HITACHI
Inspire the Next

Why Growing Businesses Need RAID-6 Storage

Enterprise-class Technology Offers Stronger Data Protection

White Paper

April 2008

Hitachi Storage Solutions

Hitachi Data Systems

Executive Summary

In today's information driven world, cost-effective data protection is essential to the health and well being of growing businesses. In a recognition of this need, many businesses use RAID technology to help protect the data in their storage systems. But not all RAID technology is alike. Different levels of RAID offer different levels of data protection. RAID-6, the focus of this paper, offers distinct data protection advantages over the common alternatives. The paper explains in simple, non-technical terms what RAID is, how RAID levels differ and why RAID-6 is important for your business.

Contents

Why Growing Businesses Need RAID-6 Storage

Enterprise-class Technology Offers Stronger Data Protection

White Paper

How Important Is Your Data?

When considering any technology that helps to protect your data, it's good to begin by taking a step back and considering the value of your business information. Could you deliver your products or services without your data? If you could not access your data for two, three or even four days could your business stay up and running? That is how long it might take to restore a host system from backup tapes if you have lost data due to a hard drive failure.

And hard drives do fail — more often than most people realize. If you look at the ratings for disk drives, it appears they seldom fail. Some are rated for more than one million hours of mean time between failures (MTBF), based on technical calculations. The reality on the ground is quite different from the impression created by MTBF ratings.

In practice, disk drives fail all the time, and the more drives you have, the more chances you have for drive failures. For example, each year about seven out of every one thousand Serial ATA (SATA) drives will fail, according to statistics published by drive vendors. The reality reported by several customer studies in large data centers shows a much higher failure rate (see Figure 1), closer to two to four failures per year per 100 SATA drives. Numbers like that underscore the need for RAID systems.

Figure 1. Hard Disk Drive Failure Rates

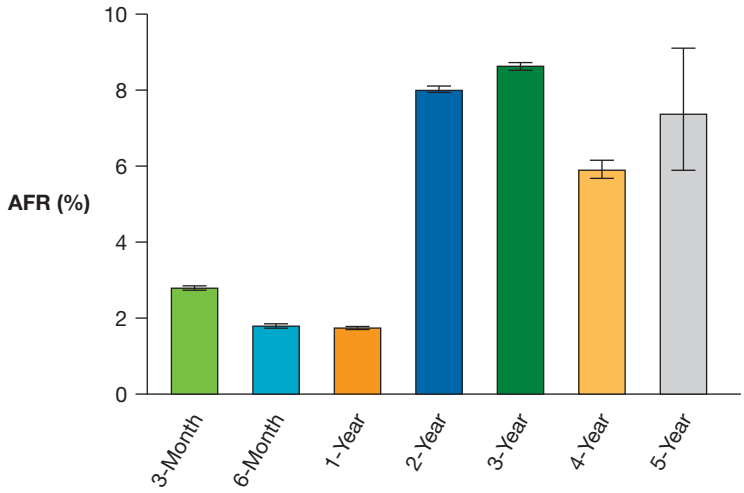

Studies of large data centers show a higher disk drive failure rate than MTBF ratings would indicate, as shown in this chart of annualized hard disk drive failure rates broken down by age groups.[1]

1 "Failure Trends in a Large Disk Drive Population," Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz Andr´e Barroso. Appears in the Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST'07), February 2007.

Understanding RAID Technology

While it may sound like something best left to computer experts, RAID is actually a simple concept. The term stands for "redundant array of independent disks." It is an approach that helps protect your data from hard disk failures. Storage systems link disks together into RAID groups in a manner that allows the disks in the group to back each other up. Ideally, when you lose a disk, the RAID group should be able to rebuild itself without losing any data.

That sounds easy enough, but here again the reality is something else. RAID comes in many levels, and not all levels are equal. Some RAID levels offer far greater protection than others. And as drive sizes increase so does the risk that you would not be able to rebuild a RAID group when a drive fails. In simple terms, the larger the disk and the larger the RAID group, the higher the probability that you would not be able to reconstruct data following a catastrophic disk failure.

To understand the value of RAID technology, it helps to first consider a technique called "parity." Parity allows you to reconstruct data from a failed drive using data that is held on the surviving drives and a parity drive in the same RAID group. The use of parity schemes enables more efficient use of storage capacity when compared to the full duplication of data, which is too expensive for many companies.

Storage administrators in large and enterprise-sized organizations understand these issues and concepts. That is why they seek out the best, most cost-effective protection for their RAID groups. That often means RAID-6. Most FORTUNE 1000® companies now use RAID-6, or a comparable level of RAID protection, in their high-end disk systems (see next section).

Even Seagate Technology and Microsoft Corporation are now recommending the use of RAID-6 (or full mirroring) with the use of large SATA drives, due to the inherent SATA bit error rate[2] (we'll explain more about big drives and bit error rates in the following discussion). And increasingly, small and medium businesses are coming to the same conclusion: they need RAID-6 to protect their data in a cost-effective manner.

Why RAID-6 Matters

Most RAID designs employ some level of parity, yet not all RAID designs are created equally. Parity refers to a code sequence for detecting errors across disk drives, either through the use of parity bits or parity blocks. With this technique in place in a storage system, a disk drive can fail and the surviving data and parity disks in the RAID group should be able to restore the information from the failed drive. That is because the surviving drives collectively have the data needed to rebuild the information from the failed drive. However, simply having a parity drive in a RAID group does not mean your data is fully and completely protected. For higher availability and nondisruptive operation without the use of full data duplication, you need a solution that makes use of dual-parity RAID technology.

In terms of reliability, RAID-6 is a big step up from RAID-5, a technology that is widely used in servers and storage systems, especially those designed for small and medium businesses. RAID-5 uses a single-parity approach, which means it is designed to recover from the failure of one drive, but not two drives, in a single RAID group. A two-drive failure in a RAID-5 configuration will lead to lost data. In contrast, in a RAID-6 configuration, any two drives can fail at the same time without any loss of data.

So why doesn't everyone go with RAID-6? The short answer is that RAID-6 configurations usually cost a little more up front, because they require an additional disk drive that would not be needed in a single-parity configuration.

2 WinHEC 2005 conference presentation from Willis Whittington, Seagate Technology, and Jeff Mastro, Microsoft Corporation.

But, in the long run, RAID-6 can actually be more economical because disk drives are relatively inexpensive, while lost data is very expensive to recover — assuming it can even be recovered.

Given the advanced level of protection provided by RAID-6, there is a growing movement to make RAID-6 a more widely adopted standard for storage systems, even those at the entry level end of the storage market. This movement is likely to gather momentum as larger SATA disk drives arrive — bringing with them a greater risk of data loss due to incomplete drive rebuilds.

Bigger Drives Bring Bigger Risks

While new 1TB (terabyte) SATA drives have entered the market, many system vendors are only now in the process of adding them to their product portfolios. With these larger disks comes an increased risk for data loss resulting from a failed drive rebuild in a RAID system.

Let's pause for a little background. As we noted above, if a drive fails in a RAID group, the other drives in that group must rebuild the failed drive from scratch. For a rebuild to succeed, all of the data on all of the sectors of the healthy drives must be readable. If any sector is unreadable, the rebuild work hits a wall and comes to a stop.

At that point, you have lost data or you have corrupted data if you are using RAID-5 or some other single-parity system. A RAID-6 configuration gives you a second chance if the drive group encounters an unreadable sector during a rebuild. This second chance comes with the second parity drive that is there to correct the errors in the process of rebuilding the RAID group.

So how likely is it that a RAID group would encounter unreadable data? It is more likely than you may think. Here's why: When a drive's read-and-write head reads data on a disk, it will rarely make a mistake. Any mistake is known as a "bit error." On average, bit errors occur at a predictable rate. Drive manufacturers can tell you the bit error rates for their products. This rate refers to the average number of bits of data that are transmitted incorrectly when data is being read on a disk. Today's server-grade SATA hard disk drives have a standard bit error rate of 1 unrecoverable read error per 125TB of data. This means that for every 1TB hard disk drive read, there is a 0.8 percent chance that one sector will be unreadable.

While bit error rates are extremely small, they become more significant with larger drives and larger RAID groups. That is because the more bits that are read on a disk, the more chances you have for bit errors — and the subsequent failure of an attempt to rebuild a drive.

So what is the risk? Using calculations by Adaptec, a storage equipment provider, we can see these probability rates (see Figure 2):

- With a typical RAID-5 configuration using 1TB SATA drives in an 11 drive RAID group, the probability of data loss due to the combination of a disk failure and the bit error rate is 7.7 percent.
- With a RAID-6 system of the same size, the probably of data loss from the combination of a disk failure and the bit error rate is virtually 0 percent.[3]

In other words, when you lose a disk drive in a RAID-5 configuration, there is a significant chance that you could also lose data.

[3] For details, see http://storageadvisors.adaptec.com/2005/11/01/raid-reliability-calculations

Figure 2. Probability of Data Loss

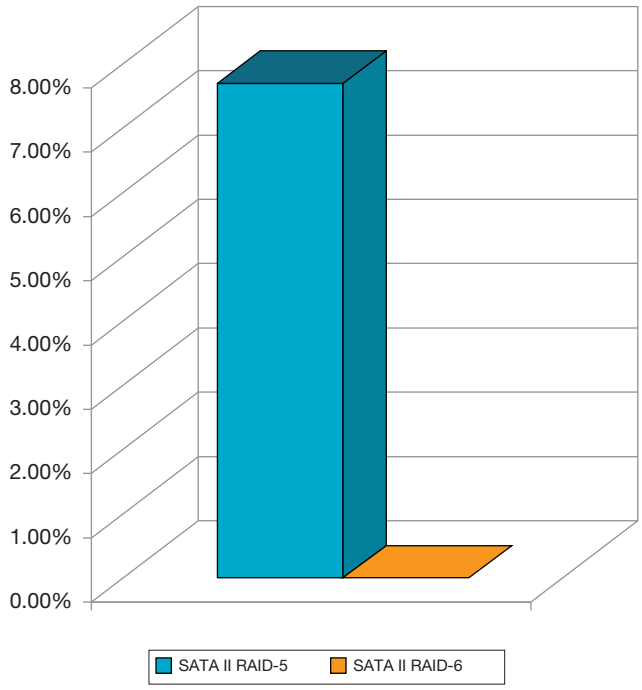

The probability of data loss due to disk failure and bit error rate (PDL_BER) is much lower with SATA RAID-6.

As drive sizes grow in the years ahead, this disparity will become even greater (see Figure 3). The SATA II road map calls for the arrival of 4TB drives in 2011, quadrupling the size of the 1TB drives that arrive in 2008.

Figure 3. Risk Probability Considering Growth in Drive Sizes

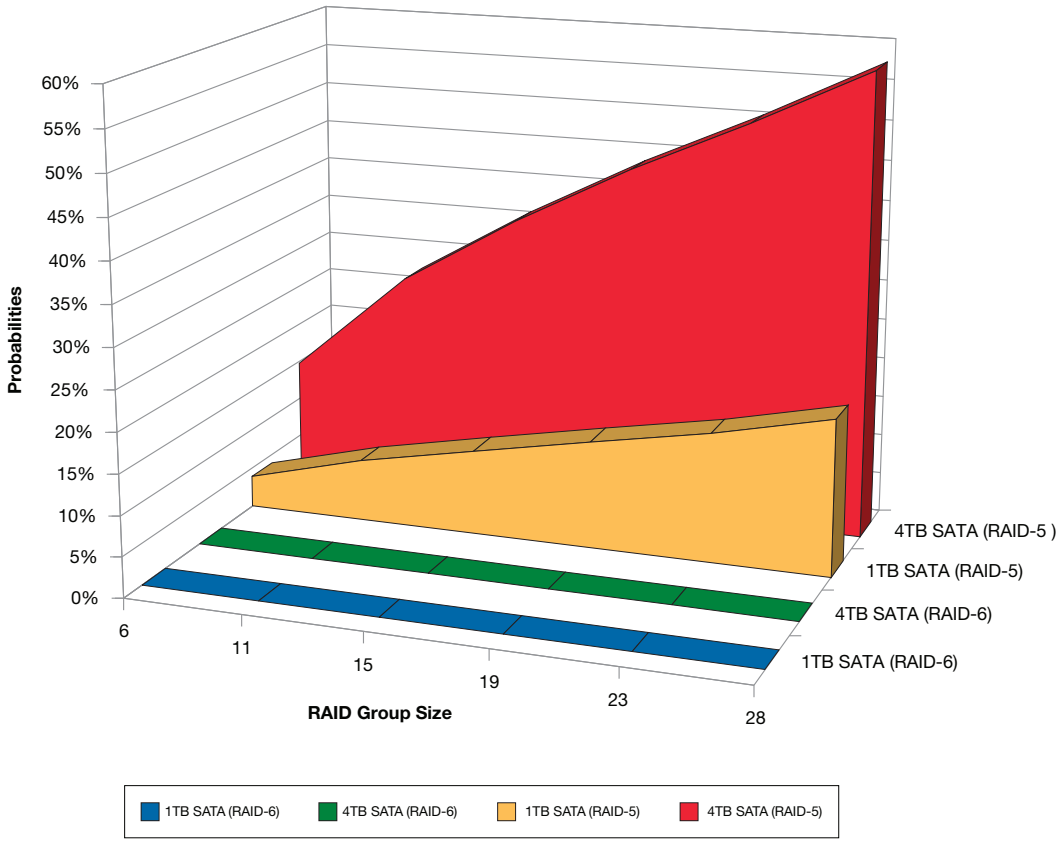

As the SATA II road map unfolds, the disparity between RAID-5 and RAID-6 groups regarding probability of data loss due to disk failure and bit error rate becomes even more evident.

An Added Risk

Want to add another piece to the puzzle? None of the above considers the risk of a second drive failure while a rebuild is in progress. If you have a RAID-5 system, then during the rebuild period you have no RAID protection at all. If a drive fails during this repair window, you have lost data.

Here again, larger drives equate to larger risk. With larger drives, it takes more time to rebuild a RAID group after a drive fails. The process of rebuilding a 1TB drive in an 11 drive RAID group takes at least 24 hours in an idle storage system. The same drive rebuild can take in excess of a week to complete in a busy system. During this time, your data is at a heightened risk if you are using RAID-5 or some other single-parity disk system.

In contrast, dual-parity RAID-6 systems protect you against data loss from a second drive failure while the RAID group is rebuilding the failed drive. In essence, during a rebuild a RAID-6 system provides RAID-5 protection for your data. When the failed drive is rebuilt, your system is back to RAID-6 protection.

While the core value of RAID-6 technology is data protection, dual-parity RAID systems also provide operational flexibility that RAID-5 cannot support. They allow a system administrator to postpone large SATA hard disk drive rebuilds to off hours when the storage system is the least busy. This reduces impact to application response times due to drive rebuild I/Os, without ever exposing the data to the risk of loss from a second drive failure.

Given that the typical rebuild time of a 1TB SATA drive can be as long as a week, it is essential to select storage systems that support dedicated hardware RAID-6 controllers over a RAID-6 software implementation. On software

RAID-6 storage systems, I/O intensive applications are exposed to crash and data corruption while a drive rebuild process is ongoing for an extended period. A drive rebuild process can saturate a storage system's processor (or processors) with I/O requests, which will prevent it from processing host I/Os. At that point, host I/O requests will time out and the host applications will crash — and consequently most likely corrupt the data.

A dedicated hardware RAID-6 storage system, on the other hand, offloads the parity computation and drive rebuild process to a separate co-processor, which frees up the controller to manage host I/O requests. This design ensures better application response time and enhances data integrity even when the storage system is under stress by a long hard disk rebuild process.

Look to Hitachi Data Systems for RAID-6 Protection

Hitachi Data Systems incorporates dedicated hardware RAID-6 technology in many of its modular storage products, including those designed for small and medium businesses and the workgroups, departments and branch offices within larger enterprises.

That is the case with Hitachi Simple Modular Storage, a product designed for organizations that want to venture into shared storage. It incorporates RAID-6 technology in a package priced for the budgets of growing businesses.

Hitachi Simple Modular Storage was designed to offer storage with no single point of failure. In simple terms, this means that multiple things can go wrong and you are still in business. In addition to RAID-6 protection, Simple Modular Storage offers dual controllers in an "active-active" architecture. This means that if one storage controller is suddenly unavailable, then the system makes all data volumes available via a second controller.

In general, Simple Modular Storage is designed to make shared storage easy. Right out of the box, easy-to-use wizards and utilities simplify installation and configuration. And should a disk drive malfunction, the Hitachi Storage Navigator Modular 2 software alerts you to the problem and directs you to contact a support center to request a replacement.

Moving Forward with Better Data Protection

While it is hard to know when a hard disk drive is going to fail, a forward-looking business can and should plan as if failure were inevitable. That, in essence, is what you do when you buy a RAID-6 storage system. You don't know when it will happen, but you do know that at some point in time a disk drive is going to fail in your storage system.

As you plan for that day and contemplate the increased risk of data loss that comes with larger SATA disk drives, RAID-6 makes all the more sense. Ultimately, RAID-6 buys you a lot of insurance for very little money — and that is value that is hard to beat.

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

WHP-280-00 LKW April 2008