RAID Technology and Data Storage Today

Jeffrey Doto

Brandon Krakowsky

April 15th, 2007

Abstract

With information generation and data transfer speed at an all time high, data storage is fast becoming one of the fastest growing industries in the world. Enterprise-level corporations, end-users, and the government, all need to deal with an ever-increasing plethora of data. RAID, or "Redundant Array of Independent Disks", provides a very flexible and reliable storage system which can both enhance performance and leverage fault tolerance. We discuss the various RAID types and their uses, as well as their individual drawbacks. Future aspects of RAID, as well as current enterprise-class storage issues are mentioned.

Table of Contents

Motivation

Information Overload

The issue of data storage is becoming increasingly important, as more and more information is stored in "digital" form. Assuming that the average digital file gets replicated three times, the technology research firm IDC determined that "the world generated 161 billion gigabytes - 161 exabytes - of digital information last year" (Brian Bergstein 2007). Data management is one of the fastest growing industries in the world. According to World Wide Learn, "The information sector has the second highest projected job growth rate. Publishing, the Internet, and telecommunications should see real output climb to $1.6 trillion by 2014" (Wendy Croix).

But where is all of this information stored? "IDC estimates that the world had 185 exabytes of storage available last year and will have 601 exabytes in 2010. But the amount of stuff generated is expected to jump from 161 exabytes last year to 988 exabytes (closing in on 1 zettabyte) in 2010" (Brian Bergstein 2007). How do companies today deal with what Linda Null and Julia Lobur refer to as the "information explosion" (Lobur 2006)?

> "Companies worldwide today face a tremendous explosion of data produced at every level of the organization, from email, databases, the Internet, and e-commerce to image-based applications that produce video and audio files. Booming trends in the audiovisual market, including digital cameras and camcorder products, are expected to produce 700 petabytes of data per year, dwarfing the amount of data produced in traditional IT applications. Movies will be produced and distributed digitally … Even broadband and new cell phone applications will contribute to significant, near-term storage growth.
>
> With the accelerated use of content-rich systems, the need to store, manage and protect the important data produced is more important than ever. Storage systems will race to increase capacities and performance…" (John Woelbern and Sony Electronics' Tape Storage Solutions Div 2003).

Proliferation of the Internet

The proliferation of the Internet has served to only increase the need for reliable data storage. Hosting companies need to deal with millions of web pages. And with the popularity of web-based audio and video, these same companies now need to deal with data storage on a whole new level. When was the last time you hit a website that did not have some kind of embedded audio or video file? It has become commonplace, because people are expecting it. A simple webpage with only text and graphics is a thing of the past.

Hosting a few video files is one thing, but what about companies who specialize in audio or video On-Demand delivery? YouTube is a popular website which allows users to upload and share their own video content. Google Video does the same. How do these companies deal with the management of all of this content? With new people turned on to these services everyday, hosting companies need to look to the future for storage and delivery options.

Perhaps video is still a fairly new concept for some. What about digital audio? Remember Napster? Of course, they still provide services and conduct business, but not in the same way as they used to. But that still hasn't changed the fact that they have to deal with hosting all of that content. iTunes, Rhapsody, and MP3.com do the same, providing users with the ability to download audio files with ease. Most of the companies in this business have expanded their services to include video delivery as well, so it is all really one in the same issue. Do not forget about podcasting and vodcasting!

People have been uploading photos for some time now. It is a pretty common way to share your favorite shots. Services like Flickr and Photobucket make it so easy to upload, create an album, and send a link to family and friends. Who even thinks about where the photos go? Again, as these kinds of services become more popular, hosting companies need to look to the future for reliable data storage.

Database driven sites rely heavily on the integrity of their backup systems. Most people do not have a local backup of their blog or wiki entries. Can you imagine if your MySpace account was accidentally deleted or lost? Companies like Amazon, eBay, and Citizens Bank have huge backbones. Their sites need to be up and running at all times. Without a doubt, any company using online transactional processes to conduct business needs to have a reliable storage and backup system, but what about the National Security Agency? Did you know that the government has a database of phone call records of "tens of millions" of Americans (Cauley 2006)? Can you imagine if they lost it?

E-mail has become the most popular mode of communication. When you send a message, where does it go? Messages in your "Sent" folder must be stored somewhere. What would happen if you lost an important work-related email? If you are like most, email is a lifeline. For larger corporations, email storage is a must.

Sarbanes-Oxley and HIPAA

In 2002, President George W. Bush signed into office the Sarbanes-Oxley (SOX or Sarbox) act. This act changed the way publicly held businesses were held responsible for their accounting practices, which directly affected their Information Technology staff. The businesses now had to adhere to strict regulations requiring longer data retention time, which would allow auditors to look back through the businesses' financial histories, as well as their email records (Randolph Kahn). It is amazing when you think of the amount of documents a business or institution can generate: business transactions, invoices, purchase orders, contracts, payments made, website records, internal documents. The Sarbox law specifically sites: "records relevant to the audit or review…memoranda, correspondence, communications…(including electronic records)" (Jones 2005)

Section 404 of the Sarbanes-Oxley Act in particular, has placed an added stress to IT managers, as well as an enormous cost of maintenance, which some experts in the industry have been critical of (Economist 2004; Jones 2005; Lallande 2005) These added regulations have certainly been a boon to data storage vendors, as practically every enterprise-class storage vendor will mention their Sarbanes-Oxley compliance in regards to their products (Consulting 2005; Consulting 2005; Corporation 2005; Dell 2007) Costs are predicted to be somewhere between $20-$28.8 billion in compliance-related technology (Economist 2004; Sullivan 2006).

Similar to the Sarbanes-Oxley Act, Title II of the Health Insurance Portability and Accountability Act (HIPAA) of 1996 has reformed the electronic aspect of Healthcare insurance plans, service providers, and employers. These regulations have the intent of promoting Electronic Data Interchange, or EDI, throughout the US healthcare system. Again, longer data retention times will be required, which greatly benefits the storage sector of the IT industry (Consulting 2005).

Data Management for the User

With all of these advances in Internet technology, where does it leave the user? Why do we need to be concerned with data storage? Did you know that every time you go to a website, all of the images and media files are "cached" or stored on your local machine for rapid access? Any idea where that directory is located? You would be surprised how quickly you amass "temporary internet files".

Downloadable content needs to be stored somewhere. CDs are a lot less popular now than they used to be. Music fans would rather download an entire album from iTunes, or maybe even just download their favorite song off of the album. Who hasn't ripped audio files from an audio CD or imported a CD directly into iTunes? Where are these files stored? If you are like most, your music collection is growing rapidly. As is the amount of space you use to store it. Can you imagine if you lost all of your music?

And while video On-Demand becomes more popular, so does the use of webcams and video editing software. Digital video cameras make it very easy to transfer video to your computer. Hard drives fill up quickly when you are editing video.

Digital cameras are commonplace since prices have come down. Taking pictures and editing them has never been easier. Everybody is a photographer. A personal library of photos could easily grow to great proportions, let alone the increasing complexity of the file directory tree.

Microprocessor Technological Advances

As microprocessor technology continues to improve, so does primary memory size. If Moore's law holds true, these things will only continue to improve. What does that mean for overall computer performance? It means that any resulting improvements will be marginal unless secondary data storage progresses at close to the same rate (Peter M. Chen 1994). This idea is governed by Amdahl's Law which predicts that overall performance enhancement is limited by the slower parts of the overall system (Lobur 2006).

Increased microprocessor speed opens the door to newer processor-intensive applications. It also makes new things possible with older applications. As mentioned earlier, things like audio recording, video editing, and photo editing have all become commonplace for the average user. This capacity creates the need for high-performance secondary storage (Peter M. Chen 1994).

Introduction

Magnetic Disk Technology

Before we can fully understand RAID, we must first understand the inner workings of a magnetic hard disk, and the problems often associated with a disk.

As we all know and take for granted that a computer has memory, and can store information on that memory. In the case of our current generation of laptop and desktop computers, a magnetic hard disk stores this information in a semi-permanent fashion. We are also familiar with the proverbial "hard drive crash", and most of us have unfortunately experienced the loss of a crucial piece of information due to one of these crashes. It is important to first understand a little background about how a hard disk works to understand how it can fail.

A "hard" disk is different from the also well-known "floppy" disk because it has a fixed, rigid platter that is coated in a magnetic material, which stores the actual information. This rigid platter has a motor that causes it to spin at speeds varying from 5400 rotations per minute (RPM), up to 15,000 RPM. An actuator arm with a read/write head something akin to the stylus of a record player, moves in and out, reading and writing data to and from the disk (Lobur 2006). Thus, there are several obvious points of failure in a hard disk: moving mechanical parts can deteriorate and break, or there can be manufacturing defects, such as surface imperfections, or dust contamination. The chamber to the disk must be hermetically sealed so no dust can interfere with the surface of the disk. Thus, overtime, the mechanics of the hard drive are bound to fail, and all hard drives are sold with a Mean Time To Failure rating, or MTTF, which means that ultimately, every user must not face the question of "will my data be lost", but "when will my data be lost". For the average home user, this may not be of terrible importance, but for a multimillion dollar business, this could be catastrophic. The need for fail-safe, highly reliable storage is omnipresent in big businesses of all kinds, government, academic institutions, and just about any large institution one can think of.

RAID Defined

"RAID" was originally coined by David A. Patterson, Garth Gibson, and Randy H. Katz of the University of California at Berkeley, in their 1988 paper, "A Case for Redundant Arrays of Inexpensive Disks" (David A Patterson 1988). The original meaning of the acronym was "Redundant Array of Inexpensive Disks", but because the term "inexpensive" is relative, "Redundant Array of *Independent* Disks" is now generally accepted as the proper meaning (Lobur 2006).

RAID is a method of combining several hard disk drives into one logical unit, called a LUN, or a "logical unit number", appearing as a single device to the host system (Technick.net; Poelker 2005). Disk protection is achieved any number of ways based on the particular implementation of RAID. "RAID technology was developed to address the fault-tolerance and performance limitations of conventional disk storage" (Technick.net).

Disks can also be combined in this fashion, but without RAID protection. This setup is known as "Just a Bunch of Disks" or JBOD (Lobur 2006).

History

By the time "A Case for Redundant Arrays of Inexpensive Disks" had been written in 1988, magnetic disks had become very large, dense, and costly pieces of equipment. Often requiring a strictly controlled environment, use of these disks was generally restricted to only the largest computer systems. With the number of personal computers on the rise and an increasing reliability on electronic data storage, disk designers began to design smaller and cheaper disks, as an alternative to these "single large expensive disks" (SLEDs) (David A Patterson 1988; Lobur 2006).

When compared to the larger expensive discs, it was found that these smaller "inexpensive" disks were at least equal to, if not better than, the larger expensive disks, in terms of both I/Os per second per actuator and price per megabyte. Most importantly, these smaller disks had the same basic functionality as the larger disks, including an embedded SCSI chip controller, which could be used as a Direct Memory Access device at the end of the SCSI bus (David A Patterson 1988).

Based on the fact that the smaller disks had the same functionality as the larger disks, storage systems could be built as arrays of these inexpensive disks. The idea was to not only overcome the problems associated with the single large disks, but to also achieve better performance and increased reliability for both small and large data storage systems. Interleaving the discs would allow for large supercomputer transfers, while keeping them independent would allow for many smaller transaction-processing transfers (David A Patterson 1988). The different levels of RAID address these scenarios to varying degrees.

Reliability

RAID addresses the issue of disk reliability by making use of "redundancy". This means that data is redundantly distributed across all (or some) of the disks for the purposes of fault tolerance and data protection. If one disk fails or if a piece of data becomes corrupted, it can be recovered from one or more of the other disks (Technick.net).

In dealing with redundancy, two things need to be considered:

1. Calculating the redundant information in the event of an error (Peter M. Chen 1994).
2. The method of distributing the information across the disks (Peter M. Chen 1994).

These two points are considered and addressed based on the type of RAID.

Performance

Disk performance is enhanced based on the fact that there are multiple disks working in parallel. Depending on the distribution of data, different pieces of information can be read from different disks at the same time (Adaptec 2007).

RAID Technology

RAID Level 0

Also known as "Non-Redundant", RAID Level 0 uses a technique called "striping" for distributing information across the disks. Data is broken down into individual blocks, each block written to a separate disk drive. Redundancy is not employed and there is no parity generated, resulting in no data protection. If one striped drive fails, they all fail. In fact, as the number of disks increases, so does the probability of failure (Peter M. Chen 1994; Lobur 2006; Corporation 2007).

RAID Level 0 provides excellent write performance since the load of storing data is spread out across multiple disks and there is no overhead in updating redundant information. This scheme is excellent for video production and editing, image editing, pre-press applications, and any application requiring high bandwidth (Corporation 2007).

RAID Level 0 requires at least 2 drives and is typically very inexpensive to implement since it uses the minimal number of disks (Corporation 2007).

RAID Level 1

Also known as “Mirrored”, RAID Level 1 uses a technique called “mirroring” or “shadowing” for distributing information across the disks. All data is written to at least two separate physical disks, so there are always two copies of the information. There is no parity generated but 100% data redundancy is provided. If one disk fails, data can easily be recovered using its mirror image (Peter M. Chen 1994; Lobur 2006; Corporation 2007).

RAID Level 1 provides excellent data protection. Write performance is compensated slightly based on the fact that all data is written twice. Read performance is better than RAID Level 0 since two discs can be read at once. This scheme is often used for accounting, payroll, financial, and any application requiring very high availability (Corporation 2007).

RAID Level 1 requires at least 2 drives and is typically expensive to implement since it uses twice the number of disks (Corporation 2007).

RAID Level 2

RAID Level 2 uses a technique similar to “striping” for distributing information across some number of disks. Strips consist of single words split at the bit level, spread across the data disks, one bit per disk. Hamming codes are generated for each word and are also spread across separate Error Correcting Code disks. When data is read, it is cross-referenced with these Hamming codes to check for errors. Error correction is performed “on the fly” by subtracting the error-free data on the other disks from the Hamming code information. Multiple disks are required to identify a failed disk (Peter M. Chen 1994; Lobur 2006; Corporation 2007).

RAID Level 2 provides high data transfer rates. However, write performance is compensated highly since every write requires the additional calculation and storage of Hamming codes. This scheme is hardly ever used and has no commercial implementations mainly because it requires considerable overhead in generating the Hamming codes. Also, since the number of ECC disks is proportionate to the log of the number of data disks, the ratio is higher for smaller words making it inefficient and expensive (Peter M. Chen 1994; Lobur 2006; Corporation 2007).

RAID Level 3

RAID Level 3 uses a technique called “bit-interleaved parity” (similar to “striping”) for distributing information across some number of disks. Data is split at the bit level and spread across the data disks. Single parity bits are generated and stored on a separate parity disk (Peter M. Chen 1994; Lobur 2006; Corporation 2007).

The read and write performance of RAID Level 3 is compensated since multiple operations are not supported at once and all disks are used for every operation. As a result, RAID Level 3 performs poorly in situations where there are many little data transfers, but performs very well when dealing with large blocks of data. This scheme is often used for video production and live streaming, image editing, video editing, prepress applications, and any application requiring high throughput (Poelker 2005; Lobur 2006; Corporation 2007).

RAID Level 3 requires at least 3 drives (Corporation 2007).

RAID Level 4: Block Interleaved Parity

The implementation of RAID 4 is similar to that of RAID 3, instead RAID 4 writes data blocks of varying size, instead of bits of data across the disk, with one disk dedicated to parity (Peter M. Chen). These blocks of data are known as the “striping unit”, and can span multiple disks, similar to that found in RAID 0 (Peter M. Chen). This usage of block striping increases the random access performance compared to RAID 3, but because only one disk functions as dedicated parity, write operations may pose a bottleneck. Bottlenecks can be found in smaller write operations, where data may not span the entire group of data disks, opposed to write operations where data spans all of the data disks (Peter M. Chen; Lobur 2006).

A large write operation's parity is computed by XOR'ing the data found across all drives, which has a relatively low computational overhead. Suppose a data block spans 4 disks, with one stripe on each disk, called stripe 1 - 4, for each disk respectively. The parity disk will have one dedicated parity block for the other four disks' content. Thus, this parity block will be equal to the XOR of all the other data stripes [1, 8].

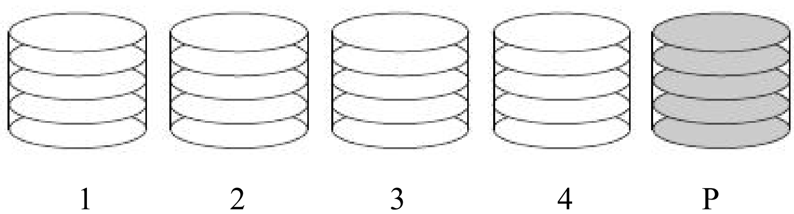

Figure 1 – RAID 4: Parity 1-4 = Stripe 1 XOR Stripe 2 XOR Stripe 3 XOR Stripe 4. Picture adapted from Chen, et al.

This is somewhat different from a small write operation where only one data disk is being accessed, and the single parity disk is being updated accordingly. The new parity for a small write is calculated as follows:

New Parity = (old data XOR new data) XOR old parity

This calculation requires 2 disks, 2 reads, and 2 writes, opposed to all disks, and one read and one write (David A. Patterson 1988). This limits one write at a time per disk. A small write is so relatively slow to a Level 1 RAID that it becomes impractical to justify its usage [9]. While RAID 4 makes reading 4 disks in parallel very efficient, its writing is so compromised that it is totally impractical, especially in high-throughput environments where multiple requests at a time need to be handled. RAID 5 overcomes this problem by distributing the parity over all disks, thus eliminating the dedicated disk bottleneck.

RAID Level 5: Block Interleaved Distributed Parity

As we previously stated, RAID 4 suffers from the impractical bottleneck of having only one parity disk. RAID 5 overcomes this problem by distributing the parity information across each disk. This allows multiple individual disk writes per group (David A Patterson 1988). The advantage of RAID 5 is most easily seen by visualizing the structure with a diagram.
Figure 2 details this.

RAID 5:

	Disk				
Sector	0	1	2	3	4
0	0	1	2	3	P0
1	5	6	7	P1	4
2	10	11	P2	8	9
3	15	P3	12	13	14
4	P4	16	17	18	19

RAID 4:

	Disk				
Sector	0	1	2	3	4
0	0	1	2	3	P0
1	4	5	6	7	P1
2	8	9	10	11	P3
3	12	13	14	15	P4
4	16	17	18	19	P5

Figure 2 – Optimal RAID 5 parity setup. Observe how in the RAID 5 depiction, the parity (P Blocks) are distributed throughout all disks, instead of being located on just one disk, as in the case with RAID 4. This means that if the machine needed to write to sector 0 of disk 3, and sector 4 of disk 1, the operation can occur simultaneously because the parity for the disks is independent. In contrast, notice how in the diagram for RAID 4, there is a conflict if you try to service the sectors 0 and 4 simultaneously: because the parity is contained on one disk, you can only write one group at a time. This is the cause for the aforementioned bottleneck.

The added benefit of distributed parity gives RAID 5 the most optimal characteristics from all of the RAID levels discussed so far. It offers data redundancy with small write times close to that of RAID 1, yet maintains the larger storage capacity of a RAID 3 or RAID 4 configuration [9]. In a typical RAID 5 setup, if you had 4 500 GB drives, totaling 1 terabyte of storage, and these disks were put into a RAID 5 array, you would still have 75% of the usable disk space, opposed to a RAID 1 configuration, you would lose 50% of the disk space, because you are strictly mirroring (Apple 2007). This additional parity reading and writing must be performed by an on-board controller, which adds some additional cost to RAID 5. We will cover controllers in a later section of the paper.

It is important to note that a RAID 5 configuration still has drawbacks, the main one being that if two disks are lost at once, you have no way of recovering the data. We will see in the next section that there are additional ways to rescue an array of disks that has had multiple simultaneous failures. RAID level 5 requires at least 3 disks to implement.

RAID Level 6: Block interleaved Striping with Dual Error Protection

Patterson, et al. envisioned the first 5 levels of RAID, but additional research has been done to develop additional RAID schemes that overcome some weaknesses of the aforementioned schemes. RAID 6 protects against multiple disk failure data loss by employing both Parity (P), and Reed-Soloman (Q) codes. This also requires additional storage compared to the RAID 5 scheme; in general, if you can store the data on N disks, RAID 6 requires N+2 disks, with a minimum of 4 disks (Peter M. Chen; Lobur 2006). The additional of the Reed-Solomon coding scheme adds in an additional overhead into the system, instead of 4 reads and writes, there are now 6: the 4 standard read-writes, plus an additional update on the P and Q blocks (Peter M. Chen). This additional computation may not be available on all RAID controllers; in fact, Apple Computer's popular X-Serve RAID, does NOT offer a RAID 6 configuration (Apple 2007). In general, RAID 6 is less common than 0, 1, 3, 5, and 01. We found only one array currently offered that supports RAID 6, the AC&NC JetStor 416iS (Guide 2001).

Hybrid RAID

RAID X+Y vs. RAID Y+X

Interestingly, you can RAID a RAID array. When we say that, we mean that if your controller permits it, you can, for instance, take a striped disk array (RAID 0), and Mirror it (RAID 1) (Guide 2001). You can think of this as "nesting" the two disk arrays. There is therefore, a difference between taking a striped set and mirroring it, and striping a mirrored set. The latter, RAID 1+0, would be implemented for more a performance gain in mirroring. The more common implementation is RAID 0+1, and is found in most enterprise-class RAID devices (AC&NC 2007; Computer 2007).

This opens interesting combinations, which can be implemented to overcome any shortcomings of other RAID levels. One interesting nested group would be RAID 5+1, which would be a mirrored RAID 5 combination, for the individual who requires ABSOLUTELY no chance of data loss. An unfortunate drawback of hybrid RAID systems is that they are very costly, but if you have the means, they are very beneficial.

Future RAID

RAID DP

Conventional RAID schemes which use a single parity drive (RAID 3 for example) can easily recover one failed disk. But what if a second disk fails before the previous has been rebuilt? This is very common considering the amount of time it can take to rebuild an entire drive (Sunstar Company 2007).

RAID DP uses a second parity disk which can rebuild a second drive in a double disk failure situation. Parity for the second drive is calculated differently than the first (Sunstar Company 2007).

RAID Z

The future may not be too far away. Sun introduced the ZFS file system with the release of their Solaris 10 operating system. The "Z" stands for Zettabyte, which to put in perspective, the IDC estimates that by 2010, there will be 988 exabytes, just under a zettabyte, in all computer storage worldwide [29]. ZFS is a 128-bit file system, which gives its user 16 billion times the capacity of 32 or 64 bit systems (Microsystems 2007).

RAID Z incorporates all of the features of RAID 5, but further benefits from the "ZFS transactional model … utilizing the ZFS checksum mechanisms to prove the integrity of the data before handing it back to the applications." RAID Z can not only recognize corrupted data, but also return it corrected, made possible by the very intelligent ZFS file system (Sun Microsystems 2007). This is all done via software, thus bypassing the need for an external controller, which, for the cost-conscious of us out there, this can be very practical.

Apple's new operating system Leopard, will feature the ability to format a drive using Sun's ZFS system, and therefore implement RAID Z (Sun Microsystems 2007).

RAID Controllers

A RAID controller acts as an interface between the host system and the RAID system. It serves to manage the physical disks and deliver the data to the host system (Freeman 2004).

Interfacing with RAID

A storage system can be connected to the host system in a number of ways. One way uses "Serial Advanced Technology Attachment", or SATA, which communicates over a serial link (Planet 2003). Another way is "Small Computer System Interface", or SCSI, which communicates using a series of commands and is typically faster, but more expensive (Planet 2004; Adaptec 2007). iSCSI and Fibre Channel are other technologies typically used (Freeman 2004).

Software-Based RAID

A type of "Direct Attached Storage", or DAS, a software-based RAID system functions through the operating system. It is typically lower in performance than hardware-based RAID, "due to the lack of dedicated hardware" (STATS 2004). It relies on the host system CPU for all operations. Software RAID is very easy to set up and is typically very flexible (STATS 2004).

Hardware-Based RAID

Another kind of DAS, hardware-based RAID, functions similarly to Software-Based RAID except that it uses an additional piece of hardware for interfacing with the external disks. These controller cards off-load processor-intensive RAID operations to enhance performance. They may use SATA, SCSI, iSCSI, or Fibre Channel technology (Engine; Freeman 2004; Adaptec 2007).

External Hardware RAID

External hardware-based RAID systems use a controller that is completely removed from the host system. All RAID functions are off-loaded to a microprocessor inside an external RAID controller, independent of the host system. Again, they may use SATA, SCSI, iSCSI, or Fibre Channel technology (Engine; Freeman 2004; Adaptec 2007).

Implementation

Enterprise Class Storage vs. Personal Storage

To the storage neophyte, every hard drive looks the same, whether the hard drive is sitting in your home computer, or in a million-dollar disk array at NASA. The fact is, an enterprise, or "big business" level hard drive, has a great deal of advanced considerations that must go into making it "enterprise worthy". The environment required by the drives' functionality, and the purposes for which the drives will be used, dictate special engineering considerations.

Most enterprise drives are typically run in an array, or grouped fashion, much like we have discussed with the RAID disk arrays. In this sort of environment, there is an added amount of heat and vibration from the multiple drives which can detrimentally affect the performance of a single hard drive (Anderson 2003). Enterprise level systems must also run constantly, 24 hours a day, 7 days a week, with little or preferably no down time. This can mean an enterprise hard drive with a 1,000,000 Hour MTTF rating, compared to a personal hard drive with an MTTF of 300,000 hours. The extra engineering that must go into this can add a significant cost. This may include adding in extra gaskets to keep unwanted vapors out, desiccants and activated charcoal to remove humidity or volatile solvents from the air, and chassis modifications to absorb vibrations and keep the unit cool (Anderson 2003).

Historically, SCSI has been the choice for high-level storage applications, with most of the aforementioned engineering going into SCSI drives (Anderson 2003). Since about 2003, though, there seems to be a trend toward SATA replacing SCSI installations. Null and Lobur said that 80% of the enterprise level drives in use at their time of writing were SCSI, but, an InfoTechTrends poll from 2005 showed that 51% of individuals making non-desktop purchases selected SATA, versus only 10% choosing Parallel SCSI, and 11% Serial Attached SCSI. So it seems the trend is going towards SATA drives, with their much cheaper price per capacity. SCSI drives have a current storage limit of 300 GB, and at that capacity level can cost close to $900, whereas one can purchase a SATA 500 GB drive for $145.00. Fast SCSI drives running at 15,000 RPM can cost over $1000 for a 147 GB drive! In general, SCSI drives are anywhere from 2 to 4 times more expensive than SATA drives, for a fraction of the storage space. This cost has its benefits, though, as seek time on a 15,000 rpm SCSI drive is just 3.5 ms, while a comparable SATA drive has an 8-9 ms seek.

Conclusion

Information is being generated at an alarmingly fast rate. The world relies heavily on data storage. RAID provides a reliable and safe way to store and manage data, but there are many options. Depending on implementation, RAID can be configured for enhanced performance, increased reliability, or some combination of both.

Data storage has far and wide-reaching applications, from securing our nation's vital business and government documents, to speeding up the process of audio and video editing. An enormous industry has grown around providing secure data retention services, customized to the end user. The Sarbanes-Oxley law has made data integrity a necessity, putting additional pressure on those responsible for overseeing the proper implementation of reliable backup.

A news article released this week detailed how the White House recently "lost" 5 years of confidential government related email. This is certainly bound to generate a great deal of controversy in the press (Steiner 2007). Imagine the repercussions! Storage is more than a convenience … it is a necessity.

Appendix

Jeffrey Doto

Abstract

Motivation

- Sarbanes-Oxley and HIPAA
- Magnetic Disk Technology

Raid Technology

- RAID Level 4: Block Interleaved Parity
- RAID Level 5: Block Interleaved Distributed Parity
- RAID Level 6: Block Interleaved Striping with Dual Error Protection
- Hybrid RAID: RAID X+Y vs. RAID Y+X
- RAID DP
- RAID Z

Implementation

- Enterprise Class Storage vs. Personal Storage

Conclusion

Brandon Krakowsky

Abstract

Motivation

- Information Overload
- Proliferation of the Internet
- Data Management for the User
- Microprocessor Technological Advances

Introduction

- RAID Defined
- History
- Reliability
- Performance

RAID Technology

- RAID Level 0
- RAID Level 1
- RAID Level 2
- RAID Level 3
- RAID DP
- RAID Z

RAID Controllers

- Interfacing with RAID
- Software-Based RAID
- Hardware-Based RAID
- External Hardware RAID

Conclusion

References

AC&NC. (2007). "http://www.acnc.com/02_01_jetstor_sata_416s.html." from http://www.acnc.com/02_01_jetstor_sata_416s.html.

Adaptec. (2007). "Let's Talk About RAID." Retrieved April 7, 2007, from http://www.adaptec.com/en-US/products/raid_tech/_education/talk_about_raid.htm.

Anderson, D. (2003). "You Don't Know Jack about Disks." Storage **1**(4).

Apple. (2007). "Apple Servers." Retrieved April 15, 2007, from http://www.apple.com/xserve/raid/performance.html.

Brian Bergstein, A. T. W. (2007) "Tech Researchers Calculate Digital Info." Wired **Volume**, DOI:

Cauley, L. (2006) "NSA has massive database of Americans' phone calls." USA TODAY **Volume**, DOI:

Computer, A. (2007). "http://www.apple.com/xserve/raid/management.html." from http://www.apple.com/xserve/raid/management.html.

Consulting, K. (2005). "HIPAA Compliance Brief."

Consulting, K. (2005). "Sarbanes-Oxley Compliance Brief."

Corporation, A. C. N. (2007). "RAID Tutorial." Retrieved April 8, 2007, from http://www.acnc.com/raid.html.

Corporation, H. P. (2005) "HP Solutions for Sarbanes-Oxley Compliance." **Volume**, DOI:

David A. Patterson, G. G., and Randy H. Katz (1988). "A Case for Redundant Arrays of Inexpensive Disks (RAID)." ACM.

David A Patterson, G. G., and Randy H Katz (1988). A Case for Redundant Arrays of Inexpensive Disks (RAID).

Dell (2007) "Dell Data Archiving Solutions." **Volume**, DOI:

Economist, T. (2004). "File That." Economist (March 4, 2004).

Engine, G. S.-T. E. S. "About RAID Controllers." Retrieved April 15, 2007, from http://industrial-computers.globalspec.com/LearnMore/Industrial_Computers_Embedded_Computer_Components/RAID_Products/RAID_Controllers.

Freeman, B. (2004) "Storage Basics: Choosing a RAID Controller." EnterpriseStorageForum.com **Volume**, DOI:

Guide, P. (2001, April 17, 2001). "Multiple (Nested) RAID Levels." Retrieved April 15, 2007, from http://www.pcguide.com/ref/hdd/perf/raid/levels/mult.htm.

Guide, P. (2001, April 17, 2001). "RAID Level 6." Retrieved April 15, 2007, from http://www.pcguide.com/ref/hdd/perf/raid/levels/singleLevel6-c.html.

John Woelbern, d. o. O. m. s. and S. J. Sony Electronics' Tape Storage Solutions Div (2003) "Does Tape Backup Have a Future? - SAIT." **Volume**, DOI:

Jones, K. C. (2005). "Survey: IT, Other Execs Want Better Sarbox Tools." Information Week.

Lallande, A. (2005). "Adding Up Sarbanes-Oxley Costs; Millions, countless staff hours spent on compliance." Business Matters

Lobur, L. N. a. J. (2006). The Essentials of Computer Organization and Architecture, Jones and Bartlett Publishers.

Lobur, L. N. a. J. (2006). "The Essentials of Computer Organization and Architecture."

Microsystems, S. (2007). "ZFS: The last word in file systems." from http://www.sun.com/2004-0914/feature/.

Peter M. Chen, E. K. L., Garth A. Gibson, Randy H. Katz, David A. Patterson "RAID: High-Performance, Reliable Secondary Storage." ACM Computing Surveys: 62.

Peter M. Chen, E. K. L., Garth A. Gibson, Randy H. Katz, David A. Patterson (1994). "RAID: High-Performance, Reliable Secondary Storage." ACM Computing Surveys **26**(2): 145-185.

Planet, E. N. (2003, January 27th, 2003). "Serial ATA." Retrieved April 15, 2007, from http://networking.webopedia.com/TERM/S/Serial_ATA.html.

Planet, E. N. (2004, September 20th, 2004). "SCSI." Retrieved April 15, 2007, from http://networking.webopedia.com/TERM/S/SCSI.html.

Poelker, C. (2005, December 9, 2005). "RAID explained." Ask The Storage Expert: Questions & Answers Retrieved April 7, 2007, from http://searchstorage.techtarget.com/expert/KnowledgebaseAnswer/0,289625,sid5_gci1126527,00.html.

Randolph Kahn, E. "The Sarbanes-Oxley Act Understanding the Implications for Information and Records Management."

STATS, P. (2004) "Beginners Guides: Installing RAID on Desktop PCs." **Volume**, DOI:

Steiner, N. S. (2007). "Without A Trace: The Missing White House Emails and Violations of the PRA." Citizens For Responsibility and Ethics in Washington.

Sullivan, L. (2006). "Compliance Spending To Reach $28 Billion by 2007." Information Week.

Sun Microsystems, I. (2007). "Solaris 10." ZFS FAQs Retrieved April 15, 2007, from http://www.sun.com/software/solaris/faqs/zfs.xml.

Sunstar Company, I. (2007). "RAID-DP." Retrieved April 15, 2007, from http://www.sunstarco.com/Storage/StoreVault/Software%20&%20Technologies/RAID_DP.htm.

Technick.net. "RAID Technology." Retrieved April 7, 2007, from http://www.technick.net/public/code/cp_dpage.php?aiocp_dp=guide_raid.

Wendy Croix, P. D., freelance writer, university professor, and former retail manager "Quickly Growing Employment Sectors Offer Job Opportunities
Which Careers are Projected to Increase in Demand?" WorldWideLearn
The World's Premier Online Directory of Education **Volume**, DOI: