



USING RAID 6 ON THE SUN™ STORAGETEK™ 6140 ARRAY INCREASING RELIABILITY FOR STORAGE ARRAYS

Said A. Syed, NPI and OEM Array Engineering Group

Sun BluePrints™ Online

Part No 820-5825-10
Revision 1.0, 8/6/08

Table of Contents

Sun's RAID 6 Solution: Reliability, Performance, Economy	1
Disk Failure Rates, Reliability Needs Promote RAID 6 Usage	3
RAID 6 Implementation	4
Best Practices for Implementing RAID 6	8
Conclusions	11
About the Author	11
Acknowledgements	11
References	12
Ordering Sun Documents	12
Accessing Sun Documentation Online	12

Using RAID 6 on the Sun™ StorageTek™ 6140 Array: Increasing Reliability for Storage Arrays

Many business applications requiring high availability, reliability, and performance have historically chosen either RAID 5 or RAID 1+0¹ technology for their data storage needs. While both of these RAID algorithms provide increased data availability, they are not without certain disadvantages. For example, disk groups using RAID 5 are only able to sustain a single drive failure. And disk groups using RAID 1+0 algorithms must use 50 percent of the available storage capacity for mirroring. These disadvantages are exacerbated as disk drive capacity increases. With larger disks, the RAID 1+0 implementation becomes increasingly costly. And as number and size of disks increase, the risk of a disk error also increases making RAID 5 disk groups inherently less reliable.

RAID 6 is cheaper to implement than RAID 1+0 and is more reliable than RAID 5.

RAID 6, a RAID algorithm that is recently gaining popularity, helps solve these problems. By using dual parity calculations, a RAID 6 disk group can tolerate two concurrent disk failures. Sun™ StorageTek™ 6140 arrays using RAID 6 protection are an ideal solution to meet the needs of applications requiring highly available, high-performance data storage.

This article addresses the following topics:

- “Sun’s RAID 6 Solution: Reliability, Performance, Economy” on page 1 describes the Sun RAID 6 solution using the StorageTek 6140 array.
- “RAID 6 Implementation” on page 4 provides information on how RAID 6 is implemented and describes how it increases data reliability by protecting against multiple failures.
- “Best Practices for Implementing RAID 6” on page 8 discusses best practices when implementing RAID 6 on the StorageTek 6140 array.

Sun’s RAID 6 Solution: Reliability, Performance, Economy

Sun StorageTek 6140 arrays with RAID 6 technology feature enterprise class data protection for the cost conscious. This solution can provide enhanced reliability for applications with large data sets requiring both high performance and high availability.

Sun StorageTek 6140 Array

The Sun StorageTek 6140 Array provides performance, high availability, and reliability in an economical package. The array includes dual intelligent RAID controllers, with a dedicated processor to perform all RAID calculations for enhanced performance. The redundant design and non-disruptive upgrade procedures provide exceptional data

Intelligent RAID controllers and a dedicated processor for performing RAID calculations provide enhanced performance in the StorageTek 6140 array.

1. RAID 1+0, mirroring and striping, is also referred to as RAID 10 in some documentation.

protection for business-critical applications. All components in the array's data path are redundant; if one component fails, the StorageTek 6140 array automatically fails-over to the alternate component providing continuous uptime and uninterrupted data availability. Hot-swappable components enable maintenance to occur without system disruption. Support for up to 112 SATA and Fibre Channel (FC) drives helps meet growing data storage needs.



Figure 1. The Sun StorageTek 6140 Array.

Key features of the Sun StorageTek 6140 Array include:

- Dual Fibre Channel RAID controllers support RAID levels 0, 1, 3, 5, 1+0, and 6
- Scales to 112 disk drives and 112 TB of data in a small footprint
- Supports high-performance FC drives and high-performance SATA drives
- Redundant, hot-swappable components protect against any single point of failure

RAID 6 Overview

RAID 6 was developed to fulfill a specific need: provide data availability for applications requiring extremely high fault tolerance at a reasonable cost. RAID 6 is highly fault tolerant, having the ability to sustain two consecutive drive failures within the same RAID group. Furthermore, RAID 6 provides this increased data availability at a fraction of the cost of a RAID-1+0 disk group, which uses 50 percent of the storage subsystem's capacity for mirroring.

Some potential uses for RAID 6 include the following:

- *Email Servers*

Email usage has grown exponentially — there were an estimated 1.4 billion email accounts in 2007, and the number is expected to double in the next few years.

Email servers such as Sun Java™ Communications Suite and Microsoft Exchange typically require high storage availability. With an estimated 516 million corporate email users depending on them for reliable communications, email servers are ideal candidates for RAID 6 implementation.

Potential uses for RAID 6 include applications requiring extremely high fault tolerant storage, including:

- Email servers
- Web servers
- File servers
- Database servers
- Financial applications

- *Web Servers*

The Internet is a crucial source of information flow, and users depend on it for diverse needs including distributed business applications, instant messaging, online auctions, gaming, news, and video on-demand. In many environments, Web servers and the data they share need to be available to users at all times. The requirement for highly available data mandates highly fault tolerant storage. RAID 6, with its built-in availability and redundancy and low cost, is a good solution for Web server implementations.

- *File Servers*

File servers are an essential and integral part of virtually all organizations. For many businesses, user home directories need to be available at all times for all users. The need for highly available, always online data and storage at a reasonable cost makes file servers a perfect candidate for RAID 6 implementation.

- *Other Uses*

Other candidates for RAID 6 include databases, such as MySQL™ and Oracle database servers; financial applications such as SAP; and many other similar applications that require high availability of data without sacrificing performance and while keeping cost at a minimum.

The need for RAID 6 comes as:

- Drive capacity continues to increase, with 2 TB drives planned for 2009.
- Annual Replacement Rate for disk drives can be 2 - 4% on average, and as high as 13% in certain circumstances.
- Unrecoverable Read Error (URE) rates for all drive types, and especially SATA drives, are within range of expected URE every 12 TBs of data read.
- The likelihood of a second URE occurring during single drive reconstruction is all but guaranteed for RAID groups using very large capacity disk drives.

Disk Failure Rates, Reliability Needs Promote RAID 6 Usage

While disks are generally considered to be reliable, they do fail — and can fail more often than might be anticipated. A scientific study performed by the Computer Science Department¹ at Carnegie Mellon University found that the *Annual Replacement Rate* (ARR) of 100,000 disk drives was significantly higher than the *Mean Time to Failure* (MTTF) for the same disk drives. The MTTF for these drives as listed in the vendor datasheet ranged from 1,000,000 to 1,500,000 hours, which suggested an annual failure rate of 0.88%. However, the study found that the ARR was typically above 1 - 4 percent on average, and as much as 13 percent in certain instances within the first three years of the disk drive's useful life. The study also found that the disk replacement rate increased constantly as the drives aged. An additional key finding of the study was that there was no significant difference in replacement rates between SCSI, Fibre Channel and SATA drives.

In addition to disk drive failure, drives can experience *Unrecoverable Read Errors* (UREs) when reading a disk block. The URE rate is defined as an unrecoverable read error per bits read. SATA drives in particular have a much lower URE per bit read than Fibre Channel and SAS, as illustrated in Table 1, and therefore are more likely to experience unrecoverable read errors during operation.

1. Schroeder, Bianca and Gibson, Garth A. "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?" *FAST '07: 5th USENIX Conference on File and USENIX Association Storage Technologies*. <http://www.usenix.org/events/fast07/tech/schroeder/schroeder.pdf>

Table 1. URE per bit read for Fibre Channel, SAS, and SATA drives.

Drive Speed and Type	URE per Bit Read
2 and 4 Gb/sec Fibre Channel	$10^{15} - 10^{16}$
3 Gb/sec SAS	$10^{15} - 10^{16}$
3 Gb/sec SATA	$10^{14} - 10^{15}$
1.5 and 3 Gb/sec SATA	10^{14}

A URE of 10^{14} means a SATA drive will experience an unrecoverable bit error on average once every one hundred trillion bits, which is approximately every 12 Terabytes. Consider a RAID 5 disk group with eight 2 TB SATA drives (7 data drives and 1 parity drive) and usable raw space of approximately 13 TB. If the RAID group is working on regenerating a failed drive, the likelihood of another URE is extremely high, if not almost guaranteed. This phenomenon, called *Media Errors During Reconstruction* (MEDR), can result in unrecoverable RAID group. The risk for MEDRs increases with the size of the RAID group; increasing disk capacities and large RAID groups increase the risk of encountering an unrecoverable disk block during reconstruction.

RAID 6 is able to sustain a second URE by using two independent parity algorithms, P+Q.

RAID 6 is able to sustain a second URE or MEDR. Unlike RAID 5, a RAID 6 disk group will continue to service read/write requests during regeneration with not one but two consecutive drive failures. This ability gives RAID 6 a significant advantage over RAID 5, particularly with higher capacity drives.

RAID 6 Implementation

RAID 6 builds on the commonly used RAID 5 single parity calculation. RAID 5 uses a striped data set with distributed parity, as shown in Figure 2. RAID 6 utilizes two independent parity schemes striped across multiple spindles for increased fault tolerance. This approach protects against data loss caused by two concurrent disk drive failures, and thereby provides a higher level of data protection than RAID 5.

RAID 6 P+Q uses RAID 5's existing P parity algorithm and a second independent Reed-Solomon Q parity algorithm.

Sun StorageTek arrays make use of the P+Q algorithm to compute the two independent parity calculations. The RAID 6 P+Q algorithm uses the existing RAID 5 single parity generation algorithm, P parity, which is striped across all the disks within the RAID group. A second completely independent parity, Q parity, is computed based on the Reed-Solomon Error Correcting Code. The Q parity is also striped across all the disks within the RAID group. Figure 2 illustrates the striping of data blocks and both independent RAID 6 P and Q parity blocks within a RAID group.

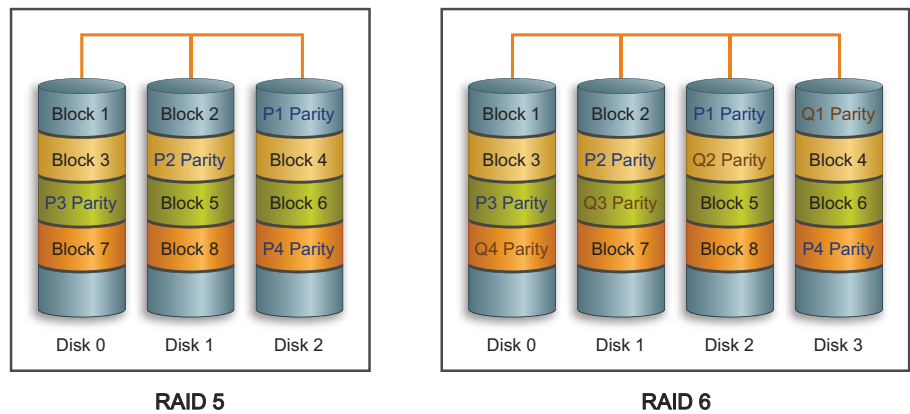


Figure 2. A RAID 6 stripe set is identical to RAID 5, with the addition of Q parity computed using Reed-Solomon Error-Correcting Code.

Data Regeneration and Availability

A RAID 5 disk group will go into degraded status after suffering a bad block (URE) or failed disk. When a disk failure occurs, data can be recalculated using the distributed parity and the remaining drives. While slower, I/O can continue without interruption to the user while the data on the failed drive is recovered. However, the data is vulnerable until the recovery completes; if there is a second failure, data loss can occur. In contrast, RAID 6 provides for a more reliable alternative than RAID 5 as it is able to sustain two complete disk failures.

RAID 6 is able to service I/O while in degraded mode and during reconstruction.

Multiple block error protection is provided during optimal state, and single block error protection is provided during degraded mode.

When a disk failure in a RAID 6 P+Q RAID group occurs, blocks are regenerated using the ordinary XOR parity (P). At the same time, both P and Q parities are recomputed for each regenerated block. If a second disk failure occurs while the first failed disk is either being recovered or still in degraded status, blocks on the second failed disk are regenerated using the second independent Q parity. With RAID 6, data availability is not affected should two drives fail concurrently: the P+Q algorithm will simply recompute the data blocks upon a read request and continue to service any reads and writes as such while data regeneration is in progress for both failed drives.

The potential for experiencing a URE or hitting a bad block increases with both the number and the size of disk drives. Also, as disk capacity increases, so does the time required to rebuild a failed drive — and it is during this rebuild time that a RAID 5 disk group is vulnerable to data loss if a second drive were to fail. Thus, the need for RAID 6 becomes more apparent when availability and large disk capacity are both important factors.

RAID 6 uses the capacity of two disk drives within the RAID group for the two independent parity algorithms.

Capacity

Due to the second independent Q parity computation and striping across all disks in the RAID group, the RAID 6 P+Q algorithm requires capacity equal to two complete disks.

Usable capacity of a RAID 6 disk group can be calculated using the formula:

$$(N - 2) \times S_{\min}$$

Where N is the total number of disks available and S_{\min} is the capacity of the smallest drive in the RAID group.

A minimum of four disk drives are required to implement a RAID 6 disk group. Given the formula above, a four disk RAID 6 disk group provides space for two data and two parity disk drives. RAID 6 P+Q parity is implemented using an addition in the Galois¹ field, which is set at GF(28). Therefore, a theoretical maximum of 257 drives can be part of any one RAID 6 disk group, 255 of which can be data drives.

Performance

RAID 6 can sustain write performance similar to that of RAID 5 when intelligent storage array controllers with dedicated processors for RAID calculations are used.

If implemented using an intelligent storage array controller with a dedicated processor for RAID computations, the RAID 6 P+Q algorithm may sustain write performance almost similar to that of RAID 5. Sun StorageTek 6140 arrays include a dedicated processor on the controller for RAID computations, for improved performance.

Performance is affected during regeneration, especially if two drive failures have occurred.

However, there may be read and write performance degradation during drive regeneration. If the RAID 6 disk group is currently in the process of regenerating a failed drive and a second drive failure does occur, the controller will need to service any reads by regenerating requested data using both P and Q parity blocks on the surviving disks. This generates a substantial overhead, possibly as much as 20 percent, for the RAID 6 disk group as compared to RAID 5. This performance degradation during degraded mode should be expected and is not significant when compared to the alternative, which is complete data loss.

A RAID 6 disk group also takes more time to rebuild after a drive failure as compared to RAID 5. The longer rebuild time is attributed to the regeneration of data using two parity algorithms instead of one.

RAID 6 Benefits and Trade-offs

Advantages of RAID 6:

- Protection against two drive failures
- Performance almost similar to RAID 5, with hardware assisted RAID computations
- Lower cost than RAID 1+0

The most substantial advantage of using RAID 6 is that it provides protection against multiple unrecoverable errors in the optimal state (normal operation with no disk failures) and against a single unrecoverable error in the degraded state. This is a huge advantage over RAID 5 which only protects against a single unrecoverable error during optimal state. Other RAID technologies offer similar or slightly better performance and

1. Anvin, H. Peter. "The Mathematics of RAID-6," online paper, July 2008.
<http://kernel.org/pub/linux/kernel/people/hpa/raid6.pdf>

reliability, however, at a very high cost. RAID 1+0 provides similar availability and redundancy as RAID 6 by mirroring the data across disk drives and then striping the mirrored data across a separate set of disks drives. However, RAID 1+0 uses half of available storage for mirroring making it a very expensive solution. The same highly available storage space for applications can be provided using RAID 6 at a fraction of the cost.

RAID 6 Trade-offs:

- Slightly reduced performance compared to RAID 5 during reconstruction
- Longer reconstruction time
- Slightly more expensive than RAID 5, due to additional use of capacity for two parity calculations

The most significant trade-off with using RAID 6 is the potential for slightly reduced write performance as compared to RAID 5 due to two independent parity calculations. However, use of intelligent RAID controllers with dedicated processors for RAID computations, like those in the StorageTek 6140 arrays, can significantly increase RAID 6 write performance to levels comparable to RAID 5. RAID 6 disk groups can also experience longer rebuild times after a drive failure, as two parity algorithms are required to regenerate the failed drive, as compared to the single RAID 5 parity calculation. RAID 6 also requires additional disk space as compared to RAID 5, as two drives are required for RAID 6 P+Q parity rather than the single drive for RAID 5 parity.

In summary, RAID 6 offers increased fault tolerance with its ability to continue operation after two failed disks. While RAID 6 can experience slightly reduced write performance and longer rebuild times after a failed disk, using dedicated RAID processors can mitigate the performance penalty. In comparison to other popular RAID algorithms, RAID 6 offers significantly improved fault tolerance when compared to RAID 5, and can provide similar availability to RAID 1+0 without requiring double the amount of storage capacity.

Sun StorageTek 6140 arrays require firmware release 7.10.x.x to support RAID 6. Existing StorageTek 6140 arrays can be upgraded; contact local Sun Services representatives for information on the upgrade process.

Best Practices for Implementing RAID 6

The latest firmware upgrade, revision 7.10.x.x, enables the Sun StorageTek 6140 array to make use of the dedicated processor for RAID P+Q parity computation and allows the creation of RAID 6 volumes. All new StorageTek 6140 arrays come pre-loaded with this latest firmware. Existing StorageTek 6140 arrays can be upgraded to the new 7.10.x.x firmware level; local Sun Services representatives can provide further information on the upgrade process.

Note – The 7.10.x.x firmware release is also available for Sun StorageTek 6540 storage arrays. However, the RAID controllers for Sun StorageTek 6540 arrays do not have the capability to perform P+Q calculations and hence cannot use the RAID 6 enhancement within the new firmware. All other enhancements are available on the StorageTek 6540 arrays, however.

With proper planning for expansion trays and disk drive considerations, a RAID 6 virtual disk in a StorageTek 6140 array can sustain up to two complete expansion tray failures, in addition to being able to sustain two consecutive disk drive failures.

Striping Across Storage Trays

Creating RAID groups (or *virtual disks*) that span multiple storage trays in the StorageTek 6140 array is a recommended best practice for enhanced availability of data. Although RAID 6 provides for higher availability than RAID 5 and can sustain two independent disk failures, the failure of a single storage array tray can cause unnecessary downtime and loss of access to data for some virtual disk configurations. This potential downtime caused by a tray failure can be avoided by creating RAID groups and virtual disks that span multiple trays. With proper planning and the right combination of StorageTek 6140 expansion trays and disk considerations, a RAID 6 virtual disk can sustain two complete expansion tray failures.

Figure 3 illustrates three example eight-spindle RAID 6 virtual disks in a Sun StorageTek 6140 storage array. Virtual Disk A is striped across two trays in a StorageTek 6140 array; Virtual Disk B is striped across 4 trays; and Virtual Disk C is striped across all seven trays in the array.

Because Virtual Disk A uses RAID 6, it provides protection against the failure of any two disk drives in the RAID group. However, Virtual Disk A will not be able to sustain a tray failure for either of the trays it resides on, as each tray contains four drives used in the virtual disk.

In contrast, Virtual Disk B will be able to continue servicing I/O during a complete tray failure. As with Virtual Disk A, Virtual Disk B provides protection against any two consecutive disk failures within the RAID group. However, by using only two disks from each tray, Virtual Disk B is also able to sustain any one complete tray failure for any of the four expansion trays it is striped across.




Figure 3. Three 8-spindle RAID 6 virtual disks spanning multiple trays in a Sun StorageTek 6140 array.

Virtual Disk C, striped across all seven trays in the StorageTek 6140 array, is the most resilient of the three example RAID groups. On top of protection against any two disk failures within the RAID group, Virtual Disk C is also able to sustain complete tray failure of any two trays in the StorageTek 6140 array, except Tray 4 (which contains two drives of the virtual disk) and Tray 0 (the controller tray). For example, if Tray 2 and Tray 3 were to both fail, Virtual Disk C would lose two disk drives from the RAID group but data would remain available. However, given the disk configuration in this example, a complete failure of Tray 4 would result in two failed disks within Virtual Disk C. In this failure scenario, Virtual Disk C would continue to service I/O, however it would not be able to sustain a second complete tray failure.

Striping Across Multiple Arrays

Striping across multiple StorageTek 6140 arrays may provide enhanced performance and reliability. When a RAID group is striped across more than one array, the multiple array controllers in the separate arrays provide multiple independent active paths to the data, which can enhance performance. As an example, assume a RAID 6 virtual disk with eight physical disks striped across eight independent StorageTek 6140 arrays. In this case, the host will use eight independent paths across eight independent array controllers to write data; the data paths are not limited to the same controller tray in a single StorageTek 6140 array.

In addition, striping against multiple arrays can help protect against loss of access to the RAID group if there is a complete controller tray failure. The loss of a controller tray affects access to all trays in that specific array. If a RAID group is striped across trays in a single array, a controller tray failure will disrupt access to the RAID group. However, if the RAID group is striped across multiple arrays with at most two drives in each array, the RAID group can sustain a failure of any one controller tray. And if the RAID groups is striped across multiple arrays with at most one drive in each array, this configuration can sustain any two complete tray failures — even two complete controller tray failures.

Because complete array controller failure is rare, striping across multiple arrays is primarily used to enhance performance. However, this configuration also provides enhanced reliability in protecting against multiple controller tray failures.

Summary of Best Practices

The following configuration details are recommended as best practices when implementing RAID 6 virtual disks on a Sun StorageTek 6140 array:

RAID 6 best practices for StorageTek 6140 arrays:

- Use no more than 8 disk drives within one RAID 6 virtual disk.
- Stripe across expansion trays for enhanced availability.
- Limit the total number of LUNs per RAID 6 virtual disk to six.
- Do not use very large, multi-terabyte LUNs due to long reconstruction times.

- *Use no more than eight disk drives per virtual disk*
As with RAID 5, RAID 6 performance starts to degrade as additional disk drives are added to a virtual disk. A maximum of 30 disks can be part of a RAID 6 virtual disk in a Sun StorageTek 6140 array. However, it is recommended to use no more than a total of eight disk drives to achieve best results for performance and number of LUNs per virtual disk, especially for very large capacity disk drives.
- *Stripe RAID 6 virtual disks across multiple trays*
For best availability, stripe the RAID 6 virtual disks across multiple trays in the StorageTek 6140 array. The ability to sustain up to two complete expansion tray failures can be achieved if proper pre-planning and considerations take place prior to creating virtual disks.
- *Limit the number of LUNs for each RAID group to six*
Under normal circumstances, it is best to limit the number the LUNs to a maximum of six for each RAID group. This configuration helps keep disk seek contention low and can provide better performance.
- *Avoid the use of very large multi-terabyte LUNs*
Since data regeneration is relatively slower for RAID 6, a multi-terabyte LUN reconstruction can take a significant amount of time, especially with very small files. Data availability is not affected during regeneration, and RAID 6 virtual disks can sustain a second drive failure without affecting data availability. However, a RAID 6 virtual disk is unable to sustain a third drive failure. Avoiding very large multi-terabyte LUNs whenever possible reduces the regeneration time and thereby reduces the risks of additional errors occurring during regeneration.

Conclusions

RAID 6 provides protection against two consecutive drive failures by using two independent parity calculations. Sun Microsystems takes advantage of the better performing P+Q algorithm implementation of RAID 6 and uses dedicated processors for RAID computations on intelligent RAID controllers in the StorageTek 6140 arrays for improved performance.

RAID 6 offers highly available storage for applications and implementations requiring extremely high fault tolerance while keeping cost down and performance at an acceptable level during normal circumstances. If best practices are followed, a RAID 6 disk group is capable of performing at rates very close to RAID 5 with a slightly higher cost. RAID 6 is also capable of providing reliability close to RAID 1+0 at a fraction of the cost.

Sun StorageTek 6140 storage arrays running firmware revision 7.10.x.x have RAID 6 capabilities and dedicated RAID computation processors which allow for enhanced RAID 6 performance. Using best practices configuration guidelines, a RAID 6 virtual disk is able to sustain up to two complete expansion tray failures within a Sun StorageTek 6140 array. Data availability may be substantially increased when RAID 6 virtual disks are striped across multiple StorageTek 6140 storage arrays.

About the Author

Said A. Syed has over 13 years of industry experience, including over 7 years with Sun. Said started with Sun as a System Support Engineer in Chicago supporting high-end and mid-range servers and Sun storage products. Said joined Sun's Storage Product Technical Support group in 2004 as the Sun Support Services global lead for Brocade SAN products. In this position, Said managed the Sun Support Services relationship with Brocade Support Services directly and supported world-wide Sun customers on high visibility, high severity escalations involving SAN infrastructure products and Sun's high-end and low-cost storage products, the Sun StorageTek 3000 and 9000 series arrays. In 2008, Said was promoted to Staff Engineer role within Sun's NPI and OEM array engineering group and is currently chartered with gaining in-depth understanding of how virtualization applications such as VMware ESX server, the Sun xVM platform, Sun Logical Domains (LDoms), Solaris™ Containers and other similar applications interact with Sun's modular and high-end storage arrays, the Sun StorageTek 6000 and 9000 series arrays.

Acknowledgements

The author would like to thank Andrew Ness, a Staff Engineer within the NPI and OEM array engineering group at Sun, for reviewing technical details outlined in this document for accuracy. Thanks also to Michael Jeffries for reviewing the document for

technical content, flow and verbiage. Further more, the author would like to extend his thanks to the Sun BluePrints team for assisting in technical writing, final review and publication.

References

Sun StorageTek 6140 array:

http://www.sun.com/storagetek/disk_systems/midrange/6140/

Anvin, H. Peter. "The Mathematics of RAID-6," online paper, July 2008.

<http://kernel.org/pub/linux/kernel/people/hpa/raid6.pdf>

Schroeder, Bianca and Gibson, Garth A. "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?" *FAST '07: 5th USENIX Conference on File and USENIX Association Storage Technologies*.

<http://www.usenix.org/events/fast07/tech/schroeder/schroeder.pdf>

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is

<http://docs.sun.com/>

To reference Sun BluePrints Online articles, visit the Sun BluePrints Online Web site at:

<http://www.sun.com/blueprints/online.html>

