to customize • Do not have access to a high level of storage-specific expertise
Seek Additional Assistance	When you . . . • Have stringent availability or performance requirements • Need mainframe connectivity • Know the target utilization of storage will be above 60–70% of capacity

Sequence of Steps

The remainder of this document consists of a sequence of steps designed to guide you through the storage configuration decisions necessary to select a satisfactory solution. Table 2 outlines this sequence and describes the necessary inputs for each step.

Table 2. Configuration Process Sequence

Step	Description	Inputs	Outputs
1	*Gather Information*	• Necessary requirements • Description of environment	• Information for next steps
2	*Select an Application Area*	• Information on application or business challenges	• Application area
3	*Select a Storage Platform*	• Application area • List of desired features	• Preferred platform
4	*Select the RAID level*	• Application area • Preferred platform	• RAID level • Configuration and tuning ideas
5	*Select Stripe Unit Size*	• Application area • RAID level • Preferred platform	• Stripe unit size
6	*Identify Required Capacities*	• Estimate of server CPU power • Storage size	• Preferred platform capabilities • Recommended data size
7	*Identify Suggested Configuration*	• Preferred platform • RAID level • Preferred platform capabilities • Recommended data size	• Configuration and logical volume layout
CONFIGURATION FULLY SPECIFIED			

Step 1: Starting Out with the Right Information

To make good configuration decisions, you need to have the right information. The best place to get it is from the people who, or group which, will eventually use the storage. Before beginning a configuration design, ask the users or system engineers the questions listed in Table 3. Another good place to learn about obtaining and understanding the correct information is Sheri Silverstein's Evaluating Data Management Requirements, which complements these guidelines. If you wish to better understand this topic, it may be the best place to look.

Table 3. Gathering Information

#	Question	Tools
1	What application will this storage be primarily used for? What are your performance expectations? Is the priority higher throughput or lower latency (response time)?	• Ask the application/project expert.
2	What is the total size of the data and expected near-term (12 to 18 month) growth being configured as part of this solution?	• Ask the application/project expert. Account for stored data, logs, indices, temporary space, extents, and staging areas.
3	What server platform will run the application? What server platform will this storage be connected to?	• Ask the application/project expert • M-value table for Solaris servers
4	Which file system will the application use, or will it be using raw devices?	• Ask the application/project expert. The tradeoffs are better manageability with file systems for better performance with raw devices.
5	What fraction of the workload can be considered sequential (versus random)?	• Use TNF tracing to capture application dynamics. Analyze traces to identify sequential references. A block versus time plot gives a qualitative sense. • Ask the application/project expert if the application has not been prototyped at this stage.
6	What fraction of the workload can be considered reads (versus writes or updates)?	• Use TNF tracing. • Use iostat or sar traces. • Ask the application/project expert if not yet prototyped.

Take into account that most applications have different components whose usage patterns differ substantially from one another. If multiple components need to be configured, it may be useful to repeat the configuration process for each component separately.

Step 2: Select an Application Area

Using the information gathered in Step 1, categorize your application into one of the areas listed in Table 4. If your application does not fit neatly into one of these four general areas, select the area that most closely matches the application's storage workload characteristics. Determining the appropriate application area is an important step, upon which subsequent configuration decisions hinge.

Table 4. Characterizing Applications

Application Area	Key Workload Characteristics	Examples
• Online Transaction Processing • Internet Service Providers • Application Service Providers • Enterprise Resource Planning	• 2–8 KB Random I/O • Response times below 15 ms • +30% write component • Log devices are synchronous	Point Of Sale, Customer Service, SAP, People Soft, TPC-C benchmark, Oracle, Informix, Sybase, DB2, Netscape Server, WWW, Excite, Infoseek, World Net
• Attribute intensive NFS version 3 • File Service using NFS version 2	• 8 KB I/Os • Synchronous metadata updates • Attribute response times below 10 ms • Data response times below 50 ms • Multithreaded Random I/O • 10–30% write component	Work Group File Server, Software Development, Older Mail servers and News servers, CIFS, Samba, TAS
• Data Intensive NFS version 3	• 32KB I/Os • Multithreaded Random I/O • Operations on +1 MB files • 20–40% write component	MCAD, Scientific Visualization, Computer Graphic Render Farm, Geo-logic, FEA, and ECAD on top of NFS
• Data Warehousing • Decision Support • High-Performance Computing	• 64–2048 KB Sequential I/O • +90% read component • Bandwidth limited	Data Mining, Red Brick, Market Research, Molecular Modeling, Finite Element Modeling, Seismic Analysis, Meteorology, Geo-logic, FEA, and ECAD over raw devices. Oracle, Informix, DB2

Step 3: Select a Storage Platform

Selecting a storage platform is a complex decision, involving factors such as availability, features, and performance. The right solution is most often a careful balance among these competing factors.

Each of the suggested configurations presented in the Appendices has been designed to provide a high level of availability and performance. You will find details on the features of each platform in the appropriate product documentation; Table 5 outlines the relative performance potentials of each. If performance is an overriding consideration, this information may help you determine the best choice for your situation.

Table 5. Application Relative Performance by Platform

Relative Performance Strengths ✔=OK ✔✔✔=GOOD ✔✔✔✔✔✔✔=BEST	**OLPT ISP ASP ERP**	**Attribute Intensive NFSv3 and All NFSv2**	**Data Intensive NFSv3**	**Decison Support Data Warehouse HPC**
D1000	✔	✔	✔	✔
A1000	✔✔✔✔	✔✔✔✔	✔✔	✔✔
A3500	✔✔✔✔✔✔✔	✔✔✔✔✔✔	✔✔✔✔✔	✔✔✔✔
A3500 FiberChannel	✔✔✔✔✔✔	✔✔✔✔✔✔✔	✔✔✔✔✔✔	✔✔✔✔✔
A5100	✔✔	✔✔✔	✔✔✔✔	✔✔✔✔✔✔
A5200	✔✔✔✔	✔✔✔✔	✔✔✔✔✔	✔✔✔✔✔✔✔
T300	✔✔✔✔✔	✔✔✔✔✔	✔✔✔✔✔✔✔	✔✔✔✔✔✔✔

Step 4: Select a RAID level

In general, RAID5 is best managed by RAID controller hardware. Software-based RAID5 is most effective when the logical volume can be tuned to do full-stripe I/O. This is most feasible for workloads with heavy sequential components and larger I/O sizes. In environments with random access patterns and/or small I/O sizes (like the first two application areas), if cost considerations dictate RAID5, we recommend giving preference to hardware RAID platforms.

The guidelines shown in Table 6 should help you configure your storage. Use the most specific recommendation available for your platform and application workload environment.

Table 6. Suggestions for configuring RAID storage

Application Area	Suggestion	Hardware RAID[a,b] T300	Hardware RAID[c] A1000[d], A3500, A3500FC	Software RAID[e,f] A5100, A5200
OLTP, ISP, ERP	DO	• RAID5 • RAID1 for log devices • RAID1 for volumes with >50% writes • VxVM to create multiple slices per volume or build thin-wide stripes	• Use RAID5 • Use RAID1 for log devices • Use RAID1 for volumes with >30% writes	• Use RAID1 • Use RAID5 in read-only environments • Use experts to set up and tune the logs • FWC accelerates RAID 5 Writes
	DON'T		• Use RAID5 for log devices • Use RAID5 for volumes with >40% writes	• Use RAID5 log volumes • Use RAID5 for >10% write environments
NFS™v2, Attribute Intensive NFSv3	DO	• RAID5 • VxVM to create multiple slices per volume or build thin-wide stripes	• Use RAID5 • Use RAID1 for latency-sensitive volumes • Use RAID1 for volumes with >30% writes	• Use RAID1 • If read-mostly (>80% read) use RAID5 • FWC accelerates RAID 5 Writes
	DON'T		• Use RAID5 for volumes with >40% writes	• Use RAID5 for >20% write environments
Data Intensive NFSv3	DO	• Use RAID5 • VxVM to create multiple slices per volume or build thin-wide stripes	• Use RAID5 • Use RAID1 when <20 ms response times needed • Use RAID1 for volumes with >50% writes	• Use RAID1 • Use RAID5 for volumes with >70% reads
	DON'T		• Use RAID5 for >50% write environments	• Use RAID5 for >30% write environments
DSS, DW, HPC	DO	• Use RAID5 • VxVM to create multiple slices per volume or build thin-wide stripes	• Use RAID5	• Use RAID5 • Use RAID1 for >30% write environments • Use multiple I/O channels and adapters • Spread adapters over multiple I/O busses
	DON'T		• Use RAID5 for >50% write environments	• Use RAID5 for >30% write environments

a. These systems are characterized by a hardware controller which presents RAID volumes to the host rather than to individual disks. The controller also helps accelerate I/O by leveraging nonvolatile memory cache.

b. T300 is mission critical only when using Partner Group configurations.

c. These systems are characterized by a hardware controller which presents RAID volumes to the host rather than to individual disks. The controller also helps accelerate I/O by leveraging nonvolatile memory cache.

d. The A1000 array is **not** considered Mission Critical due to several single points of failure, most notabley the single RAID controller in the enclosure. Multiple A1000 arrays may be combined using Software RAID (i.e., RAID1 mirroring or RAID1+0) into a highly available configuration.

e. These systems are characterized by fast access to the underlying data disk, high-bandwidth channels connecting the array to the server, and no hardware RAID to assist (or get in the way!) on the array.

f. The A5X00 family of arrays are **not** considered Mission Critical in RAID5 configurations. Since the A5X00 enclosure has a single power source and is connected to a single power sequencer, failure of a sequencer or power source will disable the entire array. RAID5 configurations are only recommended of non-mission-critical applications.

Configurations highlighted with yellow crosshatching should be implemented with caution. These configurations are considered *Non*-Mission-Critical, and may experience situations where data becomes unavailable.

Step 5: Select a Stripe Unit Size

Use the values shown in Table 7 to select a Stripe Unit size. These values have been empirically derived by testing the storage platforms with similar workloads and varying the stripe unit size. The stripe unit sizes published here are those that resulted in the best performance for each platform. Combining this stripe unit size with the configuration suggestion from Step 6 provides all the starting information necessary to put together a workable storage configuration.

Table 7. Recommended Stripe Unit Size

Application Area	RAID Level	Storage Platform	Stripe Unit Size (KB)
OLTP,ISP,ERP	1	T300, A3X00, A5X00	16
		A1000	128
	5	T300, A3X00	16
		A1000	128
		A5X00	64
NFSv2, Attribute Intensive NFSv3	1	T300	16
		A1000, A3X00	128
		A5100	32
		A5200	64
	5	T300	16
		A1000, A5200	64
		A3X00	8
		A5100	16
Data Intensive NFSv3	1	T300	32
		A3X00, A5100	64
		A1000, A5200	128
	5	T300	32
		A1000, A3X00, A5100	64
		A5200	128
DSS, DW, HPC	1	T300, A3X00, A5100	64
		A1000	128
		A5200	256
	5	T300, A1000, A3X00, A5100	64
		A5200	128

Step 6: Identify Required Capacities

A storage subsystem needs to match the server in three dimensions. First and second, the total data storage capacity of the storage and number of I/O operations per second that the subsystem can deliver (throughput) needs to be sufficient to meet application requirements. If one of these dimensions falls short of requirements, the storage will either run out of space or become a bottleneck for the system. Lastly, the number of I/O slots required by the storage configuration needs to be vacant on the server. Without sufficient I/O slots the storage may not have the desired availability or performance characteristics.

The total data storage capacity was estimated as part of Step 1. The number of vacant slots needed for the configuration relates to the appropriate storage configuration selected in the next step. If the storage is being added to an existing server, it may be necessary to upgrade the server or purchase more I/O boards.

A number of methods exist for estimating throughput requirements. One of the simplest is to use the M-value metric[1] developed for MVS and adapted to Open Systems by Brian Wong. The M-value for a server indicates the processing potential of the system, and is measured in units of “quanta”. The corresponding notion of an R-value refers to the relative I/O content of an application running on the server. For each quanta of CPU resources used on the server, there is an R-value worth of I/O done by the storage subsystem. For example, if a server has an M-value of 5000 and an R-value of 0.1, then one would expect the server to require on average 500 I/O’s per second from the storage subsystem when operating at peak capacity on the CPU.

In the context of the configuration guidelines, one can use the above methodology by matching the M-value of the application server with the I/O capabilities of the storage configuration via the R-value. Multiply the M-value by an R-value of 0.2 to get the estimated[2] maximum number of I/O’s per second required by the server. Compare this value to the I/O capacity values for the storage configurations listed in Table 8 through Table 10, and make sure that the sum of I/O capacities for the storage configurations exceeds the estimated server requirements.

1. Wong, B. *Characterizing Open Systems Workloads and Comparing Them with MVS*, 1998 Proceedings of the Computer Measurement Group. An internal version with detailed tables may be found on SWAN at **http://maji-poor.ebay/blw/scpm/**

2. An R-value of 0.2 is a reasonable estimate based on the study mentioned above. If the application is known to require little storage I/O, use a value of 0.1. Alternatively, use an R-value of 0.3 to be very conservative, or an even higher value (e.g., 0.5) if the application is very I/O intensive.

To illustrate this process, let us assume a server with an M-value around 50000 (e.g., Sun Enterprise[tm] E6500 Server with 14 CPU's) requires 570GB of Sun StorEdge A5200 FiberChannel array storage configured as RAID1 for frequently updated tables, and 1200GB of A3500 storage configured as RAID5 for the rest of the database. Assuming an R-value of 0.2, 10000 I/O's per second of storage throughput would be needed to meet peak system demands. The data requirements can be met by combining a rack of Sun Storage A5200 arrays and a 3x15 Sun StorEdge A3500 array configuration. From Table 9 and Table 8 the approximate R-estimates for these configurations are 4851 and 10164, respectively. The combined throughput capability of the storage configurations is 15015, which exceeds the estimated requirement of 10000. That means the recommended storage is sufficient to meet both the data capacity and throughput requirements of the server.

Lastly, we recommend having 10-15% excess capacity on new storage configurations to relocate hot spots. These are areas of disk where frequently used data objects cause heavy demands on the storage subsystem. By planning additional space for hot spots, we can isolate the hot spots through monitoring and migrate them to their own devices. This improves the overall response time by giving more resources to the hot device, and isolating other devices from the hot spot. In our example above, since the estimated capability of 15015 IOPS is approximately 50% above the estimated requirement of 10000, we have plenty of space to reallocate hotspots and grow.

Step 7: Identify Suggested Configuration

Table 8 through Table 10 make some general recommendations for configuring the array and dividing it into multiple RAID logical volumes. Each configuration is subsequently covered in its own Appendix, which:

- Diagrams the configuration, including host connections
- Presents a detailed logical volume layout designed for availability and performance
- Details the configuration features, considerations, and components

Locate the appropriate Table and Appendix by matching the desired array platform and recommended RAID level with the desired storage capacity. Use Table 8 for the A1000 and A3500 family, Table 9 for the A5X00 family, and Table 10 for the T300.

For the A5X00 platform you have the additional choice of directly connecting the arrays to the host or going through hubs. The tradeoff is higher performance and greater stability for I/O slots on the host. If there are enough I/O slots available on the host, it is advantageous to select the Direct configuration. For the other platforms we currently only detail the direct attached option.

TABLE 8. RAID Geometry Suggestions for the A1000 and the A3500 family

Array	RAID1[a]		RAID5	
Size	Small/Medium	Medium/Large	Small/Medium	Medium/Large
A1000[b]	App. A (p.20) **55GB, 462 IOPS[c]** R1: 2 x (3+3)	App. B (p.22) **491GB, 4158 IOPS** R1: 18 x (3+3)	App. C (p.26) **91GB, 770 IOPS** R5: 2 x (5+1)	App. D (p.28) **819GB, 6930 IOPS** R5: 18 x (5+1)
A3500	App. E (p.31) **264GB, 2233 IOPS** R1: 5 x (5+5) R1: 1 x (4+4) 2 Hot Spares	App. F (p.34) **792GB, 6699 IOPS** R1: 15 x (5+5) R1: 3 x (4+4) 6 Hot Spares	App. G (p.38) **400GB, 3388 IOPS** R5: 10 x (4+1) R1: 4 x (1+1) 2 Hot Spares	App. H (p.41) **1201GB, 10164 IOPS** R5: 30 x (4+1) R1: 12 x (1+1) 6 Hot Spares
A3500 FiberChannel	App. I (p.45) **264GB, 2233 IOPS** R1: 5 x (5+5) R1: 1 x (4+4) 2 Hot Spares	App. J (p.49) **792GB, 6699 IOPS** R1: 15 x (5+5) R1: 3 x (4+4) 6 Hot Spares	App. K (p.53) **400GB, 3388 IOPS** R5: 10 x (4+1) R1: 4 x (1+1) 2 Hot Spares	App. L (p.57) **1201GB, 10164 IOPS** R5: 30 x (4+1) R1: 12 x (1+1) 6 Hot Spares

a. RAID1 implementation depends on platform and volume management implementation. The A1000, A3500, Solstice DiskSuite, and VERITAS Volume Manager version 3.x implement it as RAID1+0. VERITAS Volume Manager version 2.x and earlier implement it as RAID0+1. The performance of the two is comprable, but RAID1+0 has reliability and recovery speed advantages.

b. High-Availability is achieved by mirroring across A1000 arrays using volume management software like Solstice DiskSuite or VERITAS Volume Manager.

c. The R-estimates are still being refined based on empirical studies. Currently, we estimate 70 IOPS per RAID5 data spindle and 110 IOPS per RAID1 data spindle using software RAID. We estimate 110 IOPS per data spindle for hardware RAID regardless of RAID level. These estimates are based on all spindles being used at 70% of max with multiple queued I/O's and a 60:40 read to write ratio. For detailed information about R-values and M-values, see **`http://majipoor.ebay/blw/scpm/`** .

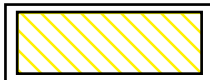

Configurations highlighted with yellow crosshatching should be implemented with caution. These configurations are considered *Non*-Mission-Critical, and may experience situations where data becomes unavailable.

TABLE 9. RAID Geometry Suggestions for the A5X00 family

Array	RAID1[a]				RAID5[b]			
Size	Small/Medium		Medium/Large		Small/Medium		Medium/Large	
Connectivity	Direct	Hubs[c]	Direct	Hubs	Direct	Hubs	Direct[d]	Hubs
A5100	App. M (p.61) **237GB** **1001 IOPS[e]** R1: 4 x (3+3) R1: 1 x (1+1) 2 Hot Spares	Growth Config.	App. N (p.63) **710 GB** **3003 IOPS** R1: 13 x (3+3) 6 Hot Spares	App. O (p.66) **710 GB** **3003 IOPS** R1: 13x (3+3) 6 Hot Spares	App. P (p.69) **1165 GB** **3248 IOPS** R5: 12 x (5+1) R1: 4 x (1+1) 4 Hot Spares	App. Q (p.72) **1165 GB** **3248 IOPS** R5: 12 x (5+1) R1: 4 x (1+1) 4 Hot Spares	Config. May Not Be Practical	App. R (p.75) **3494 GB** **9744 IOPS** R5: 36 x (5+1) R1: 12 x (1+1) 12 Hot Spares
A5200	App. S (p.78) **191 GB** **1617 IOPS** R1: 7x (3+3) 2 Hot Spares		App. T (p.80) **573 GB** **4851 IOPS** R1: 21 x (3+3) 6 Hot Spares	App. U (p.83) **573 GB** **4851 IOPS** R1: 21 x (3+3) 6 Hot Spares	App. V (p.86) **946 GB** **5208 IOPS** R5: 20 x (5+1) R1: 4 x (1+1) 4 Hot Spares	App. W (p.89) **946 GB** **5208 IOPS** R5: 20 x (5+1) R1: 4 x (1+1) 4 Hot Spares		App. X (p.92) **2839 GB** **15624 IOPS** R5: 60 x (5+1) R1: 12 x (1+1) 12 Hot Spares

a. RAID1 implementation depends on platform and volume management implementation. The Solstice DiskSuite, and VERITAS Volume Manager version 3.x implement it as RAID1+0. VERITAS Volume Manager version 2.x and earlier implement it as RAID0+1. The performance of the two is comprable, but RAID1+0 has reliability and recovery speed advantages.

b. RAID5 configurations using the A5X00 family of arrays are considered ***not*** Mission-Critical. The A5X00 enclosures require uninterrupted power to the sequencers to remain available. In addition, smaller configurations may experience unavailability in the rare instance of a sequencer failure.

c. Hub or switch connectivity for Small/Medium RAID1 configurations is not explicitly covered in this document. The configuration is practical to implement with direct connections to the host. If hub or switch connections to the host are desired for future growth, begin with the Medium/Large Hubs configuration and include only Arrays 0 and 3 from the Appendix. The host connections and logical volume layouts remain unchanged for those arrays.

d. Medium/Large RAID5 configurations connected directly to the host would require a large number of I/O slots on the host, and would not deliver significant advantage over hub configurations. As hub technology matures, using hubs will become more advantageous.

e. The R-estimates are still being refined based on empirical studies. Currently, we estimate 70 IOPS per RAID5 data spindle and 110 IOPS per RAID1 data spindle using software RAID. We estimate 110 IOPS per data spindle for hardware RAID regardless of RAID level. These estimates are based on all spindles being used at 70% of max with multiple queued I/O's and a 60:40 read to write ratio. For detailed information about R-values and M-values, see **`http://majipoor.ebay/blw/scpm/`** .

Configurations highlighted with yellow crosshatching should be implemented with caution. These configurations are considered *Non*-Mission-Critical, and may experience situations where data becomes unavailable.

TABLE 10. RAID Geometry Suggestions for the T300 family

Array	**RAID1[a]**		**RAID5**	
Size	**Small/Medium**	**Medium/Large**	**Small/Medium**	**Medium/Large**
T300 (with Hot Spares)	App. Y (p.95) **144GB, 616 IOPS[b]** R1: 2 x (4+4) 2 Hot Spares	App. Z (p.97) **579GB, 2464 IOPS** R1: 8 x (4+4) 8 Hot Spares	App. AA (p.100) **253GB, 1078 IOPS** R5: 2 x (7+1) 2 Hot Spares	App. AB (p.102) **1013GB, 4312 IOPS** R5: 8 x (7+1) 8 Hot Spares
T300 (***no*** Hot Spares)	App. Y (p.95) **162GB, 693 IOPS** R1: 2 x (4.5+4.5)	App. Z (p.97) **651GB, 2772 IOPS** R1: 8 x (4.5+4.5)	App. AA (p.100) **289GB, 1232 IOPS** R5: 2 x (8+1)	App. AB (p.102) **1158GB, 4928 IOPS** R5: 8 x (8+1)

a. The T300 uses a variation of RAID1+0 (sometimes refered to as "diagonal" RAID1) which provides the reliability and recovery advantages similar to traditional RAID1+0, but allows for an odd number of disks in the logical volume. For more information see the T300 Architecture Whitepaper.

b. The R-estimates are still being refined based on empirical studies. Currently, we estimate 70 IOPS per RAID5 data spindle and 110 IOPS per RAID1 data spindle using software RAID. We estimate 110 IOPS per data spindle for hardware RAID regardless of RAID level. These estimates are based on all spindles being used at 70% of max with multiple queued I/O's and a 60:40 read to write ratio. For detailed information about R-values and M-values, see `http://majipoor.ebay/blw/scpm/` .

Layout Naming Conventions

- The naming convention used for the logical volume layout is to label each logical volume member. Generally, this is in terms of *V<volume#>.<member#>*.
- RAID5 volumes use the letter "P" in place of one of the member drives to indicate a Parity drive as part of the logical volume.

  For example, a RAID5 volume may consist of *V7.1, V7.2, V7.3, V7.4, V7.5,* and *V7.P* .

  In reality the which drive in the logical volume holds the parity data changes for each stripe of data. The parity information rotates among the member drives. In the naming conventions the letter "P" for parity is assigned to a single drive to simplify notation.
- RAID1 volumes use the letter "M" to indicate a Mirror drive. (Except T300)

  For example, a RAID1 volume may consist of *V3.1, V3.2, V3.3, V3.1M, V3.2M,* or *V3.3M*.
- Small RAID1 volumes where one disk is mirrored to another disk use the naming convention of a 0 member number and an "M" for Mirror (e.g., *V61.0* and *V61.M*).
- RAID1 volumes use striped and mirrored Dirty Region Logs (DRLs), and are named *Log<volume#>.<member#>*, where the volume number matches the logical volume number of volume being logged, and *Log<volume#>.<member#>M* is the DRL mirror volume (e.g., *Log2.1* and *Log2.1M*).
- All layouts use "HS" to refer to a spindle designated as a Hot Spare.

Combining Hardware and Software RAID

Often, it is useful to combine software and hardware RAID for a single logical volume, and has at times been called using “Thin-Wide Stripes” or “Plaids”. This technique has a number of advantages, and surprisingly few shortcomings. Advantages include higher availability, higher performance, and simplified configuration of large systems. The primary shortcomings are the added conceptual complexity of the system and slightly more complicated management and troubleshooting.

The two most common implementations combining hardware and software RAID are:

1. Using volume management software to mirror two RAID0 logical volumes, where each volume is constructed from multiple physical disks using hardware RAID. This option has the benefit of adding redundancy and availability to a system by making a mirror copy of the data on two very fast devices, where each is by itself inherently risky. If one of the mirrored devices fails, it may take a while to reconstruct it from the other copy, but the data will still be available from the other volume, and no data will be lost. The volume management software can control the recovery rate, and make the tradeoff between recovery time and continuing performance. This option is often chosen for configuring log devices in high-performance transaction processing environments. The underlying RAID0 hardware volumes have very low latency for writes, and archive reads from the log can read from either side of the mirror, reducing archive time. Using this approach it is possible to configure A1000 arrays in a highly-available manner. The performance impact of the additional software RAID layer is minor, since the host can very efficiently issue two write I/O’s in parallel.
2. Using volume management software to create stripes of multiple RAID5 logical volumes, where each volume is constructed from multiple physical disks using hardware RAID. This option has the benefit of spreading I/O to the volume across many physical disks, and is very advantageous in environments where data access is almost entirely random. The theory is that since all the spindles in such an environment are busy doing seeks anyway, it is better to have each transaction spread its I/O requests among as many spindles as possible and parallelize the I/O. This option is often used to configure the bulk of the database tables in large transaction processing environments. The performance impact of the additional software RAID layer depends on how many hardware RAID volumes are included in the stripe. This technique has been successfully employed when spanning hundereds of physical disks and tens of hardware RAID volumes. Generally, stripping across 10-20 hardware RAID volumes will have negligable performance impact on the host CPU.

Covering this topic in depth is beyond the scope of this document. For more information, please see some of the sources listed below:

- Brian Wong’s SUPerG Spring 2000 paper titled A New Methodology for Sizing Storage Configurations
- Bob Larson’s SUPerG Spring 2000 paper and presentation may be found online internally at **`http://dhpg.west/SAE/Presentations/`** titled Wide-Thin Disk Striping for Big Systems.
- Chuck Wenner’s paper titled Plaid Storage Configuration: Combining Hardware and Software RAIDs in One Logical Volume

Appendix A A1000 RAID1 Small/Medium Configuration

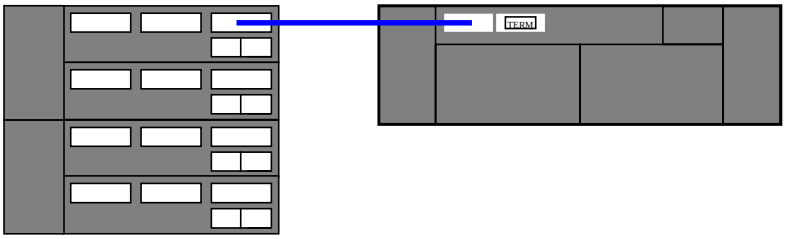

Logical Volume Layout[1] (A1000 RAID1 Small/Medium Configuration)

A1000	V1.1	V1.2	V1.3	V2.1 M	V2.2 M	V2.3 M	V1.1 M	V1.2 M	V1.3 M	V2.1	V2.2	V2.3

1. The layout naming convention is explained on page 18

Details (A1000 RAID1 Small/Medium Configuration)

Configuration

RAID Layout

- 2 3+3 RAID1 logical volumes
- ***NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- Single connection between array and host
- RAID Manager 6.2.x software manages A1000 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager™ (VxVM) used to build additional logical layers on top of hardware RAID volumes

Capacity

- 2 3+3 RAID1 logical volumes @ 27.3 GB (9.1-GB disks)
- **Total data capacity is 54.6 GB (9.1-GB disks)**
- **R-value = 462**

Considerations

Availability

- ***Single RAID controller***
- ***Single connection between array and host***
- ***MULTIPLE SINGLE POINTS OF FAILURE: This configuration is aimed at the Workgroup level and is not designed to have complete redundancy. While the A1000 comes with a number of features to aid availability, certain events may cause loss of access or even data inconsistencies.***
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- 40-MB/s Ultra-SCSI to host
- 10,000-RPM drives for high performance
- 64-MB accelerator cache per controller

Components

Hardware

(1) Ultra™ Fast/Wide Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(1) 68-pin differential terminators (150-1890)
(1) 68-pin differential 2 meter P-cable (530-1885)

Subsystem

(1) A1000 SG-XARY151A-218G (Tabletop/Deskside with 24-MB data cache and 12x18.2-GB disks)
(12) 18.2-GB drives total within array (X5238A)
(1) A1000 SG-XXARY161A-291G (Tabletop/Deskside with 24-MB data cache and 8x36.4-GB disks)
(8) 36.4-GB drives total within array (X5240A)

Software

Sun StorEdge™ RAID Manager 6.2.x or later release
Solaris™ 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager™ 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix B A1000 RAID1 Medium/Large Configuration

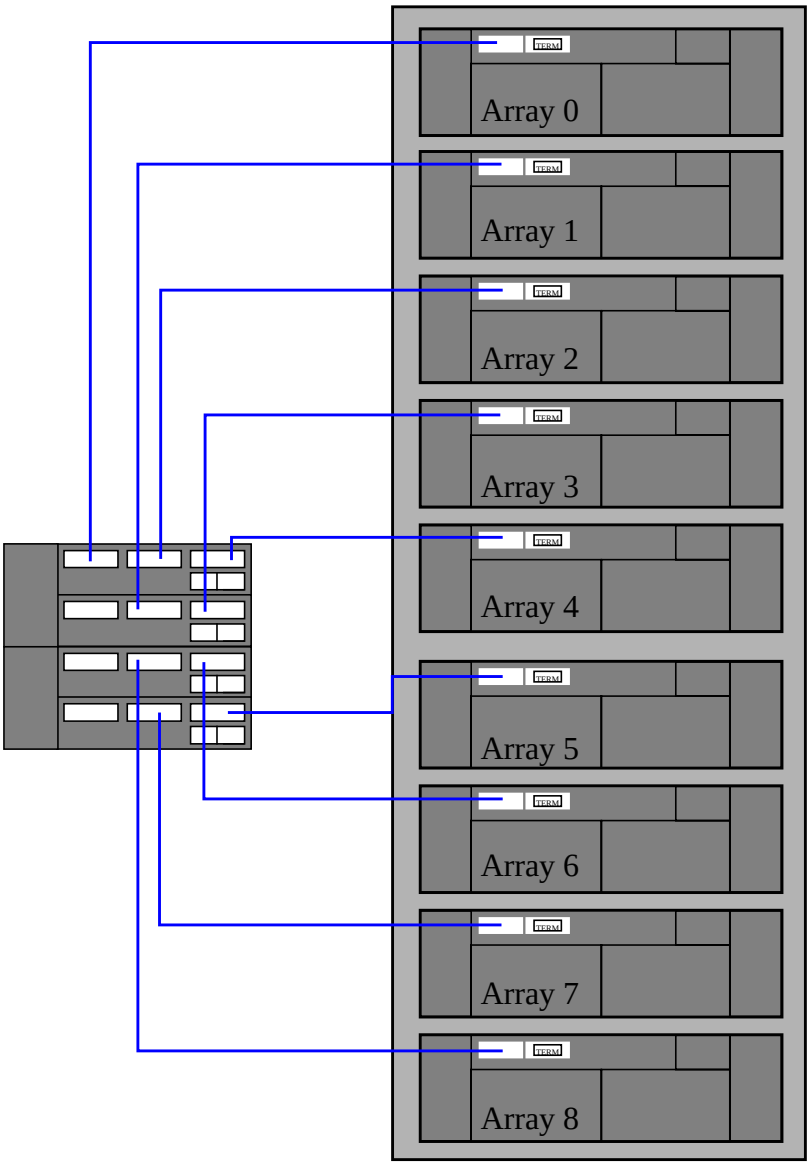

Logical Volume Layout[1] (A1000 RAID1 Medium/Large Configuration)

A0	V1.1	V1.2	V1.3	V2.1M	V2.2M	V2.3M	V1.1M	V1.2M	V1.3M	V2.1	V2.2	V2.3
A1	V3.1	V3.2	V3.3	V4.1M	V4.2M	V4.3M	V3.1M	V3.2M	V3.3M	V4.1	V4.2	V4.3
A2	V5.1	V5.2	V5.3	V6.1M	V6.2M	V6.3M	V5.1M	V5.2M	V5.3M	V6.1	V6.2	V6.3
A3	V7.1	V7.2	V7.3	V8.1M	V8.2M	V8.3M	V7.1M	V7.2M	V7.3M	V8.1	V8.2	V8.3
A4	V9.1	V9.2	V9.3	V10.1M	V10.2M	V10.3M	V9.1M	V9.2M	V9.3M	V10.1	V10.2	V10.3
A5	V11.1	V11.2	V11.3	V12.1M	V12.2M	V12.3M	V11.1M	V11.2M	V11.3M	V12.1	V12.2	V12.3
A6	V13.1	V13.2	V13.3	V14.1M	V14.2M	V14.3M	V13.1M	V13.2M	V13.3M	V14.1	V14.2	V14.3
A7	V15.1	V15.2	V15.3	V16.1M	V16.2M	V16.3M	V15.1M	V15.2M	V15.3M	V16.1	V16.2	V16.3
A8	V17.1	V17.2	V17.3	V18.1M	V18.2M	V18.3M	V17.1M	V17.2M	V17.3M	V18.1	V18.2	V18.3

High-Availability Layout (Software RAID using VxVM or SDS on top of Hardware RAID to mirror volumes for availability. See “Combining Hardware and Software RAID” on page 19 for details.)	**Source A1000 DiskGroup**	**Mirror A1000 DiskGroup**
	V1	V4
	V3	V6
	V5	V8
	V7	V10
	V9	V12
	V11	V14
	V13	V16
	V15	V18
	V17	V2

1. The layout naming convention is explained on page 18

Details (A1000 RAID1 Medium/Large Configuration)

Configuration

RAID Layout

- 18 3+3 RAID1 logical volumes
- ***NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- Single connection between each array and host
- RAID Manager 6.2.x software manages A1000 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) to build additional logical layers on top of hardware RAID volumes

Capacity

- 18 3+3 RAID1 logical volumes @ 27.3 GB (9.1-GB disks)
- **Total data capacity is 491.4 GB (9.1-GB disks) or 245.7 GB with High-Availability Layout**
- **R-value = 4158**

Considerations

Availability

- ***Single RAID controller per array***
- ***Single connection between array and host***
- ***USE SOFTWARE MIRRORS TO AVOID MULTIPLE SINGLE POINTS OF FAILURE: The A1000 array is aimed at the Workgroup level and is not designed to have complete redundancy. While the A1000 comes with a number of features to aid availability, certain events may cause loss of access or even data inconsistencies. Full redundancy may be achieved by layering Software Mirrors on top of the Hardware RAID disk groups (at the cost of redundant data). Follow the High-Availability Layout recommendations if high availability is desired.***
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- The High-Availability Layout configuration offers the highest level of Hardware RAID assistance and will result in the highest level of performance for certain workloads.
- 40-MB/s Ultra-SCSI to host per array (360 MB/s aggregate bandwidth)
- 10,000-RPM drives for high performance
- 64-MB accelerator cache per controller

Components

Hardware

(9) Ultra Fast/Wide Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(9) 68-pin differential terminators (150-1890)
(9) 68-pin differential 12 meter Ultra-SCSI cable (530-1886)
(1) 72-inch Expansion Rack with 2 power sequencers and cables (SG-XARY030A)
(2) Cabinet Power Cables (X3858A)

Subsystem

(9) A1000 SG-XARY152A-72G (Rack mountable with 24-MB data cache and 4 x 18.2-GB)
(9) 64-MB Add-on Cache Memory (X7040A)
(72) 18.2-GB Expansion drives (X5238A) for a total of 108 drives
(9) A1000 SG-XARY164A-145G (Rack mountable with 24-MB data cache and 4 x 36.4-GB)
(9) 64-MB Add-on Cache Memory (X7040A)
(36) 36.4-GB Expansion drives (X5240A) for a total of 72 drives

Configuration

Software	Sun StorEdge RAID Manager 6.2.x or later release Solaris 2.5.1 (8/97) or above with required OS patches (36-GB drives not supported on Solaris 2.5.1)
Other	Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix C A1000 RAID5 Small/Medium Configuration

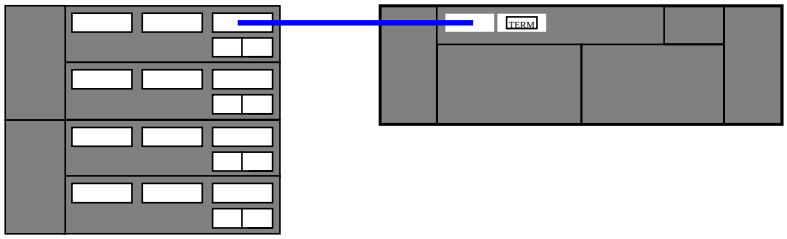

Logical Volume Layout[1] (A1000 RAID5 Small/Medium Configuration)

A1000	V1.1	V1.2	V1.3	V2.1	V2.2	V2.3	V1.4	V1.5	V1.P	V2.4	V2.5	V2.P

1. The layout naming convention is explained on page 18

Details (A1000 RAID5 Small/Medium Configuration)

Configuration

RAID Layout

- 2 5+1 RAID5 logical volumes
- ***NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- Single connection between array and host
- RAID Manager 6.2.x software manages A1000 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes

Capacity

- 2 5+1 RAID5 logical volumes @ 45.5 GB (9.1-GB disks)
- **Total data capacity is 91 GB (9.1-GB disks)**
- **R-value = 770**

Considerations

Availability

- ***Single RAID controller***
- ***Single connection between array and host***
- ***MULTIPLE SINGLE POINTS OF FAILURE: This configuration is aimed at the Workgroup level and is not designed to have complete redundancy. While the A1000 comes with a number of features to aid availability, certain events may cause loss of access or even data inconsistencies.***
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- 40-MB/s Ultra-SCSI to host
- 10,000-RPM drives for high performance
- 64-MB accelerator cache per controller

Components

Hardware

(1) Ultra Fast/Wide Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(1) 68-pin differential terminators (150-1890)
(1) 68-pin differential 2 meter P-cable (530-1885)

Subsystem

(1) A1000 SG-XARY151A-218G (Tabletop/Deskside with 24-MB data cache and 12x18.2-GB disks)
(12) 18.2-GB drives total within array (X5238A)
(1) A1000 SG-XXARY161A-291G (Tabletop/Deskside with 24-MB data cache and 8x36.4-GB disks)
(8) 36.4-GB drives total within array (X5240A)

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager™ 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix D A1000 RAID5 Medium/Large Configuration

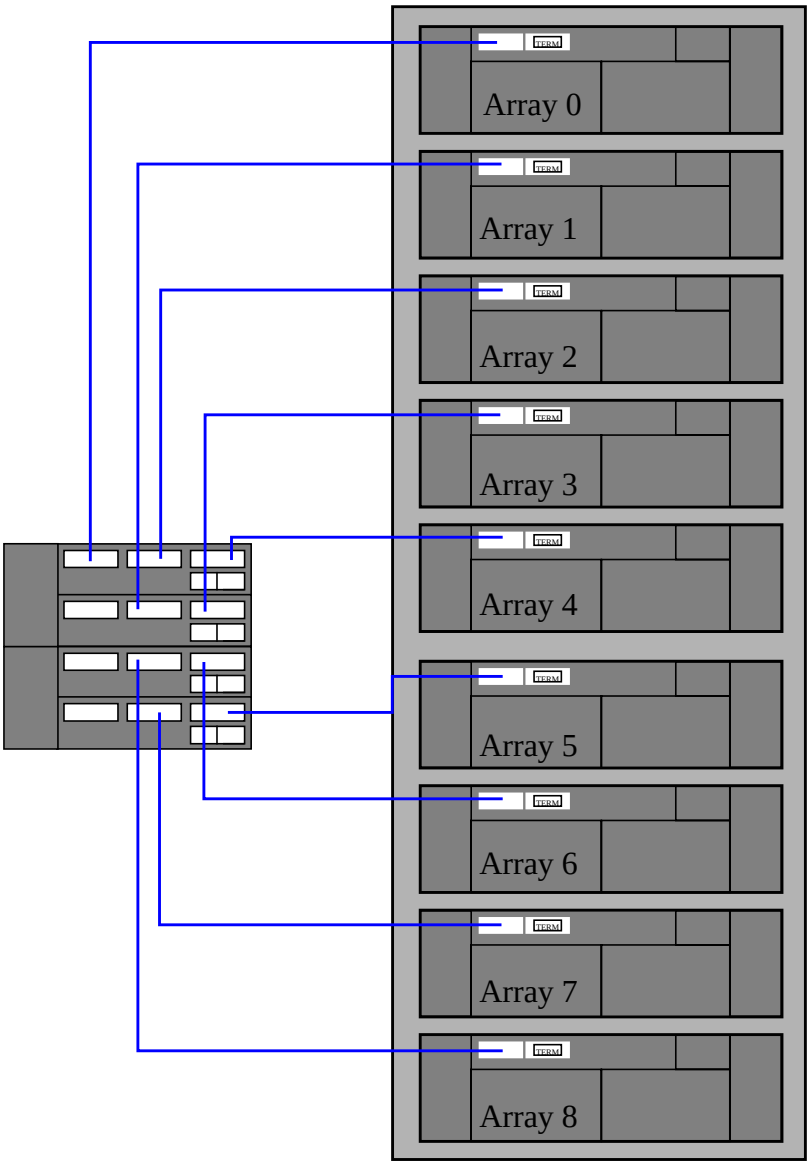

Logical Volume Layout[1] (A1000 RAID5 Medium/Large Configuration)

A0	V1.1	V1.2	V1.3	V2.1	V2.2	V2.3	V1.4	V1.5	V1.P	V2.4	V2.5	V2.P
A1	V3.1	V3.2	V3.3	V4.1	V4.2	V4.3	V3.4	V3.5	V3.P	V4.4	V4.5	V4.P
A2	V5.1	V5.2	V5.3	V6.1	V6.2	V6.3	V5.4	V5.5	V5.P	V6.4	V6.5	V6.P
A3	V7.1	V7.2	V7.3	V8.1	V8.2	V8.3	V7.4	V7.5	V7.P	V8.4	V8.5	V8.P
A4	V9.1	V9.2	V9.3	V10.1	V10.2	V10.3	V9.4	V9.5	V9.P	V10.4	V10.5	V10.P
A5	V11.1	V11.2	V11.3	V12.1	V12.2	V12.3	V11.4	V11.5	V11.P	V12.4	V12.5	V12.P
A6	V13.1	V13.2	V13.3	V14.1	V14.2	V14.3	V13.4	V13.5	V13.P	V14.4	V14.5	V14.P
A7	V15.1	V15.2	V15.3	V16.1	V16.2	V16.3	V15.4	V15.5	V15.P	V16.4	V16.5	V16.P
A8	V17.1	V17.2	V17.3	V18.1	V18.2	V18.3	V17.4	V17.5	V17.P	V18.4	V18.5	V18.P

High-Availability Layout (Software RAID using VxVM or SDS on top of Hardware RAID to mirror volumes for availability. See “Combining Hardware and Software RAID” on page 19 for details)	**Source A1000 DiskGroup**	**Mirror A1000 DiskGroup**
	V1	V4
	V3	V6
	V5	V8
	V7	V10
	V9	V12
	V11	V14
	V13	V16
	V15	V18
	V17	V2

1. The layout naming convention is explained on page 18

Details (A1000 RAID5 Medium/Large Configuration)

Configuration

RAID Layout

- 18 5+1 RAID1 logical volumes
- ***NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- Single connection between each array and host
- RAID Manager 6.2.x software manages A1000 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) to build additional logical layers on top of hardware RAID volumes

Capacity

- 18 5+1 RAID5 logical volumes @ 45.5 GB (9.1-GB disks)
- **Total data capacity is 819.0 GB (9.1-GB disks) or 409.5 GB with High-Availability Layout**
- **R-value = 6930**

Considerations

Availability

- ***Single RAID controller per array***
- ***Single connection between array and host***
- ***USE SOFTWARE MIRRORS TO AVOID MULTIPLE SINGLE POINTS OF FAILURE: The A1000 array is aimed at the Workgroup level and is not designed to have complete redundancy. While the A1000 comes with a number of features to aid availability, certain events may cause loss of access or even data inconsistencies. Full redundancy may be achieved by layering Software Mirrors on top of the Hardware RAID disk groups (at the cost of redundant data). Follow the High-Availability Layout recommendations if high availability is desired.***
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- The High-Availability Layout configuration offers the highest level of Hardware RAID assistance and will result in the highest level of performance for certain workloads.
- 40-MB/s Ultra-SCSI to host per array (360 MB/s aggregate bandwidth)
- 10,000-RPM drives for high performance
- 64-MB accelerator cache per controller

Components

Hardware

(9) Ultra Fast/Wide Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(9) 68-pin differential terminators (150-1890)
(9) 68-pin differential 12 meter Ultra-SCSI cable (530-1886)
(1) 72-inch Expansion Rack with 2 power sequencers and cables (SG-XARY030A)
(2) Cabinet Power Cables (X3858A)

Subsystem

(9) A1000 - SG-XARY146A-36G (Rack-mountable with 24-MB data cache and 4 x 9.1-GB disks)
(9) 64-MB Add-on Cache Memory (X7040A)
(72) 9.1-GB Expansion drives (X5235A) for a total of 108 drives

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches (36-GB drives not supported on Solaris 2.5.1)

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix E A3500 RAID1 Small/Medium Configuration

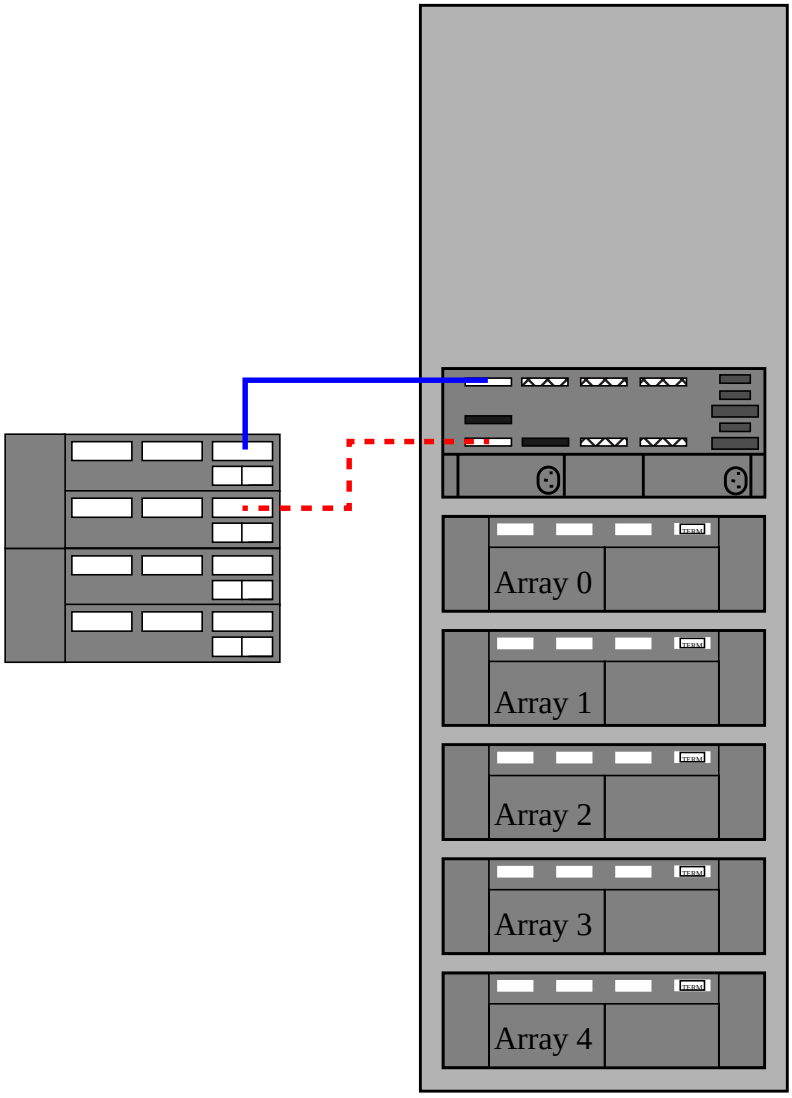

Logical Volume Layout[1] (A3500 RAID1 Small/Medium Configuration)

A0	V1.1	V1.5M	V2.1	V2.5M	V3.1	V3.5M	V4.1	V4.5M	V5.1	V5.5M	V6.1	HS
A1	V1.1M	V1.2	V2.1M	V2.2	V3.1M	V3.2	V4.1M	V4.2	V5.1M	V5.2	V6.1M	HS
A2	V1.3	V1.2M	V2.3	V2.2M	V3.3	V3.2M	V4.3	V4.2M	V5.3	V5.2M	V6.2	V6.4M
A3	V1.3M	V1.4	V2.3M	V2.4	V3.3M	V3.4	V4.3M	V4.4	V5.3M	V5.4	V6.2M	V6.3
A4	V1.5	V1.4M	V2.5	V2.4M	V3.5	V3.4M	V4.5	V4.4M	V5.5	V5.4M	V6.4	V6.3M

Controllers	Controller 0	Controller 1

1. The layout naming convention is explained on page 18

Details (A3500 RAID1 Small/Medium Configuration)

Configuration

RAID Layout

- 5 5+5 RAID1 logical volumes
- 1 4+4 RAID1 logical volume
- 2 Hot Spares
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two UWDIS connected to host
- RAID Manager 6.2.x software manages A3500 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 5 5+5 RAID1 logical volumes @ 45.5 GB (9.1-GB disks)
- 1 4+4 RAID1 logical volumes @ 36.4 GB (9.1-GB disks)
- **Total data capacity is 263.9 GB (9.1-GB disks)**
- **R-value = 2233**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- Dual 40-MB/s Ultra-SCSI connections to host (80 MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(2) Ultra-SCSI Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(2) 68-pin differential terminators (150-1890)

Subsystem

(1) A3500 - SG-XARY360A-545G - (A3500 array in 1x5x12 configuration, mounted in 72-inch expansion rack. Includes 1 controller Module with 2 controllers @ 128-MB, 2 12-meter USCSI cables, redundant fans, drives, and power supplies. Populated with 60 9.1-GB 10,000-RPM drives.)
(4) 64-MB Add-on Cache Memory (X7020A)
Configuration with 18.2-GB drives:
SG-XARY380A-1092G (1 rack, 1 SCSI controller module, 5 disk trays, 60x18.2-GB 10,000 rpm drives)
Configuration with 36.4-GB drives:
SG-XARY381A-1456G (1 rack, 1 SCSI controller module, 5 disk arrays, 40x36.4-GB 10,000 rpm drives)

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix F A3500 RAID1 Medium/Large Configuration

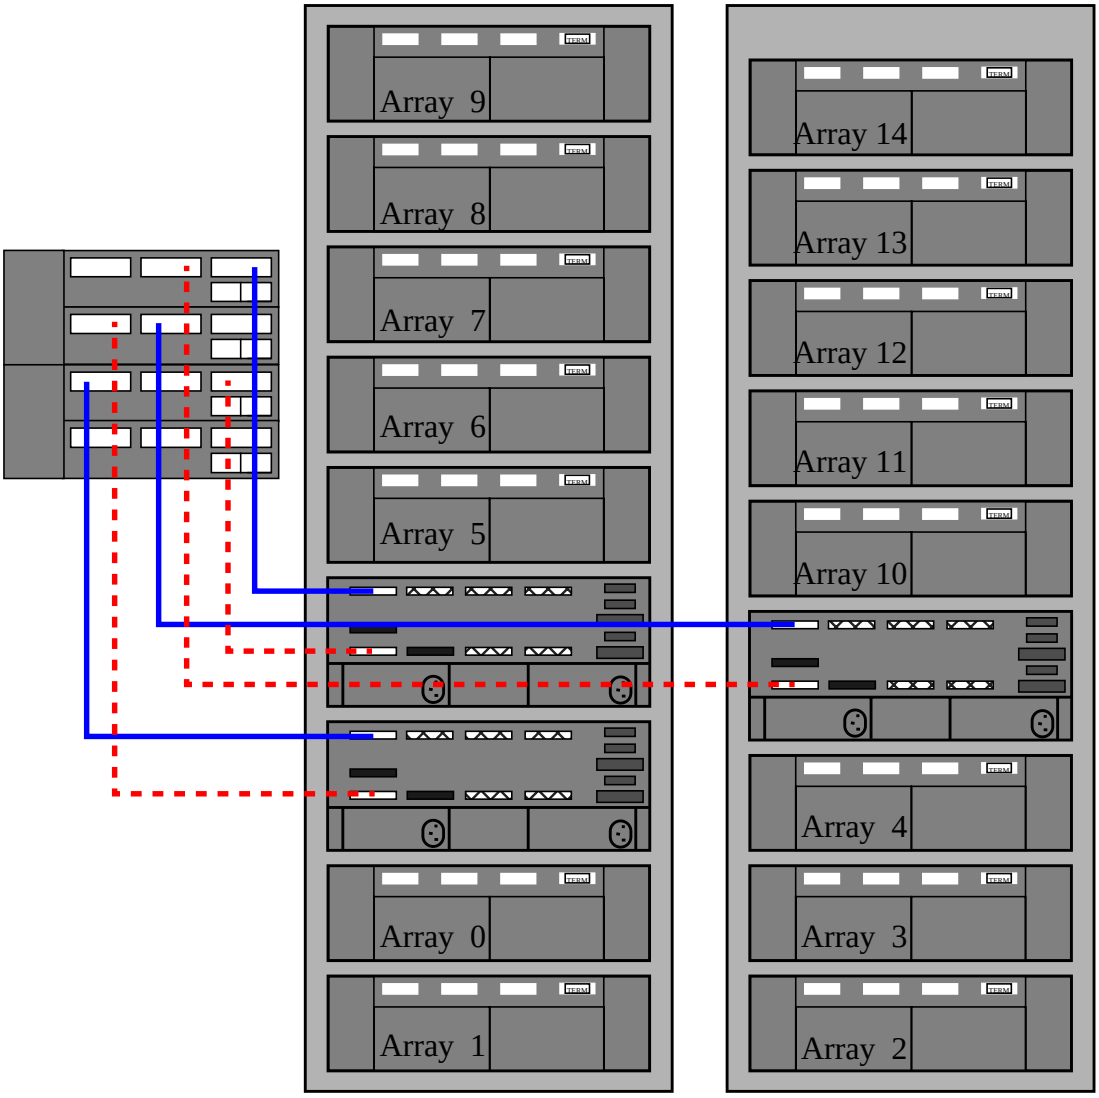

Logical Volume Layout[1] (A3500 RAID1 Medium/Large Configuration)

A0	V1.1	V1.5M	V2.1	V2.5M	V3.1	V3.5M	V4.1	V4.5M	V5.1	V5.5M	V6.1	HS
A1	V1.1M	V1.2	V2.1M	V2.2	V3.1M	V3.2	V4.1M	V4.2	V5.1M	V5.2	V6.1M	HS
A2	V1.3	V1.2M	V2.3	V2.2M	V3.3	V3.2M	V4.3	V4.2M	V5.3	V5.2M	V6.2	V6.4M
A3	V1.3M	V1.4	V2.3M	V2.4	V3.3M	V3.4	V4.3M	V4.4	V5.3M	V5.4	V6.2M	V6.3
A4	V1.5	V1.4M	V2.5	V2.4M	V3.5	V3.4M	V4.5	V4.4M	V5.5	V5.4M	V6.4	V6.3M

A5	V7.1	V7.5M	V8.1	V8.5M	V9.1	V9.5M	V10.1	V10.5M	V11.1	V11.5M	V12.1	HS
A6	V7.1M	V7.2	V8.1M	V8.2	V9.1M	V9.2	V10.1M	V10.2	V11.1M	V11.2	V12.1M	HS
A7	V7.3	V7.2M	V8.3	V8.2M	V9.3	V9.2M	V10.3	V10.2M	V11.3	V11.2M	V12.2	V12.4M
A8	V7.3M	V7.4	V8.3M	V8.4	V9.3M	V9.4	V10.3M	V10.4	V11.3M	V11.4	V12.2M	V12.3
A9	V7.5	V7.4M	V8.5	V8.4M	V9.5	V9.4M	V10.5	V10.4M	V11.5	V11.4M	V12.4	V12.3M

A10	V13.1	V13.5M	V14.1	V14.5M	V15.1	V15.5M	V16.1	V16.5M	V17.1	V17.5M	V18.1	HS
A11	V13.1M	V13.2	V14.1M	V14.2	V15.1M	V15.2	V16.1M	V16.2	V17.1M	V17.2	V18.1M	HS
A12	V13.3	V13.2M	V14.3	V14.2M	V15.3	V15.2M	V16.3	V16.2M	V17.3	V17.2M	V18.2	V18.4M
A13	V13.3M	V13.4	V14.3M	V14.4	V15.3M	V15.4	V16.3M	V16.4	V17.3M	V17.4	V18.2M	V18.3
A14	V13.5	V13.4M	V14.5	V14.4M	V15.5	V15.4M	V16.5	V16.4M	V17.5	V17.4M	V18.4	V18.3M

Controllers	Controller 0	Controller 1
	Controller 2	Controller 3
	Controller 4	Controller 5

1. The layout naming convention is explained on page 18

Details (A3500 RAID1 Medium/Large Configuration)

Configuration

RAID Layout

- 15 5+5 RAID1 logical volumes
- 3 4+4 RAID1 logical volume
- 6 Hot Spares (2 per module, 1 per controller)
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two UWDIS connected to host
- RAID Manager 6.2.x software manages A3500 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 15 5+5 RAID1 logical volumes @ 45.5 GB (9.1-GB disks)
- 3 4+4 RAID1 logical volumes @ 36.4 GB (9.1-GB disks)
- **Total data capacity is 791.7 GB (9.1-GB disks)**
- **R-value = 6699**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- 6 40-MB/s Ultra-SCSI connections to host (240 MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(6) Ultra-SCSI Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(6) 68-pin differential terminators (150-1890)

Configuration

Subsystem

(1) A3500 - SG-XARY374A-273G - (273-GB StorEdge A3500 (30 x 9.1-GB, 10K-rpm drives), with 15 trays mounted in two StorEdge expansion cabinets, redundant fans, drives and power supplies).
(1) A3500 - SG-XARY384A-546G - (546-GB StorEdge A3500 (30 x 18.2-GB, 10K-rpm drives), with 15 disk trays mounted in two StorEdge expansion cabinets, redundant fans, drives and power supplies).
(1) A3500 - SG-XARY385A-1092G - (1092-GB StorEdge A3500 (30 x 36.4-GB, 10K-rpm drives), with 15 trays mounted in two StorEdge expansion cabinets, Redundant fans, drives and power supplies).
(12) 64-MB Add-on Cache Memory (X7020A)
(3) SCSI Controller Modules (X6537A)

Drive Options:
9.1-GB, 1-inch, 10K-rpm drive (X5235A)
18.2-GB, 1-inch, 10K-rpm drive (X5238A)
36.4-GB, 1.6-inch, 10K-rpm drive (X5240A)

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix G A3500 RAID5 Small/Medium Configuration

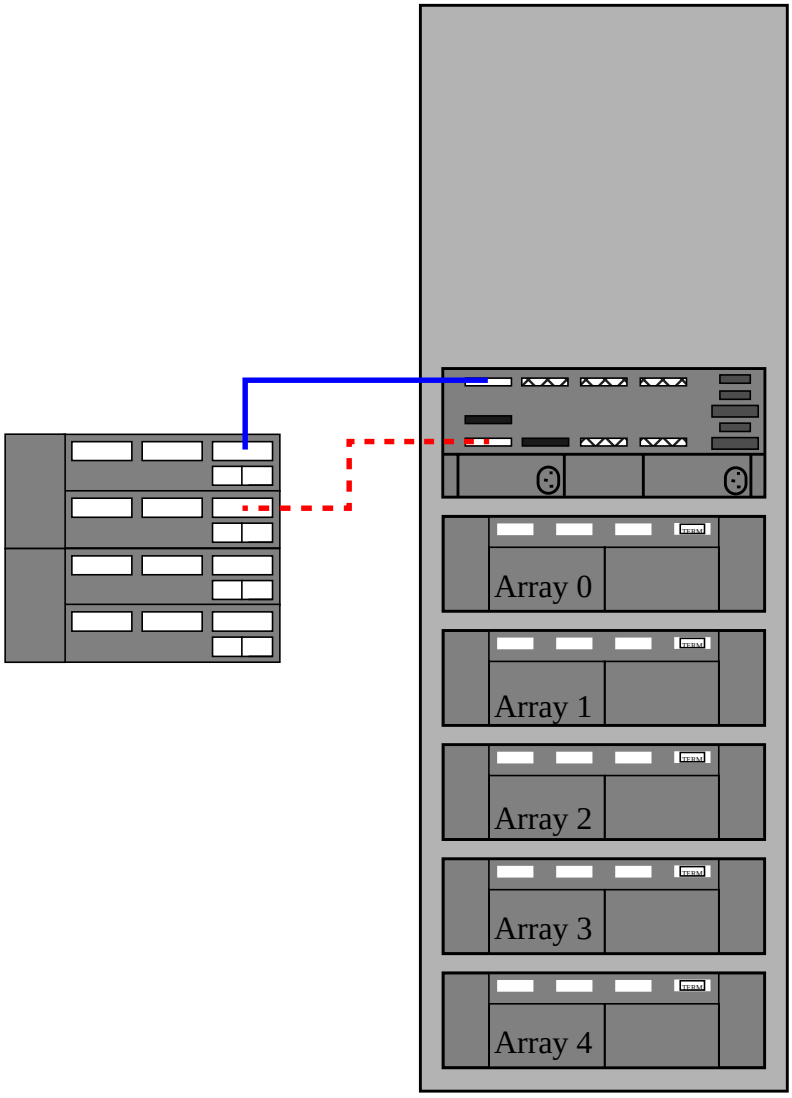

Logical Volume Layout[1] (A3500 RAID5 Small/Medium Configuration)

A0	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.0	HS
A1	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.M	HS
A2	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V12.0	V14.0
A3	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V12.M	V13.M
A4	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V13.0	V14.M

Controllers	Controller 0	Controller 1

1. The layout naming convention is explained on page 18

Details (A3500 RAID5 Small/Medium Configuration)

Configuration

RAID Layout
- 10 4+1RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 2 Hot Spares
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two UWDIS connected to host
- RAID Manager 6.2.x software manages A3500 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity
- 10 4+1 RAID5 logical volumes @ 36.4 GB (9.1-GB disks)
- 4 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks)
- **Total data capacity is 400.4 GB (9.1-GB disks)**
- **R-value = 3388**

Considerations

Availability
- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance
- Controller-based RAID
- 256-MB accelerator cache per controller
- Dual 40-MB/s Ultra-SCSI connections to host (80 MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware
(2) Ultra-SCSI Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(2) 68-pin differential terminators (150-1890)

Subsystem
(1) A3500 - SG-XARY360A-545G - (A3500 array in 1x5x12 configuration, mounted in 72-inch expansion rack. Includes 1 controller Module with 2 controllers @ 128-MB, 2 12-meter USCSI cables, redundant fans, drives, and power supplies. Populated with 60 9.1-GB 10,000-RPM drives.)
(4) 64-MB Add-on Cache Memory (X7020A)

Software
Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other
Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix H A3500 RAID5 Medium/Large Configuration

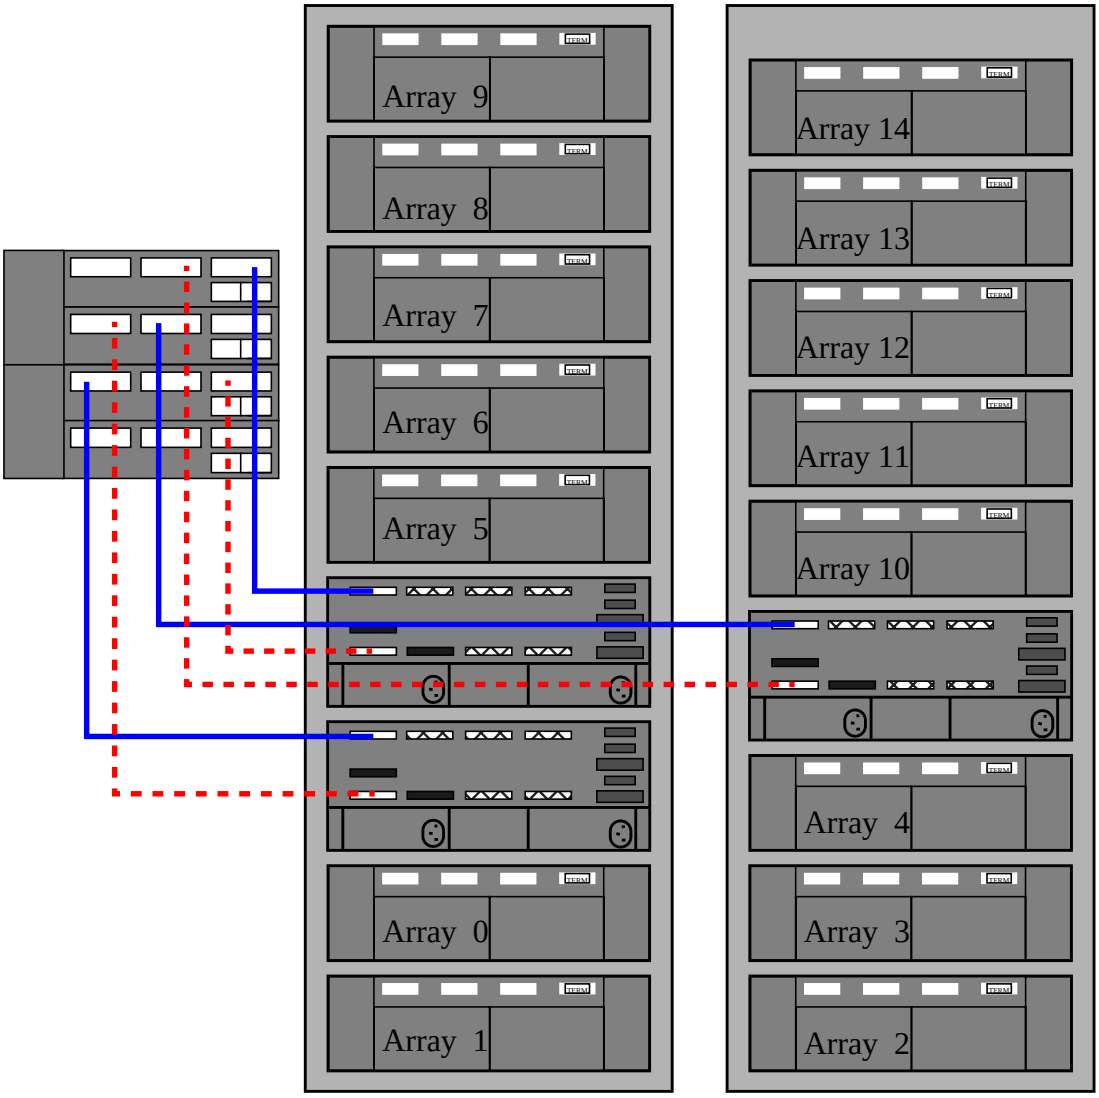

Logical Volume Layout[1] (A3500 RAID5 Medium/Large Configuration)

A0	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.0	HS
A1	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.M	HS
A2	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V12.0	V14.0
A3	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V12.M	V13.M
A4	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V13.0	V14.M

A5	V15.1	V16.1	V17.1	V18.1	V19.1	V20.1	V21.1	V22.1	V23.1	V24.1	V25.0	HS
A6	V15.2	V16.2	V17.2	V18.2	V19.2	V20.2	V21.2	V22.2	V23.2	V24.2	V25.M	HS
A7	V15.3	V16.3	V17.3	V18.3	V19.3	V20.3	V21.3	V22.3	V23.3	V24.3	V26.0	V28.0
A8	V15.4	V16.4	V17.4	V18.4	V19.4	V20.4	V21.4	V22.4	V23.4	V24.4	V26.M	V27.M
A9	V15.P	V16.P	V17.P	V18.P	V19.P	V20.P	V21.P	V22.P	V23.P	V24.P	V27.0	V28.M

A10	V29.1	V30.1	V31.1	V32.1	V33.1	V34.1	V35.1	V36.1	V37.1	V38.1	V39.0	HS
A11	V29.2	V30.2	V31.2	V32.2	V33.2	V34.2	V35.2	V36.2	V37.2	V38.2	V39.M	HS
A12	V29.3	V30.3	V31.3	V32.3	V33.3	V34.3	V35.3	V36.3	V37.3	V38.3	V40.0	V42.0
A13	V29.4	V30.4	V31.4	V32.4	V33.4	V34.4	V35.4	V36.4	V37.4	V38.4	V40.M	V41.M
A14	V29.P	V30.P	V31.P	V32.P	V33.P	V34.P	V35.P	V36.P	V37.P	V38.P	V41.0	V42.M

Controllers	Controller 0	Controller 1
	Controller 2	Controller 3
	Controller 4	Controller 5

1. The layout naming convention is explained on page 18

Details (A3500 RAID5 Medium/Large Configuration)

Configuration

RAID Layout

- 30 4+1 RAID5 logical volumes
- 12 1+1 RAID1 logical volumes
- 6 Hot Spares (2 per module, 1 per controller)
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two UWDIS connected to host
- RAID Manager 6.2.x software manages A3500 RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 30 4+1 RAID5 logical volumes @ 36.4 GB (9.1-GB disks)
- 12 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks)
- **Total data capacity is 1201.2 GB (9.1-GB disks)**
- **R-value = 10164**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- 6 40-MB/s Ultra-SCSI connections to host (240MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(6) Ultra-SCSI Differential host bus adapters (Sbus: X1065A PCI: X6541A)
(6) 68-pin differential terminators (150-1890)

Configuration

Subsystem

(1) A3500 - SG-XARY374A-273G - (273-GB StorEdge A3500 (30 x 9.1-GB, 10K-rpm drives), with 15 trays mounted in two StorEdge expansion cabinets, redundant fans, drives and power supplies).
(1) A3500 - SG-XARY384A-546G - (546-GB StorEdge A3500 (30 x 18.2-GB, 10K-rpm drives), with 15 disk trays mounted in two StorEdge expansion cabinets, redundant fans, drives and power supplies).
(1) A3500 - SG-XARY385A-1092G - (1092-GB StorEdge A3500 (30 x 36.4-GB, 10K-rpm drives), with 15 trays mounted in two StorEdge expansion cabinets, Redundant fans, drives and power supplies).
(12) 64-MB Add-on Cache Memory (X7020A)
(3) SCSI Controller Modules (X6537A)

Drive Options:
9.1-GB, 1-inch, 10K-rpm drive (X5235A)
18.2-GB, 1-inch, 10K-rpm drive (X5238A)
36.4-GB, 1.6-inch, 10K-rpm drive (X5240A)

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix I A3500FC RAID1 Small/Medium Configuration

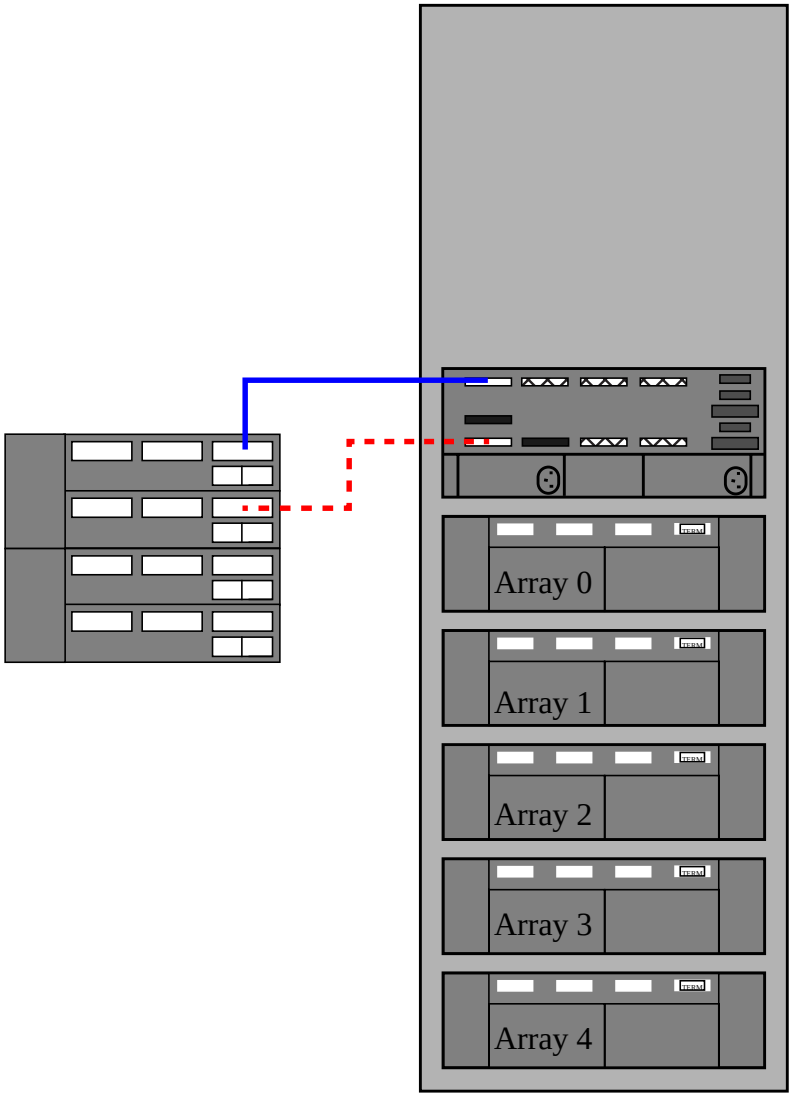

Logical Volume Layout[1] (A3500FC RAID1 Small/Medium Configuration)

A0	V1.1	V1.5M	V2.1	V2.5M	V3.1	V3.5M	V4.1	V4.5M	V5.1	V5.5M	V6.1	HS
A1	V1.1M	V1.2	V2.1M	V2.2	V3.1M	V3.2	V4.1M	V4.2	V5.1M	V5.2	V6.1M	HS
A2	V1.3	V1.2M	V2.3	V2.2M	V3.3	V3.2M	V4.3	V4.2M	V5.3	V5.2M	V6.2	V6.4M
A3	V1.3M	V1.4	V2.3M	V2.4	V3.3M	V3.4	V4.3M	V4.4	V5.3M	V5.4	V6.2M	V6.3
A4	V1.5	V1.4M	V2.5	V2.4M	V3.5	V3.4M	V4.5	V4.4M	V5.5	V5.4M	V6.4	V6.3M

Controllers	Controller 0	Controller 1

1. The layout naming convention is explained on page 18

Details (A3500FC RAID1 Small/Medium Configuration)

Configuration

RAID Layout

- 5 5+5 RAID1 logical volumes
- 1 4+4 RAID1 logical volume
- 2 Hot Spares
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two FiberChannel connections to the host
- RAID Manager 6.2.x software manages A3500FC RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 5 5+5 RAID1 logical volumes @ 45.5 GB (9.1-GB disks)
- 1 4+4 RAID1 logical volumes @ 36.4 GB (9.1-GB disks)
- **Total data capacity is 263.9 GB (9.1-GB disks)**
- **R-value = 2233**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- Dual 100-MB/s FiberChannel connections to host (200-MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(2) FCAL 100-MB host bus adapters (Sbus: X6730A PCI: X6729A)
(2) 68-pin differential terminators (150-1890)

Configuration

Subsystem

(1) A3500FC - SG-XARY360A-545G - (A3500FC array in 1x5x12 configuration, mounted in 72-inch expansion rack. Includes controller Module with 2 controllers @ 128-MB, 2 12-meter USCSI cables, redundant fans, drives, and power supplies. Populated with 60 9.1-GB 10,000-RPM drives.)

(1) A3500FC - SG-XARY360B-545G - (1 rack, 1 FCAL controller module, 5 disk trays, 60 x 9.1-GB 10K-rpm drives

Configuration with 18.2-GB drives:

SG-XARY380B-1092G (1 rack, 1 FCAL controller module, 5 disk trays, 60 x 18.2-GB 10,000-rpm drives)

Configuration with 36.4-GB drives:

SG-XARY381B-1546G (1 rack, 1 FCAL controller module, 5 disk trays, 40 x 36.4-GB 10,000-rpm drives)

Drive options:

X5235A: 9.1-GB 10,000-rpm
X5238A: 18.2-GB 10,000-rpm
X5240A: 36.4-GB 10,000-rpm

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix J A3500FC RAID1 Medium/Large Configuration

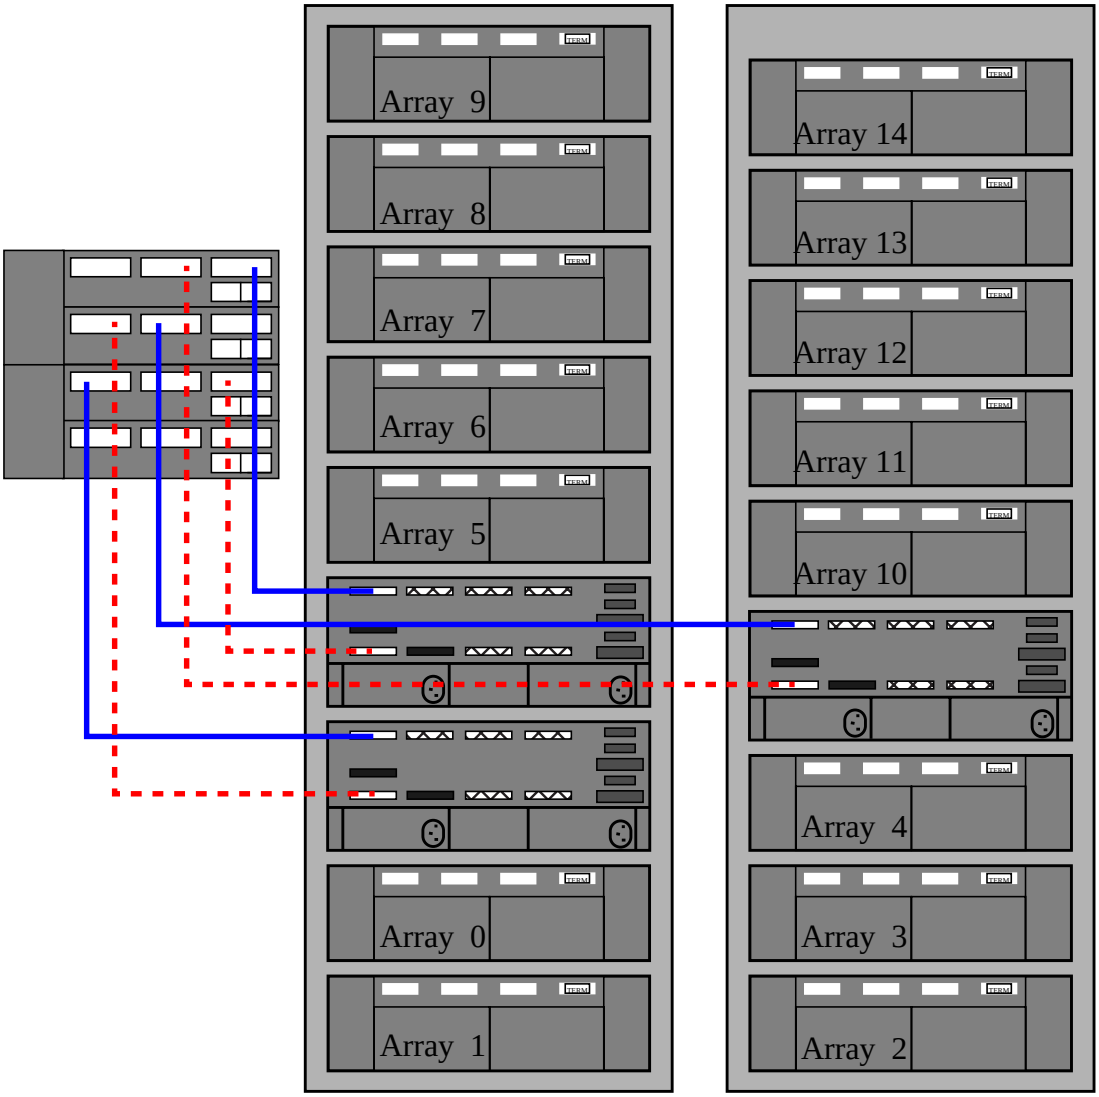

Logical Volume Layout[1] (A3500FC RAID1 Medium/Large Configuration)

A0	V1.1	V1.5M	V2.1	V2.5M	V3.1	V3.5M	V4.1	V4.5M	V5.1	V5.5M	V6.1	HS
A1	V1.1M	V1.2	V2.1M	V2.2	V3.1M	V3.2	V4.1M	V4.2	V5.1M	V5.2	V6.1M	HS
A2	V1.3	V1.2M	V2.3	V2.2M	V3.3	V3.2M	V4.3	V4.2M	V5.3	V5.2M	V6.2	V6.4M
A3	V1.3M	V1.4	V2.3M	V2.4	V3.3M	V3.4	V4.3M	V4.4	V5.3M	V5.4	V6.2M	V6.3
A4	V1.5	V1.4M	V2.5	V2.4M	V3.5	V3.4M	V4.5	V4.4M	V5.5	V5.4M	V6.4	V6.3M

A5	V7.1	V7.5M	V8.1	V8.5M	V9.1	V9.5M	V10.1	V10.5M	V11.1	V11.5M	V12.1	HS
A6	V7.1M	V7.2	V8.1M	V8.2	V9.1M	V9.2	V10.1M	V10.2	V11.1M	V11.2	V12.1M	HS
A7	V7.3	V7.2M	V8.3	V8.2M	V9.3	V9.2M	V10.3	V10.2M	V11.3	V11.2M	V12.2	V12.4M
A8	V7.3M	V7.4	V8.3M	V8.4	V9.3M	V9.4	V10.3M	V10.4	V11.3M	V11.4	V12.2M	V12.3
A9	V7.5	V7.4M	V8.5	V8.4M	V9.5	V9.4M	V10.5	V10.4M	V11.5	V11.4M	V12.4	V12.3M

A10	V13.1	V13.5M	V14.1	V14.5M	V15.1	V15.5M	V16.1	V16.5M	V17.1	V17.5M	V18.1	HS
A11	V13.1M	V13.2	V14.1M	V14.2	V15.1M	V15.2	V16.1M	V16.2	V17.1M	V17.2	V18.1M	HS
A12	V13.3	V13.2M	V14.3	V14.2M	V15.3	V15.2M	V16.3	V16.2M	V17.3	V17.2M	V18.2	V18.4M
A13	V13.3M	V13.4	V14.3M	V14.4	V15.3M	V15.4	V16.3M	V16.4	V17.3M	V17.4	V18.2M	V18.3
A14	V13.5	V13.4M	V14.5	V14.4M	V15.5	V15.4M	V16.5	V16.4M	V17.5	V17.4M	V18.4	V18.3M

Controllers	Controller 0	Controller 1
	Controller 2	Controller 3
	Controller 4	Controller 5

1. The layout naming convention is explained on page 18

Details (A3500FC RAID1 Medium/Large Configuration)

Configuration

RAID Layout

- 15 5+5 RAID1 logical volumes
- 3 4+4 RAID1 logical volume
- 6 Hot Spares (2 per module, 1 per controller)
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Six FiberChannel connections to the host
- RAID Manager 6.2.x software manages A3500FC RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 15 5+5 RAID1 logical volumes @ 45.5 GB (9.1-GB disks)
- 3 4+4 RAID1 logical volumes @ 36.4 GB (9.1-GB disks)
- **Total data capacity is 791.7 GB (9.1-GB disks)**
- **R-value = 6699**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- 6 100-MB/s FiberChannel connections to host (600-MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(6) FCAL 100-MB host bus adapter (Sbus: X6730A PCI: X6729A)
(6) 68-pin differential terminators (150-1890)

Subsystem

(1) A3500FC - SG-ARY374A-273G (2 racks, 15 disk trays, 30x9.1-GB 10K-rpm drives)

(3) FCAL Controller Module (6538A)

Configuration with 18.2-GB drives:

(1) A3500FC - SG-ARY384A-546G (2 racks, 15 disk trays, 30x18.2-GB 10K-rpm drives)
(3) FCAL Controller Module (6538A)

Configuration with 36.4-GB drives:

(1) A3500FC - SG-ARY385A-1092G (2 racks, 15 disk trays, 30x36.4-GB 10K-rpm drives)
(3) FCAL Controller Module (6538A)

Configuration

Software	Sun StorEdge RAID Manager 6.2.x or later release Solaris 2.5.1 (8/97) or above with required OS patches
Other	Sun StorEdge Volume Manager 2.4, 2.5 and 2.5.x, and VxVm 3.x Solstice DiskSuite Version 4.1 for Solaris 2.5.1 systems Solstice DiskSuite Version 4.2 for Solaris 2.6 systems and later releases

Appendix K A3500FC RAID5 Small/Medium Configuration

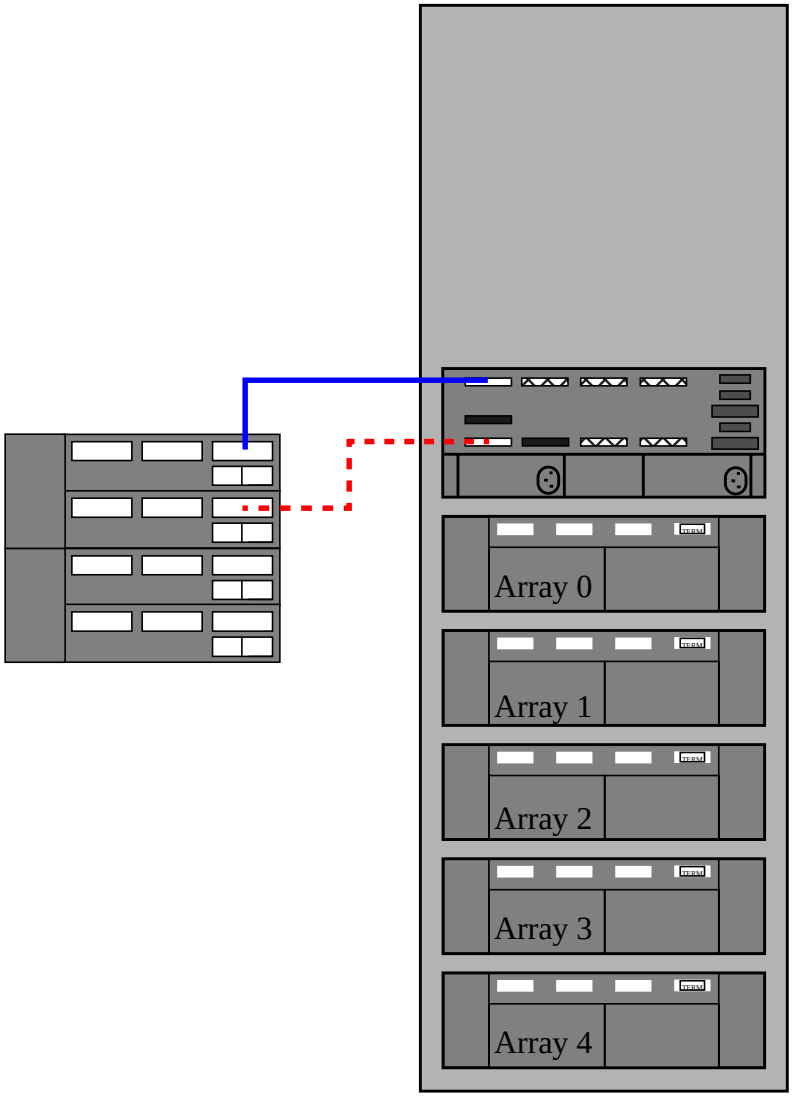

Logical Volume Layout[1] (A3500FC RAID5 Small/Medium Configuration)

A0	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.0	HS
A1	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.M	HS
A2	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V12.0	V14.0
A3	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V12.M	V13.M
A4	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V13.0	V14.M

Controllers	Controller 0	Controller 1

1. The layout naming convention is explained on page 18

Details (A3500FC RAID5 Small/Medium Configuration)

Configuration

RAID Layout

- 10 4+1RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 2 Hot Spares
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Two FiberChannel connections to the host
- RAID Manager 6.2.x software manages A3500FC RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 10 4+1 RAID5 logical volumes @ 36.4 GB (9.1-GB disks)
- 4 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks)
- **Total data capacity is 400.4 GB (9.1-GB disks)**
- **R-value = 3388**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- Dual 100-MB/s FiberChannel connections to host (200-MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(2) FCAL 100-MB host bus adapters (Sbus: X6730A PCI: X6729A)
(2) 68-pin differential terminators (150-1890)

Configuration

Subsystem

(1) A3500FC - SG-XARY360A-545G - (A3500FC array in 1x5x12 configuration, mounted in 72-inch expansion rack. Includes controller Module with 2 controllers @ 128-MB, 2 12-meter USCSI cables, redundant fans, drives, and power supplies. Populated with 60 9.1-GB 10,000-RPM drives.)

(1) A3500FC - SG-XARY360B-545G - (1 rack, 1 FCAL controller module, 5 disk trays, 60 x 9.1-GB 10K-rpm drives

Configuration with 18.2-GB drives:

SG-XARY380B-1092G (1 rack, 1 FCAL controller module, 5 disk trays, 60 x 18.2-GB 10,000-rpm drives)

Configuration with 36.4-GB drives:

SG-XARY381B-1546G (1 rack, 1 FCAL controller module, 5 disk trays, 40 x 36.4-GB 10,000-rpm drives)

Drive options:

X5235A: 9.1-GB 10,000-rpm
X5238A: 18.2-GB 10,000-rpm
X5240A: 36.4-GB 10,000-rpm

Software

Sun StorEdge RAID Manager 6.2.x or later release
Solaris 2.5.1 (8/97) or above with required OS patches

Other

Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x
Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers
Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix L A3500FC RAID5 Medium/Large Configuration

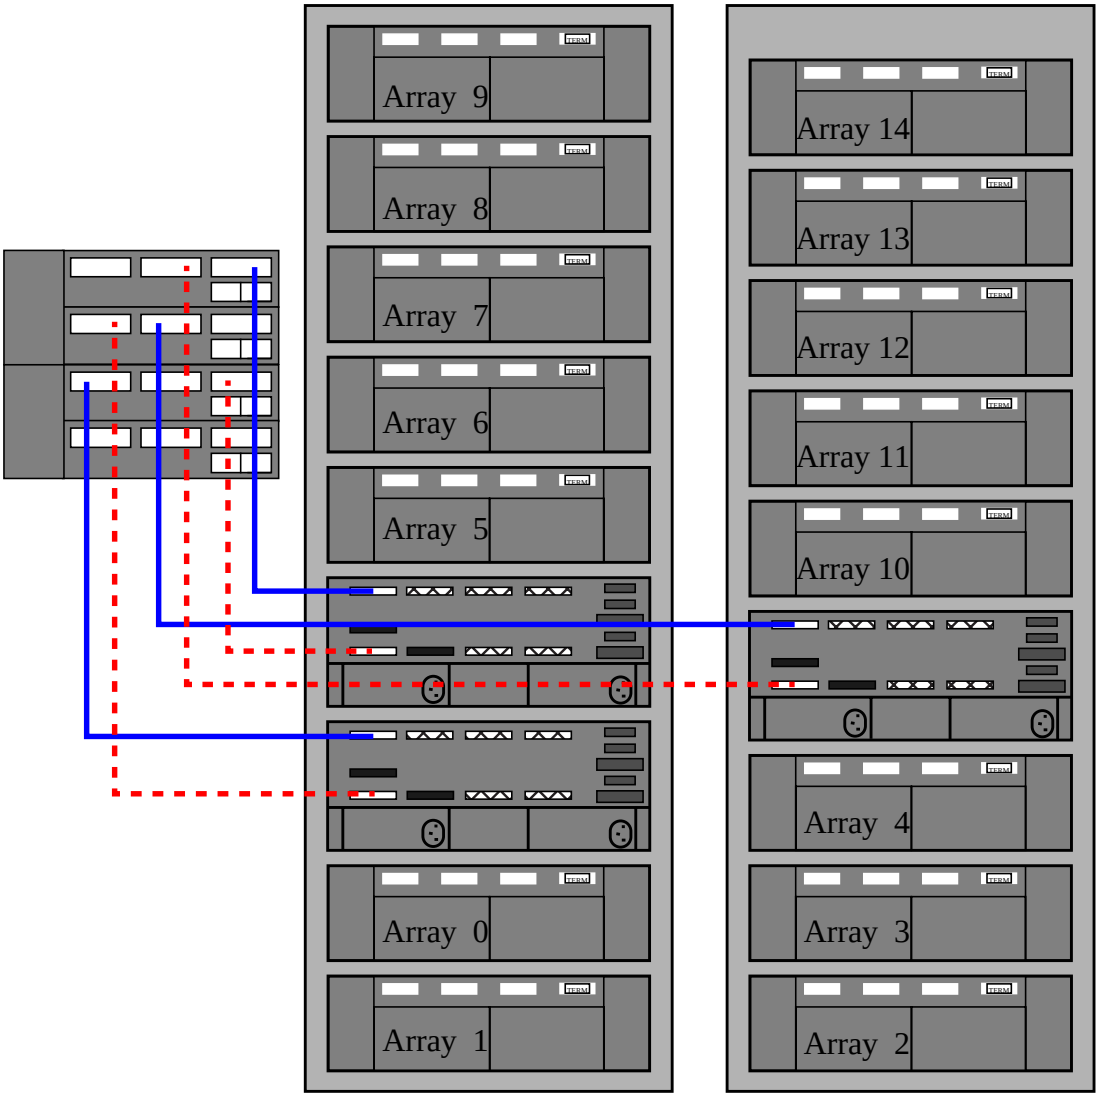

Logical Volume Layout[1] (A3500FC RAID5 Medium/Large Configuration)

A0	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.0	HS
A1	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.M	HS
A2	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V12.0	V14.0
A3	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V12.M	V13.M
A4	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V13.0	V14.M

A5	V15.1	V16.1	V17.1	V18.1	V19.1	V20.1	V21.1	V22.1	V23.1	V24.1	V25.0	HS
A6	V15.2	V16.2	V17.2	V18.2	V19.2	V20.2	V21.2	V22.2	V23.2	V24.2	V25.M	HS
A7	V15.3	V16.3	V17.3	V18.3	V19.3	V20.3	V21.3	V22.3	V23.3	V24.3	V26.0	V28.0
A8	V15.4	V16.4	V17.4	V18.4	V19.4	V20.4	V21.4	V22.4	V23.4	V24.4	V26.M	V27.M
A9	V15.P	V16.P	V17.P	V18.P	V19.P	V20.P	V21.P	V22.P	V23.P	V24.P	V27.0	V28.M

A10	V29.1	V30.1	V31.1	V32.1	V33.1	V34.1	V35.1	V36.1	V37.1	V38.1	V39.0	HS
A11	V29.2	V30.2	V31.2	V32.2	V33.2	V34.2	V35.2	V36.2	V37.2	V38.2	V39.M	HS
A12	V29.3	V30.3	V31.3	V32.3	V33.3	V34.3	V35.3	V36.3	V37.3	V38.3	V40.0	V42.0
A13	V29.4	V30.4	V31.4	V32.4	V33.4	V34.4	V35.4	V36.4	V37.4	V38.4	V40.M	V41.M
A14	V29.P	V30.P	V31.P	V32.P	V33.P	V34.P	V35.P	V36.P	V37.P	V38.P	V41.0	V42.M

Controllers	Controller 0	Controller 1
	Controller 2	Controller 3
	Controller 4	Controller 5

1. The layout naming convention is explained on page 18

Details (A3500FC RAID5 Medium/Large Configuration)

Configuration

RAID Layout

- 30 4+1 RAID5 logical volumes
- 12 1+1 RAID1 logical volumes
- 6 Hot Spares (2 per module, 1 per controller)
- Dual power paths to the array, the controller module and the disk trays
- Dual hot-plug power supplies and cooling units in the controller module and the disk trays
- Battery backup for data cache
- Automatic dispatch failover between controllers
- Six FiberChannel connections to the host
- RAID Manager 6.2.x software manages A3500FC RAID controllers
- Solstice DiskSuite (SDS) or VERITAS Volume Manager (VxVM) used to build additional logical layers on top of hardware RAID volumes
- Box-sharing across two independent Solaris hosts
- Multi-initiator support for Sun Enterprise Clusters
- 72-inch Expansion Rack for good footprint and room for growth

Capacity

- 30 4+1 RAID5 logical volumes @ 36.4 GB (9.1-GB disks)
- 12 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks)
- **Total data capacity is 1201.2 GB (9.1-GB disks)**
- **R-value = 10164**

Considerations

Availability

- Single host connection per controller with automatic controller failover
- Controller cache mirrored between controllers (can be disabled by user for higher performance)
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache
- Environmental Service Module in each tray for rapid troubleshooting

Performance

- Controller-based RAID
- 256-MB accelerator cache per controller
- 6 100-MB/s FiberChannel connections to host (600-MB/s total bandwidth)
- Ultra-SCSI between controllers for fast cache mirroring
- Ultra-SCSI between controllers and disks for increased RAID performance
- 10,000-RPM drives for high performance

Components

Hardware

(6) FCAL 100-MB host bus adapters (Sbus: X6730A PCI: X6729A)
(6) 68-pin differential terminators (150-1890)

Subsystem

(1) A3500FC - SG-ARY374A-273G (2 racks, 15 disk trays, 30x9.1-GB 10K-rpm drives)

(3) FCAL Controller Module (6538A)

Configuration with 18.2-GB drives:

(1) A3500FC - SG-ARY384A-546G (2 racks, 15 disk trays, 30x18.2-GB 10K-rpm drives)
(3) FCAL Controller Module (6538A)

Configuration with 36.4-GB drives:

(1) A3500FC - SG-ARY385A-1092G (2 racks, 15 disk trays, 30x36.4-GB 10K-rpm drives)
(3) FCAL Controller Module (6538A)

Configuration

Software	Sun StorEdge RAID Manager 6.2.x or later release Solaris 2.5.1 (8/97) or above with required OS patches
Other	Sun StorEdge Volume Manager 2.4, 2.5, 2.5.x, and VxVm 3.x Solstice DiskSuite Version 4.1 for Solaris 2.5.1 servers Solstice DiskSuite Version 4.2 for Solaris 2.6 servers and later releases

Appendix M A5100 RAID1 Small/Medium Configuration, Direct Connect

Hardware Layout Diagram

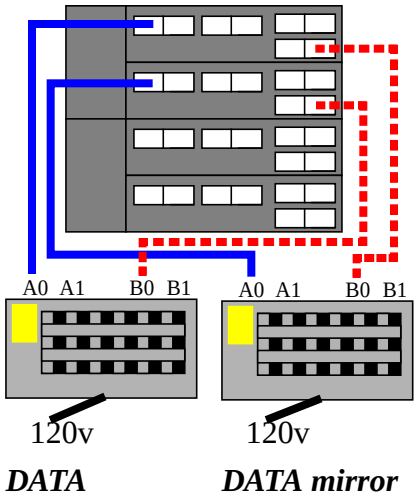

Logical Volume Layout[1] (A5100 RAID1 Small/Medium Configuration, Direct Connect)

Array 0	Front	V1.1	V1.2	V1.3	V2.1	V2.2	V2.3	HS
		Log3.1	Log3.2	Log3.3	Log4.1	Log4.2	Log4.3	
	Back	V3.1	V3.2	V3.3	V4.1	V4.2	V4.3	V5.1
		Log1.1	Log1.2	Log1.3	Log2.1	Log2.2	Log2.3	Log5.1M

Array 1	Front	V1.1M	V1.2M	V1.3M	V2.1M	V2.2M	V2.3M	HS
		Log3.1M	Log3.2M	Log3.3M	Log4.1M	Log4.2M	Log4.3M	
	Back	V3.1M	V3.2M	V3.3M	V4.1M	V4.2M	V4.3M	V5.1M
		Log1.1M	Log1.2M	Log1.3M	Log2.1M	Log2.2M	Log2.3M	Log5.1

1. The layout naming convention is explained on page 18

Details (A5100 RAID1 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout
- 4 3+3 RAID1 logical volumes
- 1 1+1 RAID1 logical volumes
- 2 Hot Spare drives (can be reduced to 1, other 1 used for misc.)
- Striped and Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each array from data center UPS

Capacity
- 4 3+3 RAID1 logical volumes @ 54.6GB (18.2GB disks)
- 1 1+1 RAID1 logical volumes @ 18.2GB (18.2GB disks)
- **Total data capacity is 236.6 GB (18.2-GB disks)**
- **R-value = 1001**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Direct connect to host simplifies fault isolation in case of problems
- ***Failure of power grid will take out array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A
(4) 2-meter cable - X973A (X978A for 15-m cable)

Subsystem
(2) A5100 - SG-XARY550A-509G (28 disk drives in total, raw capacity is 1019.2-GB with 36.4-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, and 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix N A5100 RAID1 Medium/Large Configuration, Direct Connect

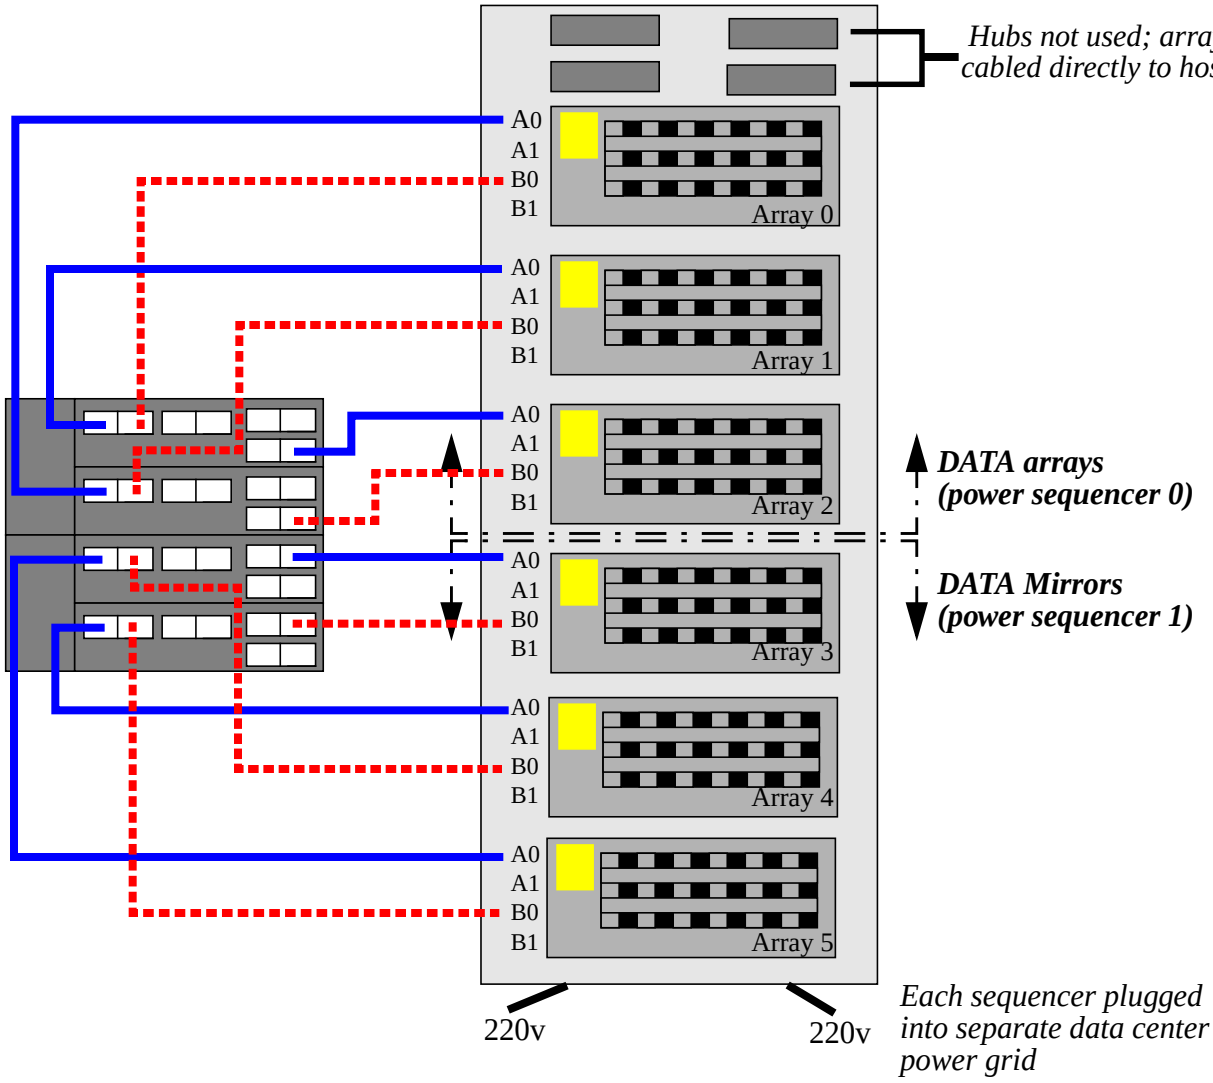

Logical Volume Layout[1] (A5100 RAID1 Medium/Large Configuration, Direct Connect)

Array 0	Front	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1
		Log8.1	Log9.1	Log10.1	Log11.1	Log12.1	Log13.1	Log1.1
	Back	V8.1	V9.1	V10.1	V11.1	V12.1	V13.1	HS
		Log2.1	Log3.1	Log4.1	Log5.1	Log6.1	Log7.1	

Array 1	Front	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2
		Log8.2	Log9.2	Log10.2	Log11.2	Log12.2	Log13.2	Log1.2
	Back	V8.2	V9.2	V10.2	V11.2	V12.2	V13.2	HS
		Log2.2	Log3.2	Log4.2	Log5.2	Log6.2	Log7.2	

Array 2	Front	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3
		Log8.3	Log9.3	Log10.3	Log11.3	Log12.3	Log13.3	Log1.3
	Back	V8.3	V9.3	V10.3	V11.3	V12.3	V13.3	HS
		Log2.3	Log3.3	Log4.3	Log5.3	Log6.3	Log7.3	

Array 3	Front	V1.1M	V2.1M	V3.1M	V4.1M	V5.1M	V6.1M	V7.1M
		Log8.1M	Log9.1M	Log10.1M	Log11.1M	Log12.1M	Log13.1M	Log1.1M
	Back	V8.1M	V9.1M	V10.1M	V11.1M	V12.1M	V13.1M	HS
		Log2.1M	Log3.1M	Log4.1M	Log5.1M	Log6.1M	Log7.1M	

Array 4	Front	V1.2M	V2.2M	V3.2M	V4.2M	V5.2M	V6.2M	V7.2M
		Log8.2M	Log9.2M	Log10.2M	Log11.2M	Log12.2M	Log13.2M	Log1.2M
	Back	V8.2M	V9.2M	V10.2M	V11.2M	V12.2M	V13.2M	HS
		Log2.2M	Log3.2M	Log4.2M	Log5.2M	Log6.2M	Log7.2M	

Array 5	Front	V1.3M	V2.3M	V3.3M	V4.3M	V5.3M	V6.3M	V7.3M
		Log8.3M	Log9.3M	Log10.3M	Log11.3M	Log12.3M	Log13.3M	Log1.3M
	Back	V8.3M	V9.3M	V10.3M	V11.3M	V12.3M	V13.3M	HS
		Log2.3M	Log3.3M	Log4.3M	Log5.3M	Log6.3M	Log7.3M	

1. The layout naming convention is explained on page 18

Details (A5100 RAID1 Medium/Large Configuration, Direct Connect)

Configuration

RAID Layout
- 13 3+3 RAID1 logical volumes
- 6 Hot Spare drives (can be reduced to 2, other 4 used for misc.)
- Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Plug each sequencer into separate data center UPS power grid

Capacity
- 13 3+3 RAID1 logical volumes @ 54.6 GB (18.2-GB disks)
- **Total data capacity is 709.8 GB (18-GB disks)**
- **R-value = 3003**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Direct connect to host simplifies fault isolation in case of problems
- ***Failure of a power sequencer or power grid will take out 3 array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI host adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(1) A5100 - SG-ARY533A-509G
(5) A5100 - SG-ARY551-509GR5
(84) disk drives in total (raw capacity is 3057-GB with 36.4-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix O A5100 RAID1 Medium/Large Configuration, Hubs and Sequencers

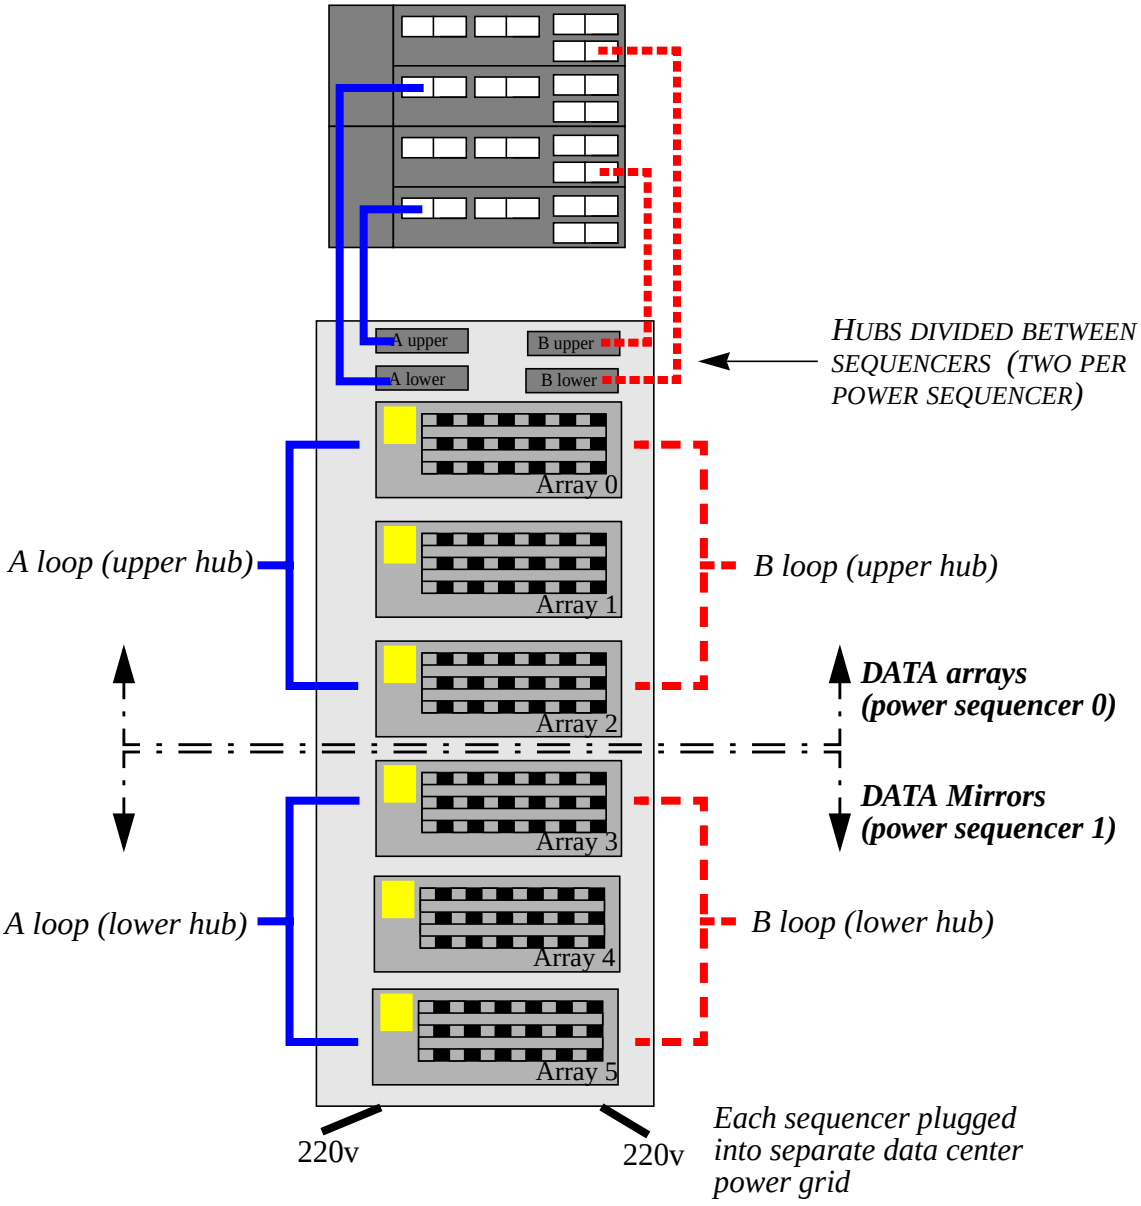

Logical Volume Layout[1] (A5100 RAID1 Medium/Large Configuration, Hubs and Sequencers)

Array 0	Front	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1
		Log8.1	Log9.1	Log10.1	Log11.1	Log12.1	Log13.1	Log1.1
	Back	V8.1	V9.1	V10.1	V11.1	V12.1	V13.1	HS
		Log2.1	Log3.1	Log4.1	Log5.1	Log6.1	Log7.1	

Array 1	Front	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2
		Log8.2	Log9.2	Log10.2	Log11.2	Log12.2	Log13.2	Log1.2
	Back	V8.2	V9.2	V10.2	V11.2	V12.2	V13.2	HS
		Log2.2	Log3.2	Log4.2	Log5.2	Log6.2	Log7.2	

Array 2	Front	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3
		Log8.3	Log9.3	Log10.3	Log11.3	Log12.3	Log13.3	Log1.3
	Back	V8.3	V9.3	V10.3	V11.3	V12.3	V13.3	HS
		Log2.3	Log3.3	Log4.3	Log5.3	Log6.3	Log7.3	

Array 3	Front	V1.1M	V2.1M	V3.1M	V4.1M	V5.1M	V6.1M	V7.1M
		Log8.1M	Log9.1M	Log10.1M	Log11.1M	Log12.1M	Log13.1M	Log1.1M
	Back	V8.1M	V9.1M	V10.1M	V11.1M	V12.1M	V13.1M	HS
		Log2.1M	Log3.1M	Log4.1M	Log5.1M	Log6.1M	Log7.1M	

Array 4	Front	V1.2M	V2.2M	V3.2M	V4.2M	V5.2M	V6.2M	V7.2M
		Log8.2M	Log9.2M	Log10.2M	Log11.2M	Log12.2M	Log13.2M	Log1.2M
	Back	V8.2M	V9.2M	V10.2M	V11.2M	V12.2M	V13.2M	HS
		Log2.2M	Log3.2M	Log4.2M	Log5.2M	Log6.2M	Log7.2M	

Array 5	Front	V1.3M	V2.3M	V3.3M	V4.3M	V5.3M	V6.3M	V7.3M
		Log8.3M	Log9.3M	Log10.3M	Log11.3M	Log12.3M	Log13.3M	Log1.3M
	Back	V8.3M	V9.3M	V10.3M	V11.3M	V12.3M	V13.3M	HS
		Log2.3M	Log3.3M	Log4.3M	Log5.3M	Log6.3M	Log7.3M	

1. The layout naming convention is explained on page 18

Details (A5100 RAID1 Medium/Large Configuration, Hubs and Sequencers)

Configuration

RAID Layout
- 13 3+3 RAID1 logical volumes
- 6 Hot Spare drives (can be reduced to 2, other 4 used for misc.)
- Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each sequencer from data center UPS

Capacity
- 13 3+3 RAID1 logical volumes @ 54.6 GB (18.2-GB disks)
- **Total data capacity is 709.8 GB (18-GB disks)**
- **R-value = 3003**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Host connection through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***Failure of a power sequencer or power grid will take out 3 array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A

Subsystem
(1) A5100 - SG-ARY533A-3057G (6 x 509.6-GB in a 72-inch rack)
(84) disk drives in total (raw capacity is 3057-GB with 36.4-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix P A5100 RAID5 Small/Medium Configuration, Direct Connect

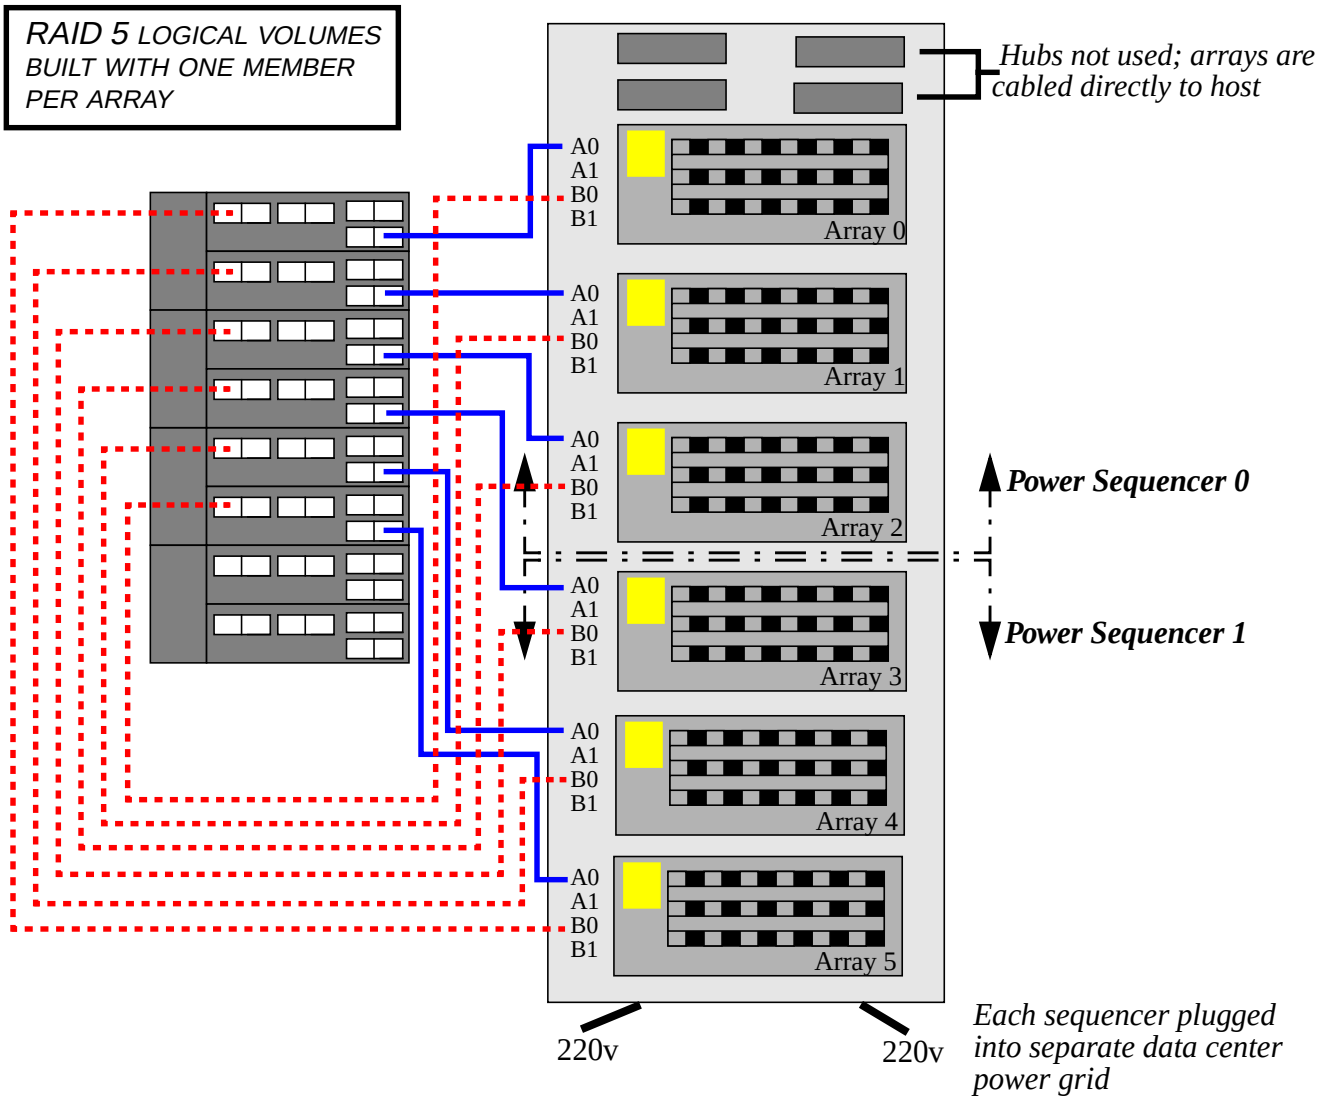

Logical Volume Layout[1] (A5100 RAID5 Small/Medium Configuration, Direct Connect)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1
	B	V8.1	V9.1	V10.1	V11.1	V12.1	V13.0	HS
A1	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3
	B	V8.3	V9.3	V10.3	V11.3	V12.3	V14.0	HS
A2	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5
	B	V8.5	V9.5	V10.5	V11.5	V12.5	V15.0	V16.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2
	B	V8.2	V9.2	V10.2	V11.2	V12.2	V13.M	HS
A4	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4
	B	V8.4	V9.4	V10.4	V11.4	V12.4	V14.M	HS
A5	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P
	B	V8.P	V9.P	V10.P	V11.P	V12.P	V15.M	V16.M

1. The layout naming convention is explained on page 18

Details (A5100 RAID5 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout
- 12 5+1 RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 4 Hot Spare drives (can be reduced to 2, other 2 used for misc.)
- 6 Arrays — one RAID5 member per array, RAID1 members mirrored between adjacent arrays
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each Sequencer from data center UPS

Capacity
- 12 5+1 RAID1 logical volumes @ 91 GB (18.2-GB disks)
- 4 1+1 RAID1 logical volumes @ 18.2 GB (18.2-GB disks)
- **Total data capacity is 1164.8 GB (18.2-GB disks)**
- **R-value = 3248**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI Host Adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(1) A5100 - SG-ARY533-509G
(5) A5100 - SG-ARY551-509GR5
(84) disk drives in total (raw capacity is 3057-GB with 36.4-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix Q A5100 RAID5 Small/Medium Configuration, Hubs and Sequencers

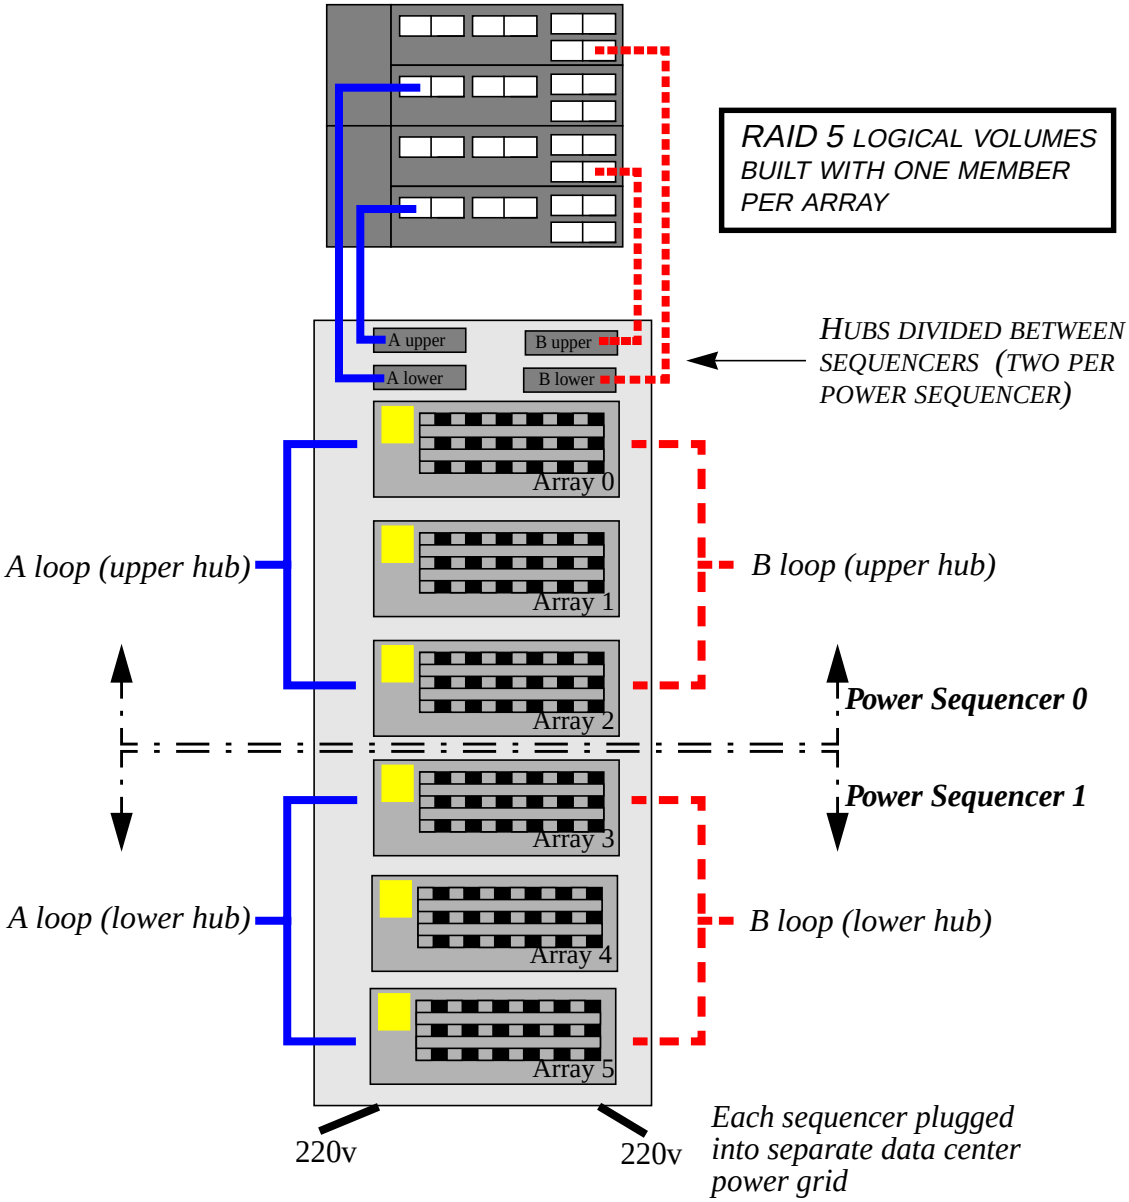

Logical Volume Layout[1] (A5100 RAID5 Small/Medium Configuration, Hubs and Sequencers)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1
	B	V8.1	V9.1	V10.1	V11.1	V12.1	V13.0	HS
A1	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3
	B	V8.3	V9.3	V10.3	V11.3	V12.3	V14.0	HS
A2	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5
	B	V8.5	V9.5	V10.5	V11.5	V12.5	V15.0	V16.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2
	B	V8.2	V9.2	V10.2	V11.2	V12.2	V13.M	HS
A4	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4
	B	V8.4	V9.4	V10.4	V11.4	V12.4	V14.M	HS
A5	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P
	B	V8.P	V9.P	V10.P	V11.P	V12.P	V15.M	V16.M

1. The layout naming convention is explained on page 18

Details (A5100 RAID5 Small/Medium Configuration, Hubs and Sequencers)

Configuration

RAID Layout

- 12 5+1 RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 4 Hot Spare drives (can be reduced to 2, other 2 used for misc.)
- 6 Arrays — one RAID5 member per array, RAID1 members mirrored between sequencers, hot spares available in all sequencers
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each sequencer from data center UPS

Capacity

- 12 5+1 RAID1 logical volumes @ 91 GB (18.2-GB disks)
- 4 1+1 RAID1 logical volumes @ 18.2 GB (18.2-GB disks)
- **Total data capacity is 1164.8 GB (18.2-GB disks)**
- **R-value = 3248**

Considerations

Availability

- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- Host connection through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available.***

Performance

- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware

(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A

Subsystem

(1) A5100 - SG-ARY533A-3057G (6 x 509.6-GB in a 72-inch rack)
(84) disk drives in total (raw capacity is 3057 GB with 36.4-GB disks)

Software

Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other

N/A

Appendix R A5100 RAID5 Medium/Large Configuration, Hubs and Sequencers

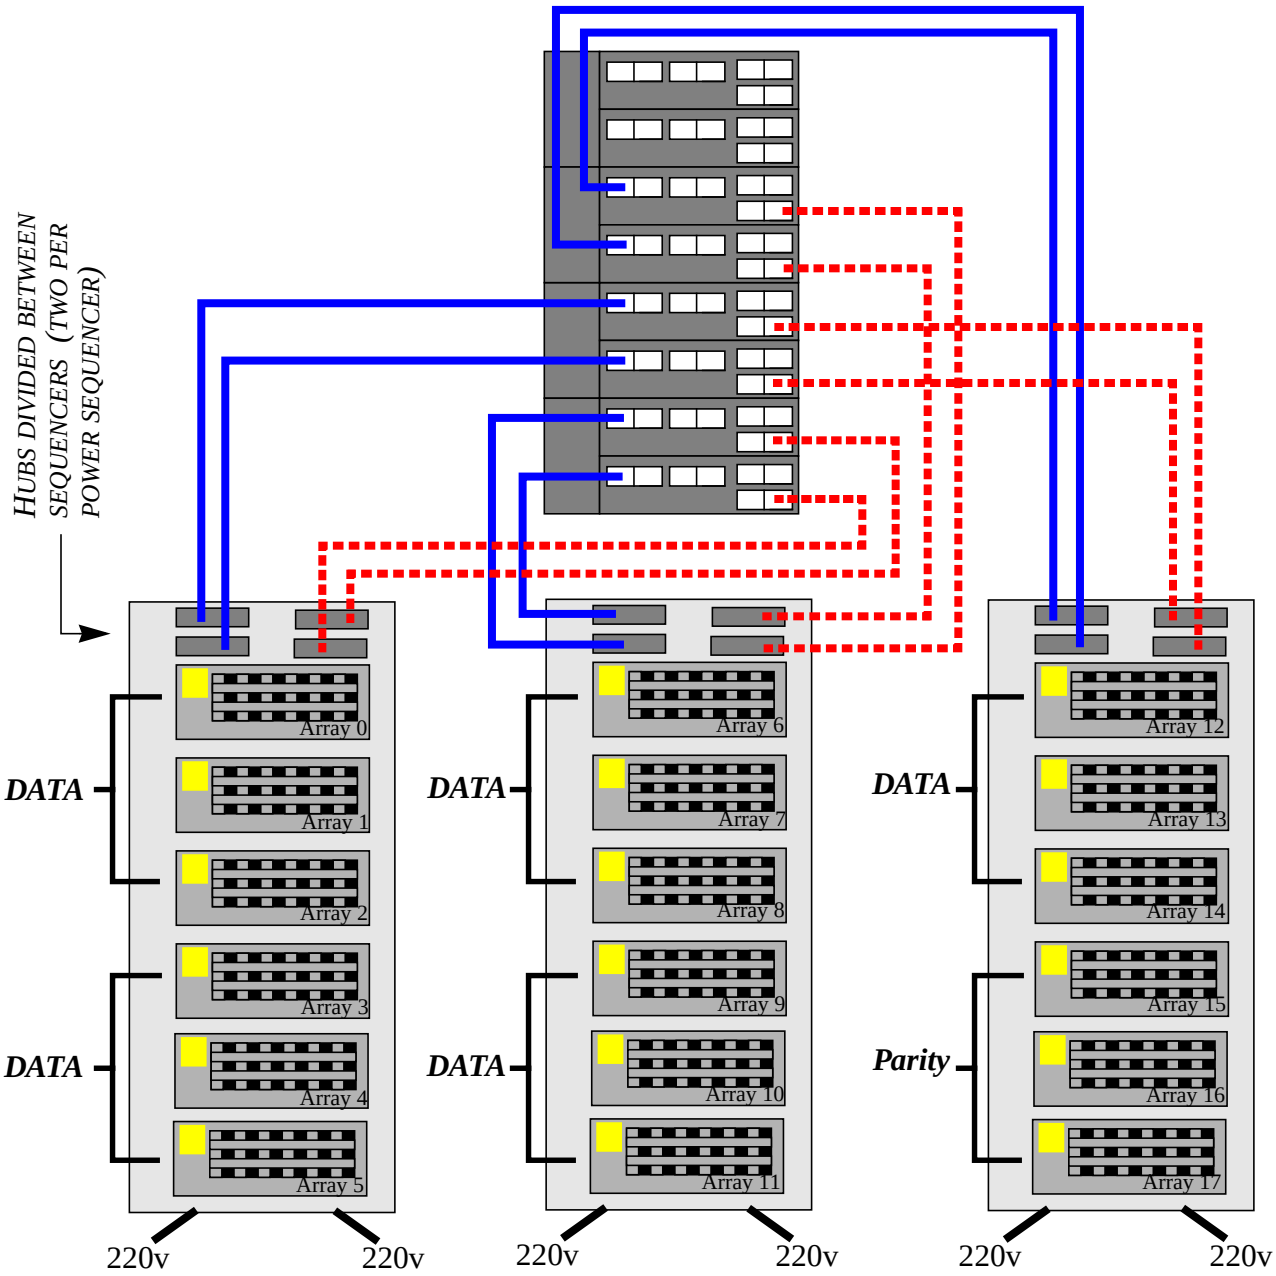

Sequencer 0 and Sequencer 1 of each rack go to a different data center power grid

Logical Volume Layout[1] (A5100 RAID5 Medium/Large Configuration, Hubs and Sequencers)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1
	B	V8.1	V9.1	V10.1	V11.1	V12.1	V37.0	HS
A1	F	V13.1	V14.1	V15.1	V16.1	V17.1	V18.1	V19.1
	B	V20.1	V21.1	V22.1	V23.1	V24.1	V38.0	HS
A2	F	V25.1	V26.1	V27.1	V28.1	V29.1	V30.1	V31.1
	B	V32.1	V33.1	V34.1	V35.1	V36.1	V39.0	V40.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2
	B	V8.2	V9.2	V10.2	V11.2	V12.2	V37.M	HS
A4	F	V13.2	V14.2	V15.2	V16.2	V17.2	V18.2	V19.2
	B	V20.2	V21.2	V22.2	V23.2	V24.2	V38.M	HS
A5	F	V25.2	V26.2	V27.2	V28.2	V29.2	V30.2	V31.2
	B	V32.2	V33.2	V34.2	V35.2	V36.2	V39.M	V40.M
A6	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3
	B	V8.3	V9.3	V10.3	V11.3	V12.3	V41.0	HS
A7	F	V13.3	V14.3	V15.3	V16.3	V17.3	V18.3	V19.3
	B	V20.3	V21.3	V22.3	V23.3	V24.3	V42.0	HS
A8	F	V25.3	V26.3	V27.3	V28.3	V29.3	V30.3	V31.3
	B	V32.3	V33.3	V34.3	V35.3	V36.3	V43.0	V44.0
A9	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4
	B	V8.4	V9.4	V10.4	V11.4	V12.4	V41.M	HS
A10	F	V13.4	V14.4	V15.4	V16.4	V17.4	V18.4	V19.4
	B	V20.4	V21.4	V22.4	V23.4	V24.4	V42.M	HS
A11	F	V25.4	V26.4	V27.4	V28.4	V29.4	V30.4	V31.4
	B	V32.4	V33.4	V34.4	V35.4	V36.4	V43.M	V44.M
A12	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5
	B	V8.5	V9.5	V10.5	V11.5	V12.5	V45.0	HS
A13	F	V13.5	V14.5	V15.5	V16.5	V17.5	V18.5	V19.5
	B	V20.5	V21.5	V22.5	V23.5	V24.5	V46.0	HS
A14	F	V25.5	V26.5	V27.5	V28.5	V29.5	V30.5	V31.5
	B	V32.5	V33.5	V34.5	V35.5	V36.5	V47.0	V48.0
A15	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P
	B	V8.P	V9.P	V10.P	V11.P	V12.P	V45.M	HS
A16	F	V13.P	V14.P	V15.P	V16.P	V17.P	V18.P	V19.P
	B	V20.P	V21.P	V22.P	V23.P	V24.P	V46.M	HS
A17	F	V25.P	V26.P	V27.P	V28.P	V29.P	V30.P	V31.P
	B	V32.P	V33.P	V34.P	V35.P	V36.P	V47.M	V48.M

1. The layout naming convention is explained on page 18

Details (A5100 RAID5 Medium/Large Configuration, Hubs and Sequencers)

Configuration

RAID Layout
- 36 5+1 RAID5 logical volumes
- 12 1+1 RAID1 logical volumes
- 12 Hot Spare drives (can be reduced to 6, other 6 used for misc.)
- 3 Racks with 6 Arrays each — one RAID5 member per sequencer, RAID1 members mirrored between sequencers, hot spares available in all sequencers
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each sequencer from data center UPS

Capacity
- 36 5+1 RAID1 logical volumes @ 91 GB (18.2-GB disks)
- 12 1+1 RAID1 logical volumes @ 18.2 GB (18.2-GB disks)
- **Total data capacity is 3494.4 GB (18.2-GB disks)**
- **R-value = 9744**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- Host connection through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available. To maintain availability through loss of power, six independent power grids are necessary!***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI Host Adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(3) A5100 - SG-ARY533A-3057G (6 x 509.6-GB in a 72-inch rack)
(252) disk drives in total (raw capacity is 9171-GB with 36.4-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix S A5200 RAID1 Small/Medium Configuration, Direct Connect

Hardware Layout Diagram

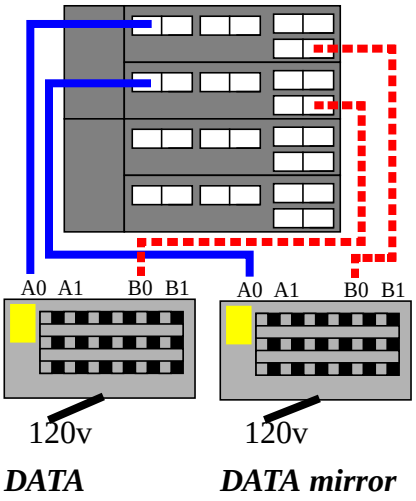

Logical Volume Layout[1] (A5200 RAID1 Small/Medium Configuration, Direct Connect)

Array 0	Front	V1.1	V1.2	V1.3	V2.1	V2.2	V2.3	V3.1	V3.2	V3.3	V7.1	HS
		Log5.1	Log5.2	Log5.3	Log6.1	Log6.2	Log6.3	Log7.1	Log7.2	Log7.3	Log4.1	
	Back	V4.1	V4.2	V4.3	V5.1	V5.2	V5.3	V6.1	V6.2	V6.3	V7.2	V7.3
		Log3.1	Log3.2	Log3.3	Log2.1	Log2.2	Log2.3	Log1.1	Log1.2	Log1.3	Log4.2	Log4.3

Array 1	Front	V1.1M	V1.2M	V1.3M	V2.1M	V2.2M	V2.3M	V3.1M	V3.2M	V3.3M	V7.1M	HS
		Log5.1M	Log5.2M	Log5.3M	Log6.1M	Log6.2M	Log6.3M	Log7.1M	Log7.2M	Log7.3M	Log4.1M	
	Back	V4.1M	V4.2M	V4.3M	V5.1M	V5.2M	V5.3M	V6.1M	V6.2M	V6.3M	V7.2M	V7.3M
		Log3.1M	Log3.2M	Log3.3M	Log2.1M	Log2.2M	Log2.3M	Log1.1M	Log1.2M	Log1.3M	Log4.2M	Log4.3M

1. The layout naming convention is explained on page 18

Details (A5200 RAID1 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout
- 7 3+3 RAID1 logical volumes
- 2 Hot Spare drives (can be reduced to 1, other 1 used for misc.)
- Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each array from data center UPS

Capacity
- 7 3+3 RAID1 logical volumes @ 27.3 GB (9.1-GB disks)
- **Total data capacity is 191.1 GB (9.1-GB disks)**
- **R-value = 1617**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Direct connect to host simplifies fault isolation in case of problems
- ***Failure of power grid will take out array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A
(4) 2-meter cable - X973A (X978A for 15-m cable)

Subsystem
(2) A5200 - SG-XARY520A-200G (200.2-GB)
(44) disk drives in total (raw capacity is 400.4-GB with 9.1-GB disks)
(2) A5200 - SG-XARY540A-400G
(44) disk drives in total (raw capacity is 800.8-GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix T A5200 RAID1 Medium/Large Configuration, Direct Connect

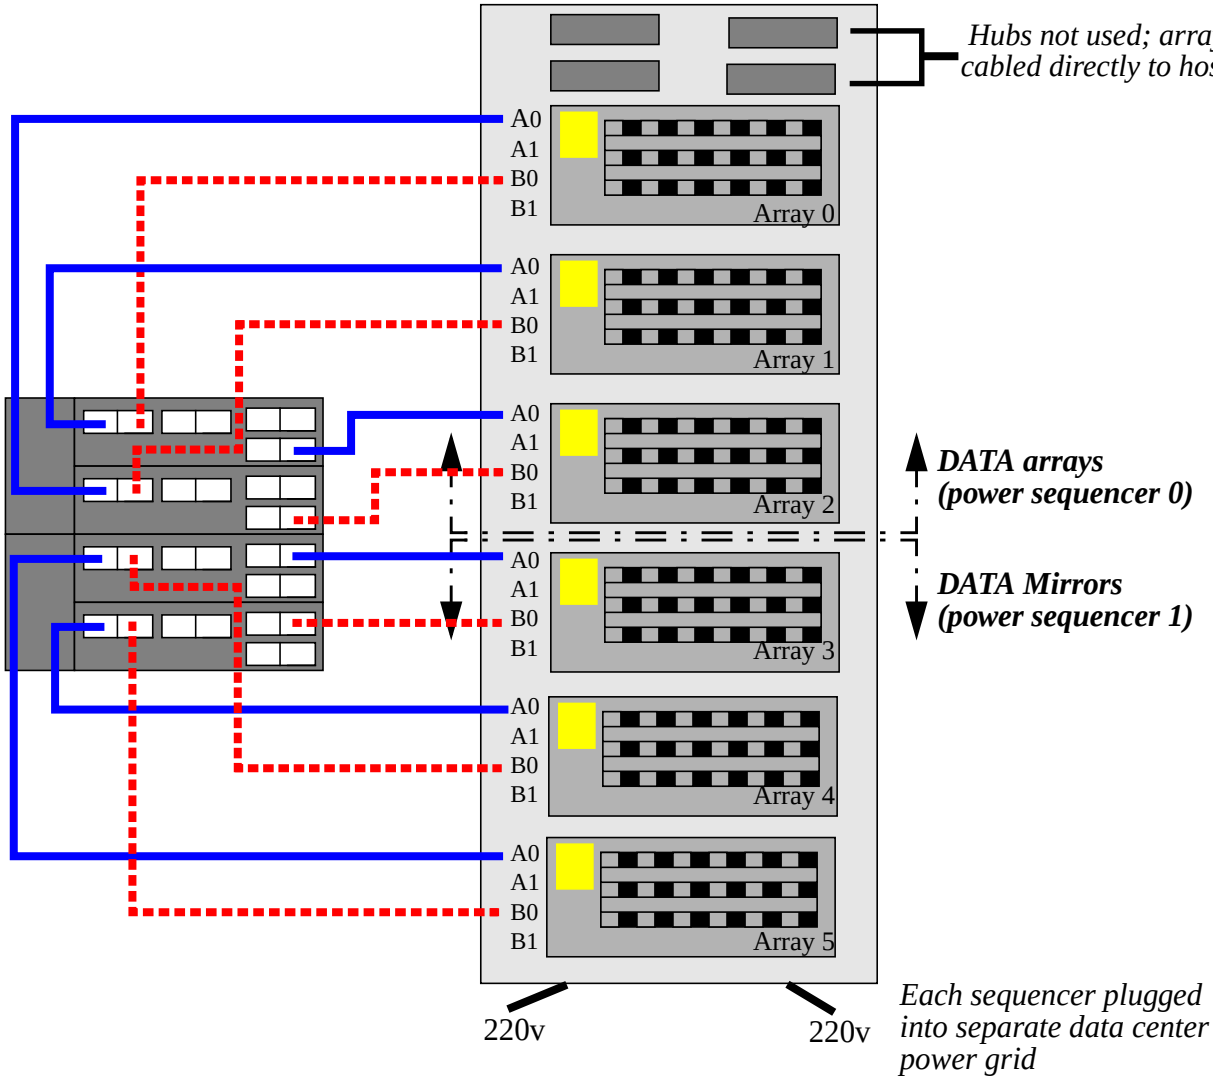

Logical Volume Layout[1] (A5200 RAID1 Medium/Large Configuration, Direct Connect)

Array 0	Front	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.1
		Log12.1	Log13.1	Log14.1	Log15.1	Log16.1	Log17.1	Log18.1	Log19.1	Log20.1	Log21.1	Log1.1
	Back	V12.1	V13.1	V14.1	V115.1	V116.1	V17.1	V18.1	V19.1	V20.1	V21.1	HS
		Log2.1	Log3.1	Log4.1	Log5.1	Log6.1	Log7.1	Log8.1	Log9.1	Log10.1	Log11.1	
Array 1	Front	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.2
		Log12.2	Log13.2	Log14.2	Log15.2	Log16.2	Log17.2	Log18.2	Log19.2	Log20.2	Log21.2	Log1.2
	Back	V12.2	V13.2	V14.2	V115.2	V116.2	V17.2	V18.2	V19.2	V20.2	V21.2	HS
		Log2.2	Log3.2	Log4.2	Log5.2	Log6.2	Log7.2	Log8.2	Log9.2	Log10.2	Log11.2	
Array 2	Front	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V11.3
		Log12.3	Log13.3	Log14.3	Log15.3	Log16.3	Log17.3	Log18.3	Log19.3	Log20.3	Log21.3	Log1.3
	Back	V12.3	V13.3	V14.3	V115.3	V116.3	V17.3	V18.3	V19.3	V20.3	V21.3	HS
		Log2.3	Log3.3	Log4.3	Log5.3	Log6.3	Log7.3	Log8.3	Log9.3	Log10.3	Log11.3	
Array 3	Front	V1.1M	V2.1M	V3.1M	V4.1M	V5.1M	V6.1M	V7.1M	V8.1M	V9.1M	V10.1M	V11.1M
		Log12.1M	Log13.1M	Log14.1M	Log15.1M	Log16.1M	Log17.1M	Log18.1M	Log19.1M	Log20.1M	Log21.1M	Log1.1M
	Back	V12.1M	V13.1M	V14.1M	V115.1M	V116.1M	V17.1M	V18.1M	V19.1M	V20.1M	V21.1M	HS
		Log2.1M	Log3.1M	Log4.1M	Log5.1M	Log6.1M	Log7.1M	Log8.1M	Log9.1M	Log10.1M	Log11.1M	
Array 4	Front	V1.2M	V2.2M	V3.2M	V4.2M	V5.2M	V6.2M	V7.2M	V8.2M	V9.2M	V10.2M	V11.2M
		Log12.2M	Log13.2M	Log14.2M	Log15.2M	Log16.2M	Log17.2M	Log18.2M	Log19.2M	Log20.2M	Log21.2M	Log1.2M
	Back	V12.2M	V13.2M	V14.2M	V115.2M	V116.2M	V17.2M	V18.2M	V19.2M	V20.2M	V21.2M	HS
		Log2.2M	Log3.2M	Log4.2M	Log5.2M	Log6.2M	Log7.2M	Log8.2M	Log9.2M	Log10.2M	Log11.2M	
Array 5	Front	V1.3M	V2.3M	V3.3M	V4.3M	V5.3M	V6.3M	V7.3M	V8.3M	V9.3M	V10.3M	V11.3M
		Log12.3M	Log13.3M	Log14.3M	Log15.3M	Log16.3M	Log17.3M	Log18.3M	Log19.3M	Log20.3M	Log21.3M	Log1.3M
	Back	V12.3M	V13.3M	V14.3M	V115.3M	V116.3M	V17.3M	V18.3M	V19.3M	V20.3M	V21.3M	HS
		Log2.3M	Log3.3M	Log4.3M	Log5.3M	Log6.3M	Log7.3M	Log8.3M	Log9.3M	Log10.3M	Log11.3M	

1. The layout naming convention is explained on page 18

Details (A5200 RAID1 Medium/Large Configuration, Direct Connect)

Configuration

RAID Layout
- 21 3+3 RAID1 logical volumes
- 6 Hot Spare drives (can be reduced to 2, other 4 used for misc.)
- Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each Sequencer from data center UPS

Capacity
- 21 3+3 RAID1 logical volumes @ 27.3 GB (9.1-GB disks)
- **Total data capacity is 573.3 GB (9.1-GB disks)**
- **R-value = 4851**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Direct connect to host simplifies fault isolation in case of problems
- ***Failure of a power sequencer or power grid will take out 3 array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI Host Adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(1) A5200 - SG-ARY523A- 200G
(5) A5200 - SG-ARY521A- 200GR5 (Adds up to 6 x 200.2 GB in a 72-inch rack)
(132) disk drives in total (raw capacity is 1201.2 GB with 9.1-GB disks)
(1) A5200 - SG-ARY543A- 400G
(5) A5200 - SG-ARY541A- 400GR5 (Adds up to 6 x 400.4 GB in a 72-inch rack)
(132) disk drives in total (raw capacity is 2400-GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix U A5200 RAID1 Medium/Large Configuration, Hubs and Sequencers

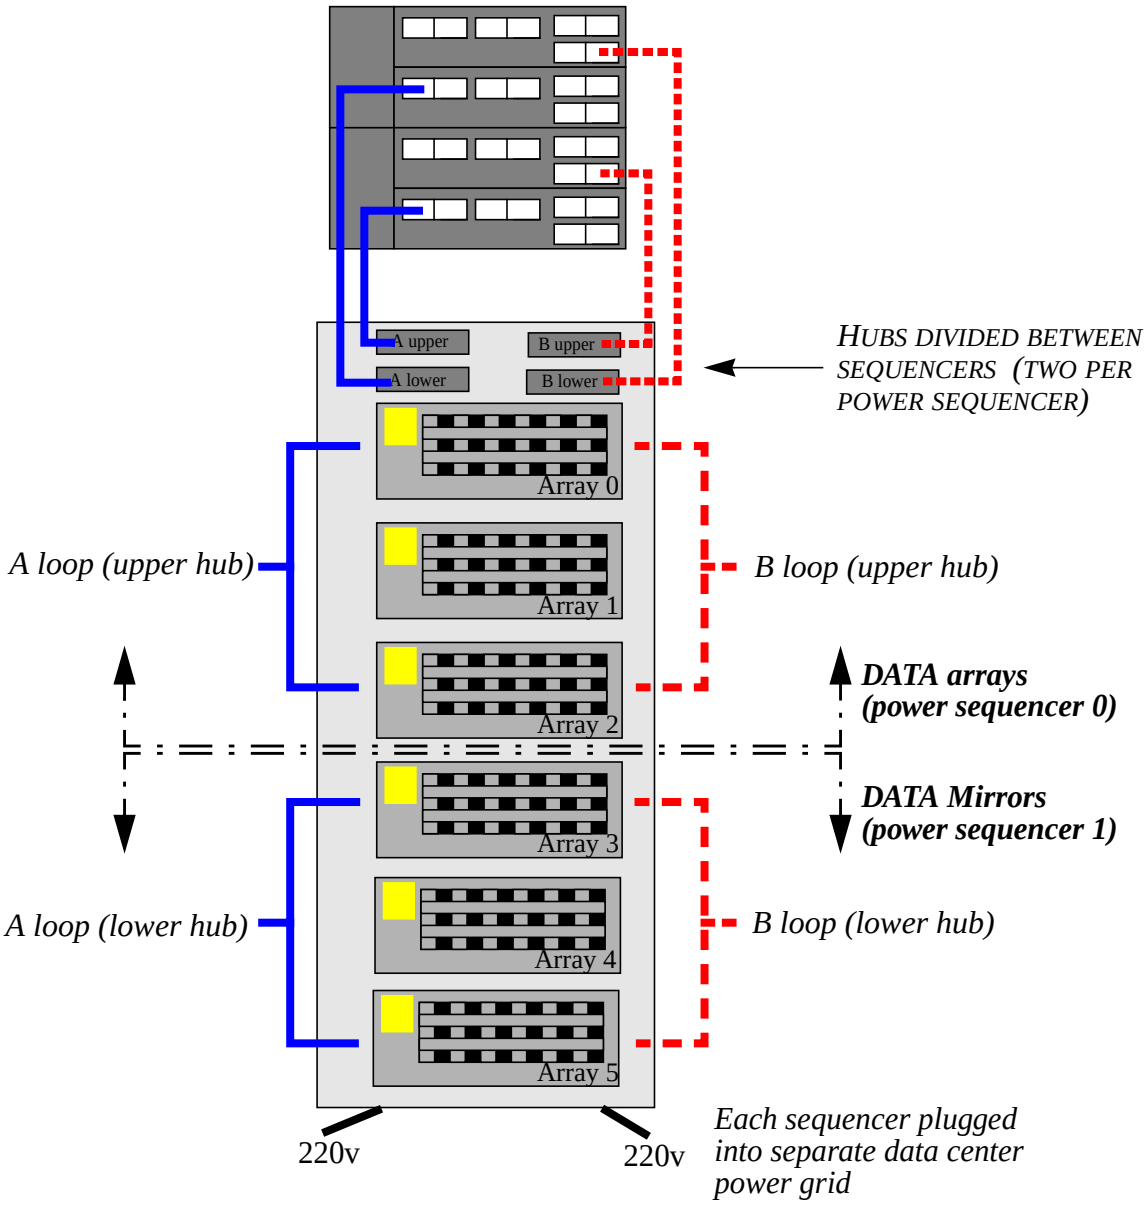

Logical Volume Layout[1] (A5200 RAID1 Medium/Large Configuration, Hubs and Sequencers)

Array 0	Front	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.1
		Log12.1	Log13.1	Log14.1	Log15.1	Log16.1	Log17.1	Log18.1	Log19.1	Log20.1	Log21.1	Log1.1
	Back	V12.1	V13.1	V14.1	V115.1	V116.1	V17.1	V18.1	V19.1	V20.1	V21.1	HS
		Log2.1	Log3.1	Log4.1	Log5.1	Log6.1	Log7.1	Log8.1	Log9.1	Log10.1	Log11.1	

Array 1	Front	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.2
		Log12.2	Log13.2	Log14.2	Log15.2	Log16.2	Log17.2	Log18.2	Log19.2	Log20.2	Log21.2	Log1.2
	Back	V12.2	V13.2	V14.2	V115.2	V116.2	V17.2	V18.2	V19.2	V20.2	V21.2	HS
		Log2.2	Log3.2	Log4.2	Log5.2	Log6.2	Log7.2	Log8.2	Log9.2	Log10.2	Log11.2	

Array 2	Front	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V11.3
		Log12.3	Log13.3	Log14.3	Log15.3	Log16.3	Log17.3	Log18.3	Log19.3	Log20.3	Log21.3	Log1.3
	Back	V12.3	V13.3	V14.3	V115.3	V116.3	V17.3	V18.3	V19.3	V20.3	V21.3	HS
		Log2.3	Log3.3	Log4.3	Log5.3	Log6.3	Log7.3	Log8.3	Log9.3	Log10.3	Log11.3	

Array 3	Front	V1.1M	V2.1M	V3.1M	V4.1M	V5.1M	V6.1M	V7.1M	V8.1M	V9.1M	V10.1M	V11.1M
		Log12.1M	Log13.1M	Log14.1M	Log15.1M	Log16.1M	Log17.1M	Log18.1M	Log19.1M	Log20.1M	Log21.1M	Log1.1M
	Back	V12.1M	V13.1M	V14.1M	V115.1M	V116.1M	V17.1M	V18.1M	V19.1M	V20.1M	V21.1M	HS
		Log2.1M	Log3.1M	Log4.1M	Log5.1M	Log6.1M	Log7.1M	Log8.1M	Log9.1M	Log10.1M	Log11.1M	

Array 4	Front	V1.2M	V2.2M	V3.2M	V4.2M	V5.2M	V6.2M	V7.2M	V8.2M	V9.2M	V10.2M	V11.2M
		Log12.2M	Log13.2M	Log14.2M	Log15.2M	Log16.2M	Log17.2M	Log18.2M	Log19.2M	Log20.2M	Log21.2M	Log1.2M
	Back	V12.2M	V13.2M	V14.2M	V115.2M	V116.2M	V17.2M	V18.2M	V19.2M	V20.2M	V21.2M	HS
		Log2.2M	Log3.2M	Log4.2M	Log5.2M	Log6.2M	Log7.2M	Log8.2M	Log9.2M	Log10.2M	Log11.2M	

Array 5	Front	V1.3M	V2.3M	V3.3M	V4.3M	V5.3M	V6.3M	V7.3M	V8.3M	V9.3M	V10.3M	V11.3M
		Log12.3M	Log13.3M	Log14.3M	Log15.3M	Log16.3M	Log17.3M	Log18.3M	Log19.3M	Log20.3M	Log21.3M	Log1.3M
	Back	V12.3M	V13.3M	V14.3M	V115.3M	V116.3M	V17.3M	V18.3M	V19.3M	V20.3M	V21.3M	HS
		Log2.3M	Log3.3M	Log4.3M	Log5.3M	Log6.3M	Log7.3M	Log8.3M	Log9.3M	Log10.3M	Log11.3M	

1. The layout naming convention is explained on page 18

Details (A5200 RAID1 Medium/Large Configuration, Hubs and Sequencers)

Configuration

RAID Layout
- 21 3+3 RAID1 logical volumes
- 6 Hot Spare drives (can be reduced to 2, other 4 used for misc.)
- Mirrored DRL logs
- Array pairs — each array is mirrored to a separate array and loop pair
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each array from data center UPS

Capacity
- 21 3+3 RAID1 logical volumes @ 27.3 GB (9.1-GB disks)
- **Total data capacity is 573.3 GB (9.1-GB disks)**
- **R-value = 4851**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Using Mirrored Dirty Region Logging to insure data integrity and decrease the resync times of mirrors after system crash, power failure or other unclean shutdown.
- Host connect through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***Failure of a power sequencer or power grid will take out 3 array mirrors, causing significant resync activity on resolution. Dirty Region Logs help reduce resynch activity.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A

Subsystem
(1) A5200 - SG-ARY543A- 2400G (6 x 400.4 GB in a 72-inch rack)
(132) disk drives in total (raw capacity is 2402.4 GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix V A5200 RAID5 Small/Medium Configuration, Direct Connect

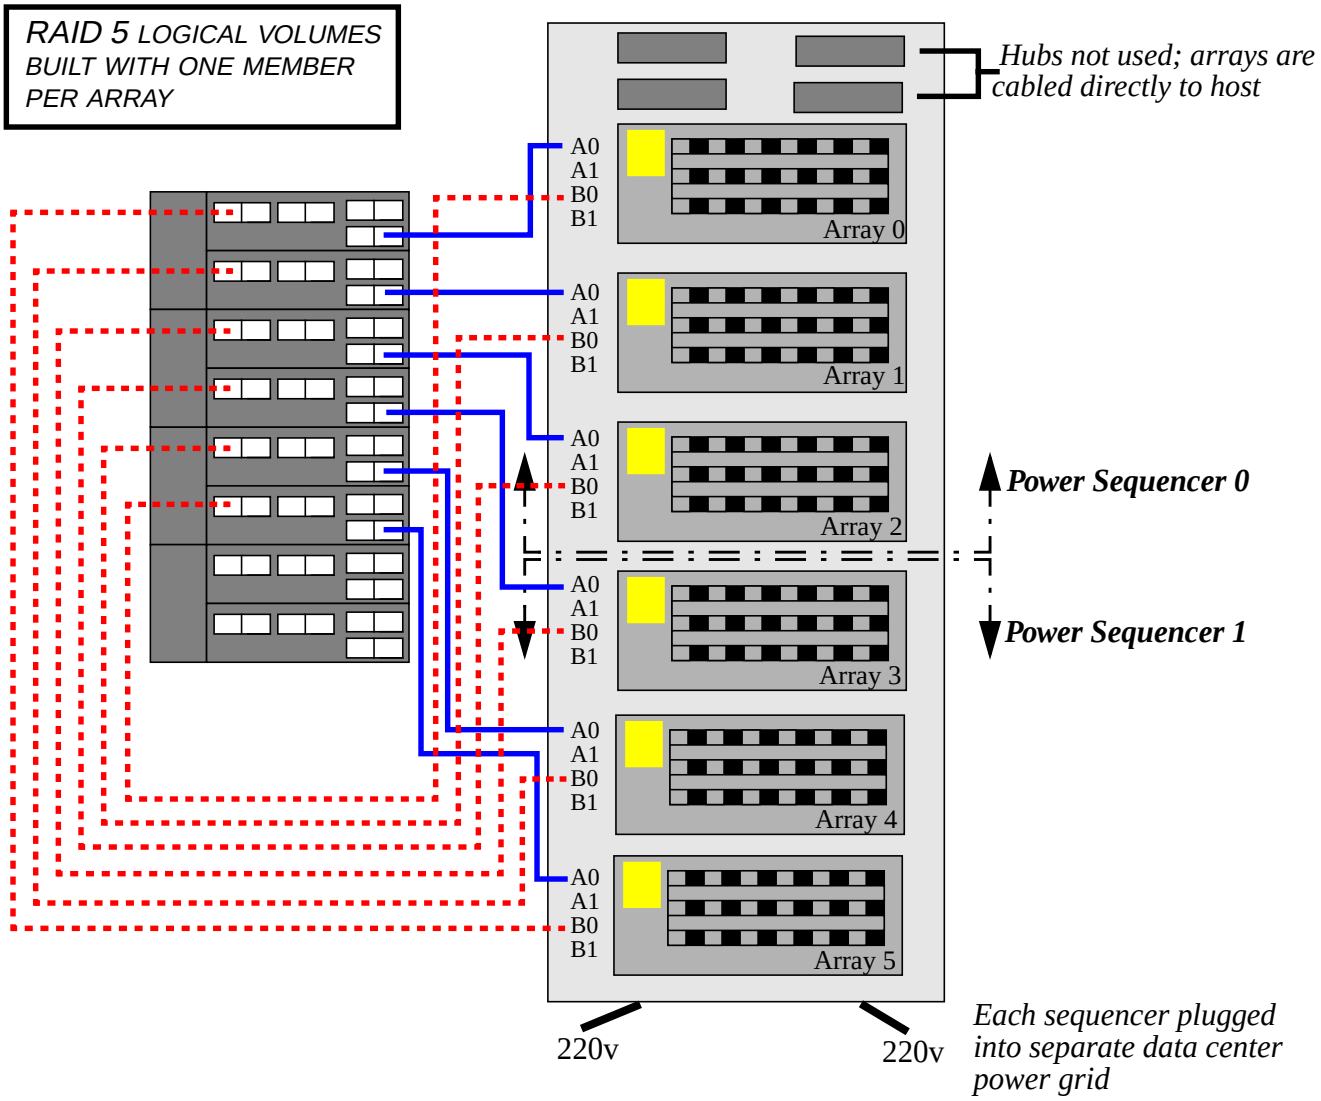

Logical Volume Layout[1] (A5200 RAID5 Small/Medium Configuration, Direct Connect)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.1
	B	V12.1	V13.1	V14.1	V15.1	V16.1	V17.1	V18.1	V19.1	V20.1	V21.0	HS
A1	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V11.3
	B	V12.3	V13.3	V14.3	V15.3	V16.3	V17.3	V18.3	V19.3	V20.3	V22.0	HS
A2	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5	V8.5	V9.5	V10.5	V11.5
	B	V12.5	V13.5	V14.5	V15.5	V16.5	V17.5	V18.5	V19.5	V20.5	V23.0	V24.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.2
	B	V12.2	V13.2	V14.2	V15.2	V16.2	V17.2	V18.2	V19.2	V20.2	V21.M	HS
A4	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V11.4
	B	V12.4	V13.4	V14.4	V15.4	V16.4	V17.4	V18.4	V19.4	V20.4	V22.M	HS
A5	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V11.P
	B	V12.P	V13.P	V14.P	V15.P	V16.P	V17.P	V18.P	V19.P	V20.P	V23.M	V24.M

1. The layout naming convention is explained on page 18

Details (A5200 RAID5 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout
- 20 5+1 RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 4 Hot Spare drives (can be reduced to 2, other 2 used for misc.)
- 6 Arrays — one RAID5 member per array, RAID1 members mirrored between sequencers
- Direct connect to host (no hubs or daisy chains)
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each Sequencer from data center UPS

Capacity
- 20 5+1 RAID1 logical volumes @ 45.5 GB (9.1GB disks) or 91-GB (18.2-GB disks)
- 4 1+1 RAID1 logical volumes @ 9.1 GB (9.1GB disks) or 18.2-GB (18.2-GB disks)
- **Total data capacity is 946.4 GB (9.1-GB disks)**
- **R-value = 5208**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI Host Adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(1) A5200 - SG-ARY523A- 200G
(5) A5200 - SG-ARY521A- 200GR5 (Adds up to 6 x 200.2 GB in a 72-inch rack)
(132) disk drives (raw capacity is 1201.2 GB with 9.1-GB disks)
(1) A5200 - SG-ARY543A- 400G
(5) A5200 - SG-ARY541A- 400GR5 (Adds up to 6 x 400.4 GB in a 72-inch rack)
(132) disk drives in total (raw capacity is 2400-GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7 or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix W A5200 RAID5 Small/Medium Configuration, Hubs and Sequencers

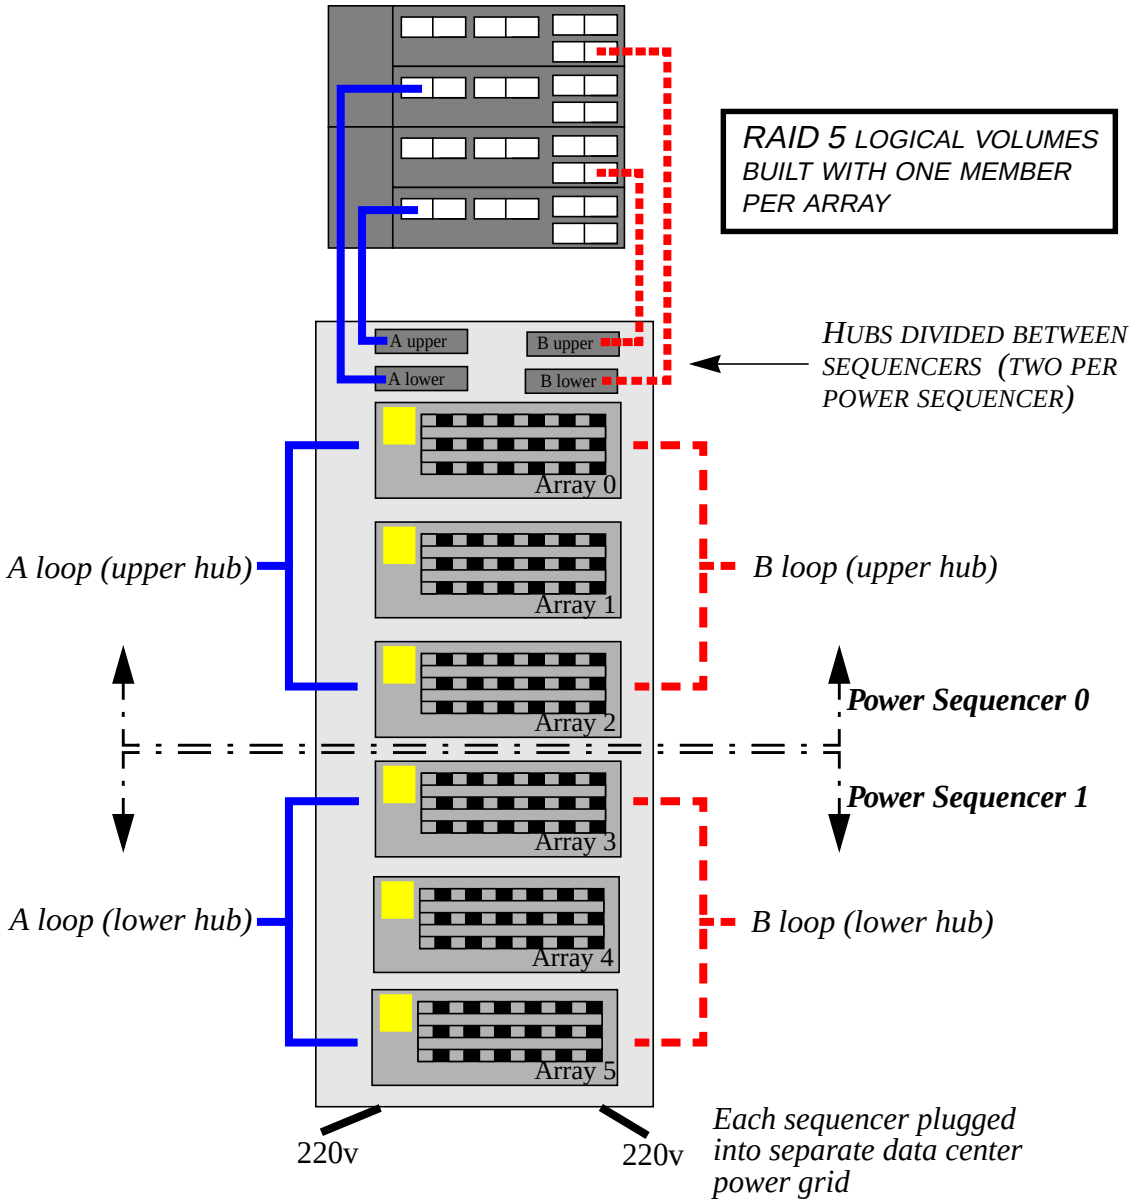

Logical Volume Layout[1] (A5200 RAID5 Small/Medium Configuration, Hubs and Sequencers)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.1
	B	V12.1	V13.1	V14.1	V15.1	V16.1	V17.1	V18.1	V19.1	V20.1	V21.0	HS
A1	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V11.3
	B	V12.3	V13.3	V14.3	V15.3	V16.3	V17.3	V18.3	V19.3	V20.3	V22.0	HS
A2	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5	V8.5	V9.5	V10.5	V11.5
	B	V12.5	V13.5	V14.5	V15.5	V16.5	V17.5	V18.5	V19.5	V20.5	V23.0	V24.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.2
	B	V12.2	V13.2	V14.2	V15.2	V16.2	V17.2	V18.2	V19.2	V20.2	V21.M	HS
A4	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V11.4
	B	V12.4	V13.4	V14.4	V15.4	V16.4	V17.4	V18.4	V19.4	V20.4	V22.M	HS
A5	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V11.P
	B	V12.P	V13.P	V14.P	V15.P	V16.P	V17.P	V18.P	V19.P	V20.P	V23.M	V24.M

1. The layout naming convention is explained on page 18

Details (A5200 RAID5 Small/Medium Configuration, Hubs and Sequencers)

Configuration

RAID Layout
- 20 5+1 RAID5 logical volumes
- 4 1+1 RAID1 logical volumes
- 4 Hot Spare drives (can be reduced to 2, other 2 used for misc.)
- 6 Arrays — one RAID5 member per array, RAID1 members mirrored between sequencers, hot spares available in all sequencers
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each sequencer from data center UPS

Capacity
- 20 5+1 RAID1 logical volumes @ 45.5 GB (9.1-GB disks) or 91 GB (18.2-GB disks)
- 4 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks) or 18.2 GB (18.2-GB disks)
- **Total data capacity is 946.4 GB (9.1-GB disks)**
- **R-value = 5208**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- Host connection through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available.***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(2) Sbus Host Adapter with 1 GBIC - X6730A / (2) PCI Host Adapter with 1 GBIC - X6729A
(2) GBIC - X6731A

Subsystem
(1) A5200 - SG-ARY543A- 2400G (6 x 400.4 GB in a 72-inch rack)
(132) disk drives in total (raw capacity is 2402.4 GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7 or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

Appendix X A5200 RAID5 Medium/Large Configuration, Hubs and Sequencers

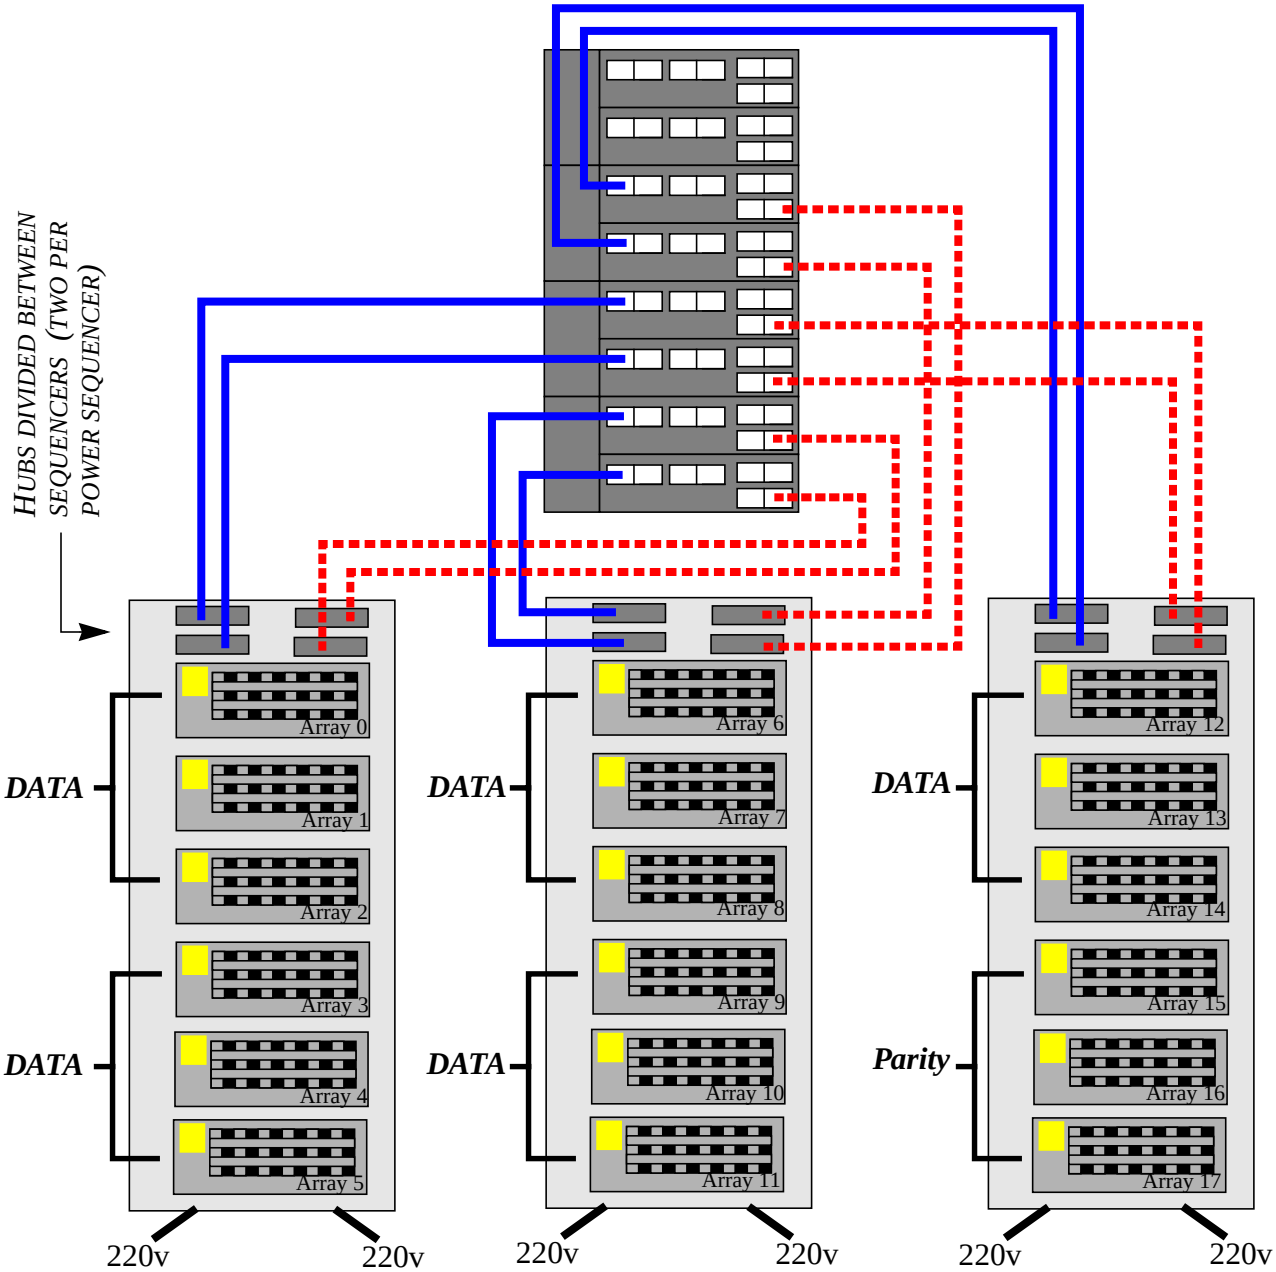

Sequencer 0 and Sequencer 1 of each rack go to a different data center power grid

Logical Volume Layout[1] (A5200 RAID5 Medium/Large Configuration, Hubs and Sequencers)

A0	F	V1.1	V2.1	V3.1	V4.1	V5.1	V6.1	V7.1	V8.1	V9.1	V10.1	V11.1
	B	V12.1	V13.1	V14.1	V15.1	V16.1	V17.1	V18.1	V19.1	V20.1	V61.0	HS
A1	F	V21.1	V22.1	V23.1	V24.1	V25.1	V26.1	V27.1	V28.1	V29.1	V30.1	V31.1
	B	V32.1	V33.1	V34.1	V35.1	V36.1	V37.1	V38.1	V39.1	V40.1	V62.0	HS
A2	F	V41.1	V42.1	V43.1	V44.1	V45.1	V46.1	V47.1	V48.1	V49.1	V50.1	V51.1
	B	V52.1	V53.1	V54.1	V55.1	V56.1	V57.1	V58.1	V59.1	V60.1	V63.0	V64.0
A3	F	V1.2	V2.2	V3.2	V4.2	V5.2	V6.2	V7.2	V8.2	V9.2	V10.2	V11.2
	B	V12.2	V13.2	V14.2	V15.2	V16.2	V17.2	V18.2	V19.2	V20.2	V61.M	HS
A4	F	V21.2	V22.2	V23.2	V24.2	V25.2	V26.2	V27.2	V28.2	V29.2	V30.2	V31.2
	B	V32.2	V33.2	V34.2	V35.2	V36.2	V37.2	V38.2	V39.2	V40.2	V62.M	HS
A5	F	V41.2	V42.2	V43.2	V44.2	V45.2	V46.2	V47.2	V48.2	V49.2	V50.2	V51.2
	B	V52.2	V53.2	V54.2	V55.2	V56.2	V57.2	V58.2	V59.2	V60.2	V63.M	V64.M
A6	F	V1.3	V2.3	V3.3	V4.3	V5.3	V6.3	V7.3	V8.3	V9.3	V10.3	V11.3
	B	V12.3	V13.3	V14.3	V15.3	V16.3	V17.3	V18.3	V19.3	V20.3	V65.0	HS
A7	F	V21.3	V22.3	V23.3	V24.3	V25.3	V26.3	V27.3	V28.3	V29.3	V30.3	V31.3
	B	V32.3	V33.3	V34.3	V35.3	V36.3	V37.3	V38.3	V39.3	V40.3	V66.0	HS
A8	F	V41.3	V42.3	V43.3	V44.3	V45.3	V46.3	V47.3	V48.3	V49.3	V50.3	V51.3
	B	V52.3	V53.3	V54.3	V55.3	V56.3	V57.3	V58.3	V59.3	V60.3	V67.0	V68.0
A9	F	V1.4	V2.4	V3.4	V4.4	V5.4	V6.4	V7.4	V8.4	V9.4	V10.4	V11.4
	B	V12.4	V13.4	V14.4	V15.4	V16.4	V17.4	V18.4	V19.4	V20.4	V65.M	HS
A10	F	V21.4	V22.4	V23.4	V24.4	V25.4	V26.4	V27.4	V28.4	V29.4	V30.4	V31.4
	B	V32.4	V33.4	V34.4	V35.4	V36.4	V37.4	V38.4	V39.4	V40.4	V66.M	HS
A11	F	V41.4	V42.4	V43.4	V44.4	V45.4	V46.4	V47.4	V48.4	V49.4	V50.4	V51.4
	B	V52.4	V53.4	V54.4	V55.4	V56.4	V57.4	V58.4	V59.4	V60.4	V67.M	V68.M
A12	F	V1.5	V2.5	V3.5	V4.5	V5.5	V6.5	V7.5	V8.5	V9.5	V10.5	V11.5
	B	V12.5	V13.5	V14.5	V15.5	V16.5	V17.5	V18.5	V19.5	V20.5	V69.0	HS
A13	F	V21.5	V22.5	V23.5	V24.5	V25.5	V26.5	V27.5	V28.5	V29.5	V30.5	V31.5
	B	V32.5	V33.5	V34.5	V35.5	V36.5	V37.5	V38.5	V39.5	V40.5	V70.0	HS
A14	F	V41.5	V42.5	V43.5	V44.5	V45.5	V46.5	V47.5	V48.5	V49.5	V50.5	V51.5
	B	V52.5	V53.5	V54.5	V55.5	V56.5	V57.5	V58.5	V59.5	V60.5	V71.0	V72.0
A15	F	V1.P	V2.P	V3.P	V4.P	V5.P	V6.P	V7.P	V8.P	V9.P	V10.P	V11.P
	B	V12.P	V13.P	V14.P	V15.P	V16.P	V17.P	V18.P	V19.P	V20.P	V69.M	HS
A16	F	V21.P	V22.P	V23.P	V24.P	V25.P	V26.P	V27.P	V28.P	V29.P	V30.P	V31.P
	B	V32.P	V33.P	V34.P	V35.P	V36.P	V37.P	V38.P	V39.P	V40.P	V70.M	HS
A17	F	V41.P	V42.P	V43.P	V44.P	V45.P	V46.P	V47.P	V48.P	V49.P	V50.P	V51.P
	B	V52.P	V53.P	V54.P	V55.P	V56.P	V57.P	V58.P	V59.P	V60.P	V71.M	V72.M

Details (A5200 RAID5 Medium/Large Configuration, Hubs and Sequencers)

Configuration

RAID Layout
- 60 5+1 RAID5 logical volumes
- 12 1+1 RAID1 logical volumes
- 12 Hot Spare drives (can be reduced to 6, other 6 used for misc.)
- 3 Racks with 6 Arrays each — one RAID5 member per sequencer, RAID1 members mirrored between sequencers, hot spares available in all sequencers
- Use dual loops through Hubs to connect to host
- Separate host adapters (on separate system boards)
- 2 loops per array (connected to separate host adapters)
- Separate power source for each sequencer from data center UPS

Capacity
- 60 5+1 RAID1 logical volumes @ 45.5 GB (9.1-GB disks)
- 12 1+1 RAID1 logical volumes @ 9.1 GB (9.1-GB disks)
- **Total data capacity is 2839.2 GB (9.1-GB disks)**
- **R-value = 15624**

Considerations

Availability
- Dynamic Multi-Pathing (DMP) provides path failover mechanism
- Direct connect to host simplifies fault isolation in case of problems
- Host connection through hubs simplifies implementation
- Host connect through hubs complicates fault isolation
- ***RACK SEQUENCERS AND POWER GRIDS ARE A SINGLE POINT OF FAILURE FOR THIS CONFIGURATION! While the sequencers are extremely reliable and have a very high Mean Time Between Failure, the loss of either power source will result in the UNAVAILABILITY of the RAID5 logical volumes. The RAID1 logical volumes will remain available. To maintain availability through loss of power, six independent power grids are necessary!***

Performance
- 100 MB/s data throughput per loop
- DMP provides higher I/O throughput by load balancing across multiple loops
- More nodes (disks) per loop may impact performance (compared to direct connect configuration)
- Two loops per host adapter may saturate host I/O bus. Use more host adapters (with one loop per adapter) for workloads with a substantial sequential components (e.g., DSS, DW, HPC, data intensive NFSv3)

Components

Hardware
(6) Sbus Host Adapter with 1 GBIC - X6730A / (6) PCI Host Adapter with 1 GBIC - X6729A
(6) GBIC - X6731A

Subsystem
(3) A5200 - SG-ARY543A- 2400G (6 x 400.4 GB in a 72-inch rack)
(396) disk drives in total (raw capacity is 7200 GB with 18.2-GB disks)

Software
Solaris 2.5.1 (8/97), 2.6, 7, or 8
Veritas Volume Manager 2.5, 2.6, or 3.x
Solstice DiskSuite version 4.1 for Solaris 2.5.1 systems
Solstice DiskSuite version 4.2 for Solaris 2.6 systems and later releases

Other
N/A

---

1. The layout naming convention is explained on page 18

Appendix Y T300 RAID1 Small/Medium Configuration, Direct Connect

Hardware Layout Diagram

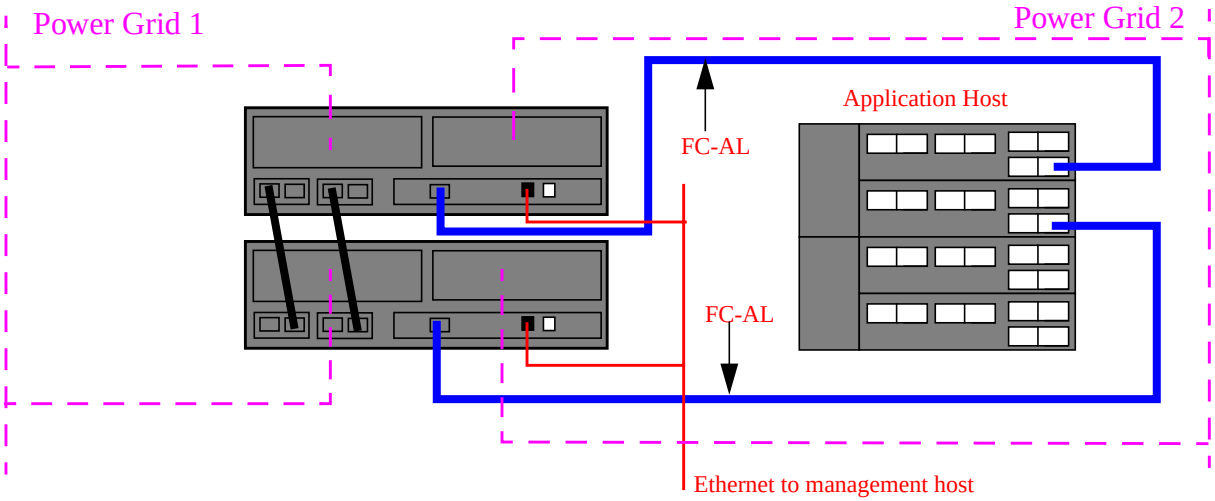

Logical Volume Layout[1] (T300 RAID1 Small/Medium Configuration, Direct Connect)

Without Hot Spares:

Unit 2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	v2.9 u2d9
Unit 1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	v1.9 u1d9

MC = Master Controller
AMC = Alternate Master Controller
MC + AMC = Partner Group

With Hot Spares:

Unit 2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	HS
Unit 1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	HS

1. The layout naming convention is explained on page 18

Details (T300 RAID1 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout

- No Hot Spare: 2, 9 disk RAID logical volumes
- With Hot Spares: 2,8 disk RAID logical volumes
- ***If NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group .***
- ***IF HOT SPARES: One hot spare per box.***
- Single connection between each array and host
- T300 boxes within a Partner Group
- CLI or GUI to interface with T300 RAID controllers
- parameter settings : mp_support=rw, cache=auto, mirror=auto
- Veritas Volume Manager™ (VxVM) and Dynamic MultiPathing (DMP) enabled to assure data path failover

Capacity

- With Hot Spares:
  - 2, 8 RAID1 logical volumes @ 72.8 GB (18.2-GB disks)
  - **Total data capacity is 145.6 GB (18.2-GB disks)**
  - **R-value = 616**
- No Hot Spare:
  - 2, 9 disk RAID1 logical volumes @ 81.9 GB (18.2-GB disks)
  - **Total data capacity is 163.8 GB (18.2-GB disks)**
  - **R-value = 693**

Considerations

Availability

- Dual RAID controllers with automatic failover
- Dual active (one path to each controller) paths to host
- No single point of failure. Even if this configuration is aimed at the Workgroup level it is designed to have complete redundancy.
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for entire array

Performance

- Controller-based RAID
- 2 x 100-MB/s FC-AL to host
- 10,000-RPM drives for high performance
- 256-MB data cache per controller (write cache mirrored between controllers)

Components

Hardware

(2) FC-AL host bus adapters (Sbus Host Adapter with 1 GBIC - X6730A; PCI: X6729A)
(2) 5 meter fiber optic cable (X9715A)

Subsystem

(2) T300 - SG-XARY630A-163G (Tabletop/Deskside with 256-MB data cache and 9x18.2-GB disks)
(1 T310 tabletop tray; T310 chassis contains: 9x18.2GB FC-AL 10k rpm drives, one RAID controller, 2 power cooling units, 2 interface cards. Features shipped with the T310 as part of the shipkit: 1 media interface adapter, 1x5 meter optic cable, 2 power cords (locking power cords in the US), 1 T300 interface cable (used to connect T300s to T300s, proprietary copper FC-AL, 0.5 meter)
(9) 18.2-GB drives total within array

Software

Solaris™ 2.6 or above with required OS patches

Other

Sun StorEdge Volume Manager™ 3.0.4 and later releases

Appendix Z T300 RAID1 Medium/Large Configuration, Direct Connect

Hardware Layout Diagram

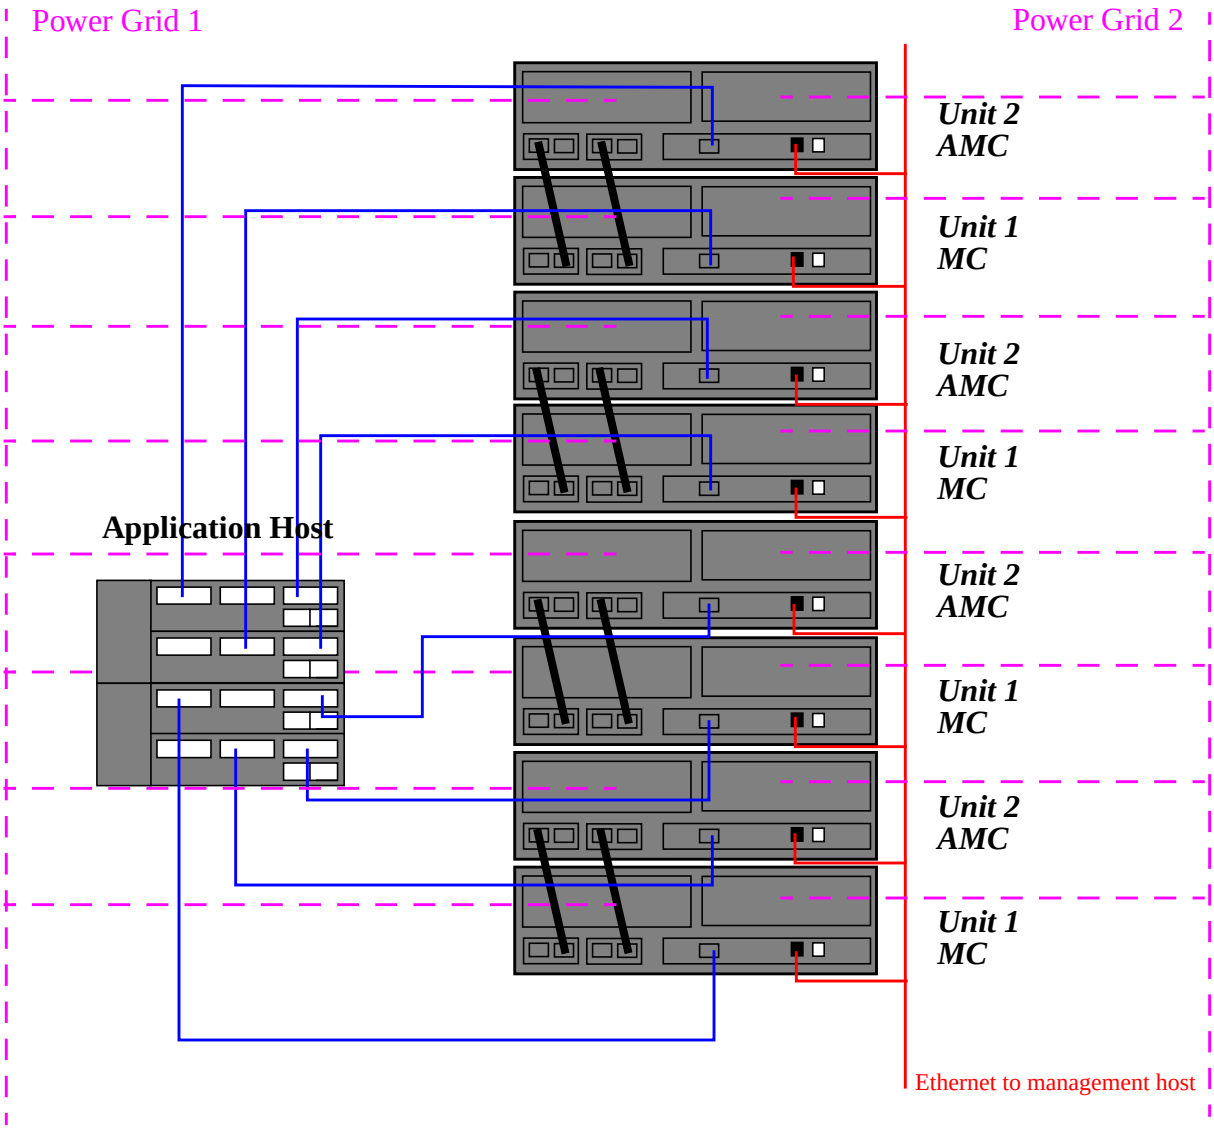

MC = Master Controller
AMC = Alternate Master Controller
MC + AMC = Partner Group

Logical Volume Layout[1] (T300 RAID1 Medium/Large Configuration)

With no Hot Spares

Unit2 AMC	v8.1 u2d1	v8.2 u2d2	v8.3 u2d3	v8.4 u2d4	v8.5 u2d5	v8.6 u2d6	v8.7 u2d7	v8.8 u2d8	v8.9 u2d9
Unit1 MC	v7.1 u1d1	v7.2 u1d2	v7.3 u1d3	v7.4 u1d4	v7.5 u1d5	v7.6 u1d6	v7.7 u1d7	v7.8 u1d8	v7.9 u1d9
Unit2 AMC	v6.1 u2d1	v6.2 u2d2	v6.3 u2d3	v6.4 u2d4	v6.5 u2d5	v6.6 u2d6	v6.7 u2d7	v6.8 u2d8	v6.9 u2d9
Unit1 MC	v5.1 u1d1	v5.2 u1d2	v5.3 u1d3	v5.4 u1d4	v5.5 u1d5	v5.6 u1d6	v5.7 u1d7	v5.8 u1d8	v5.9 u1d9
Unit2 AMC	v4.1 u2d1	v4.2 u2d2	v4.3 u2d3	v4.4 u2d4	v4.5 u2d5	v4.6 u2d6	v4.7 u2d7	v4.8 u2d8	v4.9 u2d9
Unit1 MC	v3.1 u1d1	v3.2 u1d2	v3.3 u1d3	v3.4 u1d4	v3.5 u1d5	v3.6 u1d6	v3.7 u1d7	v3.8 u1d8	v3.9 u1d9
Unit2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	v2.9 u2d9
Unit1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	v1.9 u1d9

MC = Master Controller
AMC = Alternate Master Controller
MC + AMC = Partner Group

With Hot Spares

Unit2 AMC	v8.1 u2d1	v8.2 u2d2	v8.3 u2d3	v8.4 u2d4	v8.5 u2d5	v8.6 u2d6	v8.7 u2d7	v8.8 u2d8	HS
Unit1 MC	v7.1 u1d1	v7.2 u1d2	v7.3 u1d3	v7.4 u1d4	v7.5 u1d5	v7.6 u1d6	v7.7 u1d7	v7.8 u1d8	HS
Unit2 AMC	v6.1 u2d1	v6.2 u2d2	v6.3 u2d3	v6.4 u2d4	v6.5 u2d5	v6.6 u2d6	v6.7 u2d7	v6.8 u2d8	HS
Unit1 MC	v5.1 u1d1	v5.2 u1d2	v5.3 u1d3	v5.4 u1d4	v5.5 u1d5	v5.6 u1d6	v5.7 u1d7	v5.8 u1d8	HS
Unit2 AMC	v4.1 u2d1	v4.2 u2d2	v4.3 u2d3	v4.4 u2d4	v4.5 u2d5	v4.6 u2d6	v4.7 u2d7	v4.8 u2d8	HS
Unit1 MC	v3.1 u1d1	v3.2 u1d2	v3.3 u1d3	v3.4 u1d4	v3.5 u1d5	v3.6 u1d6	v3.7 u1d7	v3.8 u1d8	HS
Unit2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	HS
Unit1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	HS

1. The layout naming convention is explained on page 18

Details (T300 RAID1 Medium/Large Configuration)

Configuration

RAID Layout

- No Hot Spares: 8, 9 disk RAID1 logical volumes
- With Hot Spares: 16, 2+2 RAID1 logical volumes
- ***If NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group .***
- ***IF HOT SPARES: Only One hot spare per box.***
- Single connection between each array and host
- arrays are organized in Partner Groups (2 arrays per Partner Group)
- CLI to interface with T300 RAID controllers
- parameter settings : mp_support=rw, cache=auto, mirror=auto
- VERITAS Volume Manager (VxVM) to build additional logical layers on top of hardware RAID volumes, and Dynamic MultiPathing (DMP) enabled for data path failover

Capacity

- With Hot Spares:
  - 8, 8 disk RAID1 logical volumes @ 72.8 GB (18.2-GB disks)
  - **Total data capacity is 582.4 GB (18.2-GB disks)**
  - **R-value = 2464**
- No Hot Spares:
  - 8, 9 disk RAID1 logical volumes @ 81.9 GB (18.2-GB disks)
  - **Total data capacity is 655.20 GB (18.2-GB disks)**
  - **R-value = 2772**

Considerations

Availability

- Dual RAID controllers with automatic failover
- Dual path to host
- No single point of failure. Even if this configuration is aimed at the Workgroup level it is designed to have complete redundancy.
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- 100-MB/s FC-AL to host
- 10,000-RPM drives for high performance
- 256-MB data cache per controller

Components

Hardware

(2) FC-AL host bus adapters (Sbus Host Adapter with 1 GBIC - X6730A; PCI: X6729A)
(2) 5 meter fiber optic cable (X9715A)

Subsystem

(8) T300 - SG-XARY630A-163G (Tabletop/Deskside with 256-MB data cache and 9x18.2-GB disks)
(1 T310 tabletop tray; T310 chassis contains: 9x18.2GB FC-AL 10k rpm drives, one RAID controller, 2 power cooling units, 2 interface cards. Features shipped with the T310 as part of the shipkit: 1 media interface adapter, 1x5 meter optic cable, 2 power cords (locking power cords in the US), 1 T300 interface cable (used to connect T300s to T300s, proprietary copper FC-AL, 0.5 meter)
(9) 18.2-GB drives total within array

Software

Solaris™ 2.6 or above with required OS patches

Other

Sun StorEdge Volume Manager™ 3.0.4 and later releases

Appendix AA T300 RAID5 Small/Medium Configuration, Direct Connect

Hardware Layout Diagram

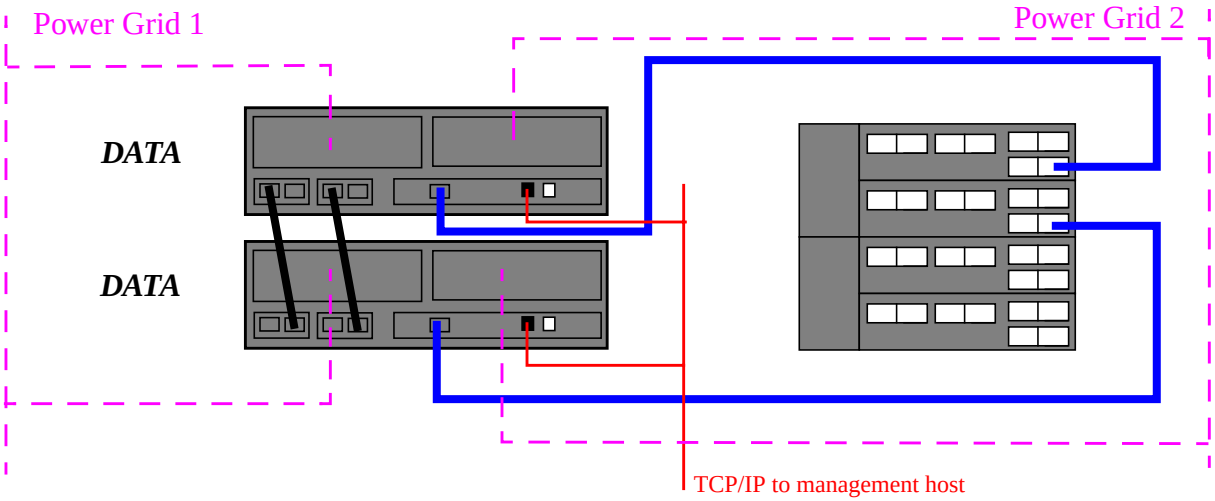

Logical Volume Layout[1] (T300 RAID5 Small/Medium Configuration, Direct Connect)

Without Hot Spares:

Unit 2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	v2.9 u2d9
Unit 1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	v1.9 u1d9

MC = Master Controller
AMC = Alternate Master Controller
MC + AMC = Partner Group

With Hot Spares:

Unit 2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	HS
Unit 1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	HS

1. The layout naming convention is explained on page 18

Details (T300 RAID5 Small/Medium Configuration, Direct Connect)

Configuration

RAID Layout

- No hot spares: 2, 9 disk RAID5 logical volumes
- With hot spares: 4, 3+1 RAID5 logical volumes
- ***If NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- ***If HOT SPARES: Only One hot spare per box.***
- Single connection between each array and host
- T300 boxes within a Partner Group
- CLI to interface with T300 RAID controllers
- parameter settings : mp_support=rw, cache=auto, mirror=auto
- VERITAS Volume Manager™ (VxVM) and Dynamic MultiPathing (DMP) enabled to assure data path failover

Capacity

- With Hot Spares:
  - 2, 7+1 RAID5 logical volumes 127.4 GB (18.2-GB disks)
  - **Total data capacity is 254.8 GB (18.2-GB disks)**
  - **R-value = 1078**
- No Hot Spares:
  - 2, 8+1 disk RAID5 logical volumes @ 145.6 GB (18.2-GB disks)
  - **Total data capacity is 291.2 GB (18.2-GB disks)**
  - **R-value = 1232**

Considerations

Availability

- Dual RAID controllers with automatic failover
- Dual path to host
- No single point of failure. Even if this configuration is aimed at the Workgroup level it is designed to have complete redundancy.
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- 100-MB/s FC-AL to host
- 10,000-RPM drives for high performance
- 256-MB data cache per controller

Components

Hardware

(2) FC-AL host bus adapters (Sbus Host Adapter with 1 GBIC - X6730A; PCI: X6729A)
(2) 5 meter fiber optic cable (X9715A)

Subsystem

(2) T300 - SG-XARY630A-163G (Tabletop/Deskside with 256-MB data cache and 9x18.2-GB disks)
(1 T310 tabletop tray; T310 chassis contains: 9x18.2GB FC-AL 10k rpm drives, one RAID controller, 2 power cooling units, 2 interface cards. Features shipped with the T310 as part of the shipkit: 1 media interface adapter, 1x5 meter optic cable, 2 power cords (locking power cords in the US), 1 T300 interface cable (used to connect T300s to T300s, proprietary copper FC-AL, 0.5 meter)
(9) 18.2-GB drives total within array

Software

Solaris™ 2.6 or above with required OS patches

Other

Sun StorEdge Volume Manager™ 3.0.4 and later releases

Appendix AB T300 RAID5 Medium/Large Configuration, Direct Connect

Hardware Layout Diagram

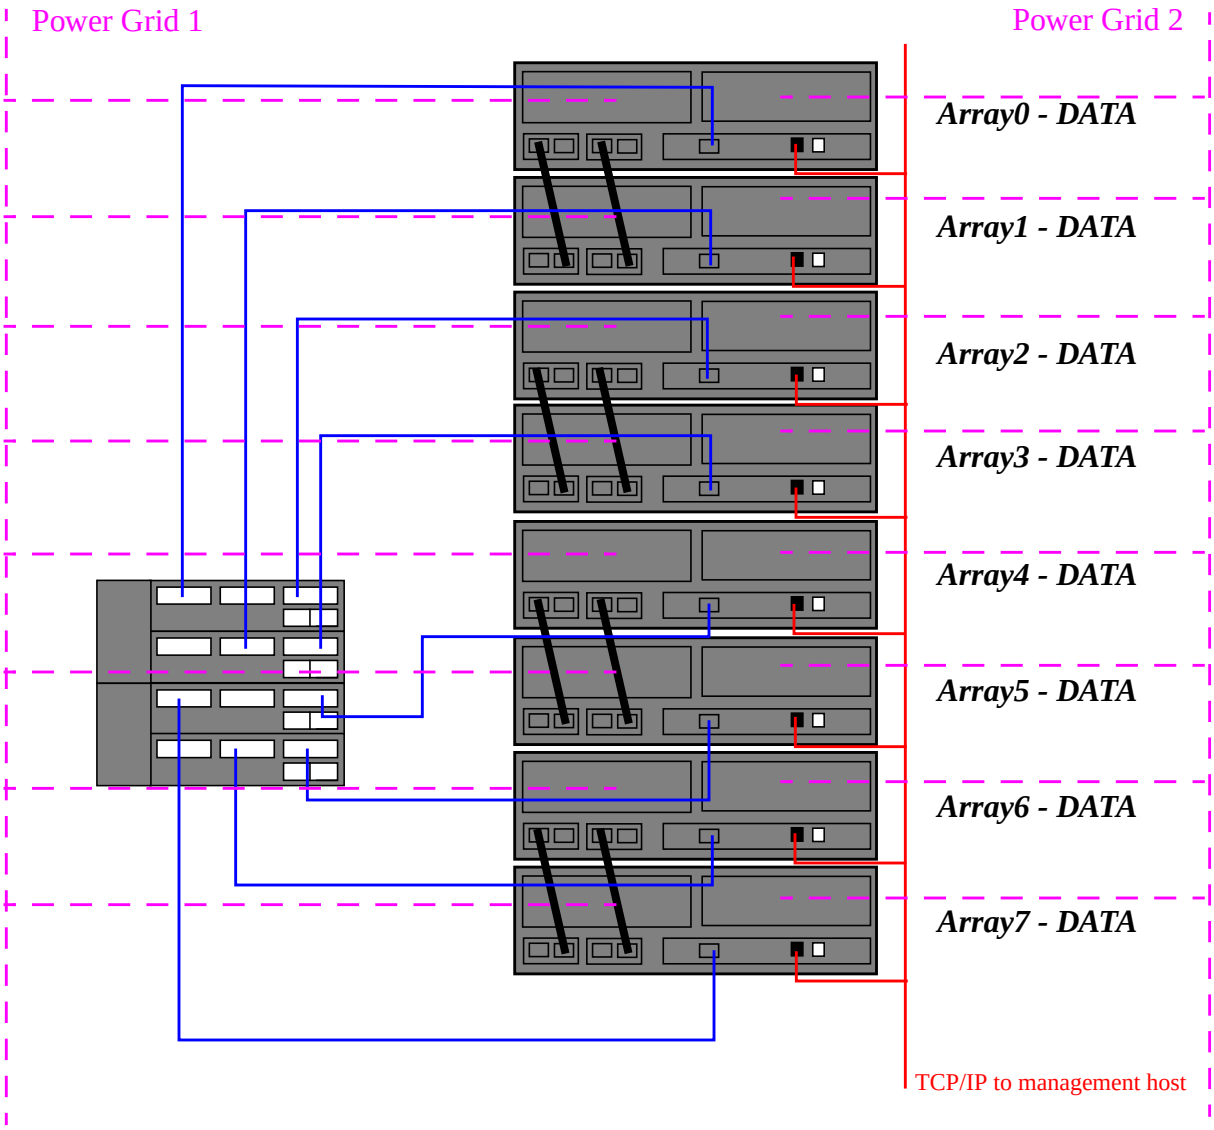

Logical Volume Layout[1] (T300 RAID5 Medium/Large Configuration)

Without Hot Spares

Unit2 AMC	v8.1 u2d1	v8.2 u2d2	v8.3 u2d3	v8.4 u2d4	v8.5 u2d5	v8.6 u2d6	v8.7 u2d7	v8.8 u2d8	v8.9 u2d9
Unit1 MC	v7.1 u1d1	v7.2 u1d2	v7.3 u1d3	v7.4 u1d4	v7.5 u1d5	v7.6 u1d6	v7.7 u1d7	v7.8 u1d8	v7.9 u1d9
Unit2 AMC	v6.1 u2d1	v6.2 u2d2	v6.3 u2d3	v6.4 u2d4	v6.5 u2d5	v6.6 u2d6	v6.7 u2d7	v6.8 u2d8	v6.9 u2d9
Unit1 MC	v5.1 u1d1	v5.2 u1d2	v5.3 u1d3	v5.4 u1d4	v5.5 u1d5	v5.6 u1d6	v5.7 u1d7	v5.8 u1d8	v5.9 u1d9
Unit2 AMC	v4.1 u2d1	v4.2 u2d2	v4.3 u2d3	v4.4 u2d4	v4.5 u2d5	v4.6 u2d6	v4.7 u2d7	v4.8 u2d8	v4.9 u2d9
Unit1 MC	v3.1 u1d1	v3.2 u1d2	v3.3 u1d3	v3.4 u1d4	v3.5 u1d5	v3.6 u1d6	v3.7 u1d7	v3.8 u1d8	v3.9 u1d9
Unit2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	v2.9 u2d9
Unit1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	v1.9 u1d9

MC = Master Controller
AMC = Alternate Master Controller
MC + AMC = Partner Group

With Hot Spares

Unit2 AMC	v8.1 u2d1	v8.2 u2d2	v8.3 u2d3	v8.4 u2d4	v8.5 u2d5	v8.6 u2d6	v8.7 u2d7	v8.8 u2d8	HS
Unit1 MC	v7.1 u1d1	v7.2 u1d2	v7.3 u1d3	v7.4 u1d4	v7.5 u1d5	v7.6 u1d6	v7.7 u1d7	v7.8 u1d8	HS
Unit2 AMC	v6.1 u2d1	v6.2 u2d2	v6.3 u2d3	v6.4 u2d4	v6.5 u2d5	v6.6 u2d6	v6.7 u2d7	v6.8 u2d8	HS
Unit1 MC	v5.1 u1d1	v5.2 u1d2	v5.3 u1d3	v5.4 u1d4	v5.5 u1d5	v5.6 u1d6	v5.7 u1d7	v5.8 u1d8	HS
Unit2 AMC	v4.1 u2d1	v4.2 u2d2	v4.3 u2d3	v4.4 u2d4	v4.5 u2d5	v4.6 u2d6	v4.7 u2d7	v4.8 u2d8	HS
Unit1 MC	v3.1 u1d1	v3.2 u1d2	v3.3 u1d3	v3.4 u1d4	v3.5 u1d5	v3.6 u1d6	v3.7 u1d7	v3.8 u1d8	HS
Unit2 AMC	v2.1 u2d1	v2.2 u2d2	v2.3 u2d3	v2.4 u2d4	v2.5 u2d5	v2.6 u2d6	v2.7 u2d7	v2.8 u2d8	HS
Unit1 MC	v1.1 u1d1	v1.2 u1d2	v1.3 u1d3	v1.4 u1d4	v1.5 u1d5	v1.6 u1d6	v1.7 u1d7	v1.8 u1d8	HS

1. The layout naming convention is explained on page 18

Details (T300 RAID5 Medium/Large Configuration)

Configuration

RAID Layout

- No Hot Spares: 8, 9 disk RAID5 logical volumes
- With Hot Spares: 16, 3+1 RAID5 logical volumes
- ***If NO HOT SPARES: System administrators need to replace failed drives as soon as possible to minimize risk of double failures in a disk group.***
- ***IF HOT SPARES: Only One hot spare per box.***
- Single connection between each array and host
- arrays are organized in Partner Groups (2 arrays per Partner Group)
- CLI to interface with T300 RAID controllers
- parameter settings : mp_support=rw, cache=auto, mirror=auto
- VERITAS Volume Manager (VxVM) to build additional logical layers on top of hardware RAID volumes, and Dynamic MultiPathing (DMP) enabled for data path failover

Capacity

- With hot spares:
  - 8, 7+1 RAID5 logical volumes @ 127.4 GB (18.2-GB disks)
  - **Total data capacity is 1019.2 GB (18.2-GB disks)**
  - **R-value = 4312**
- No hot spares:
  - 8, 8+1 disk RAID5 logical volumes @ 145.6 GB (18.2-GB disks)
  - **Total data capacity is 1164.80 GB (18.2-GB disks)**
  - **R-value = 4928**

Considerations

Availability

- Dual RAID controllers with automatic failover
- Dual path to host
- No single point of failure. Even if this configuration is aimed at the Workgroup level it is designed to have complete redundancy.
- Dual hot-plug power supplies and cooling units in the disk arrays
- Hot-swappable drives
- Battery backup for data cache

Performance

- Controller-based RAID
- 100-MB/s FC-AL to host
- 10,000-RPM drives for high performance
- 256-MB data cache per controller

Components

Hardware

(2) FC-AL host bus adapters (Sbus Host Adapter with 1 GBIC - X6730A; PCI: X6729A)
(2) 5 meter fiber optic cable (X9715A)

Subsystem

(8) T300 - SG-XARY630A-163G (Tabletop/Deskside with 256-MB data cache and 9x18.2-GB disks)
(1 T310 tabletop tray; T310 chassis contains: 9x18.2GB FC-AL 10k rpm drives, one RAID controller, 2 power cooling units, 2 interface cards. Features shipped with the T310 as part of the shipkit: 1 media interface adapter, 1x5 meter optic cable, 2 power cords (locking power cords in the US), 1 T300 interface cable (used to connect T300s to T300s, proprietary copper FC-AL, 0.5 meter)
(9) 18.2-GB drives total within array

Software

Solaris™ 2.6 or above with required OS patches

Other

Sun StorEdge Volume Manager™ 3.0.4 and later releases