DIGI-DATA CORPORATION

WHITE PAPER

Keeping Up With RAID

How the Latest RAID Versions Add New Dimensions to Storage Performance, Reliability and Usefulness

BY
Dr. Tomlinson Rauscher,
Chief Operating Officer
Digi-Data Corporation

May 2003

Keeping Up With RAID

As the information architecture paradigm shifted to distributed systems, network based computing and open systems, IT organizations turned to Redundant Arrays of Independent Disks (RAID) storage solutions to centralize and improve storage management.

RAID storage delivered improvements in:

- Reliability
- Performance
- Scalability
- Availability
- Management
- Cost

With the more recent growth of storage intensive applications such as Web, e-commerce, ERP, Email, rich media and high-performance computing (HPC), double-digit growth in storage utilization has become common. In addition, analysts such as IDC and Enterprise Storage Group predict storage to continue to double annually for several years to come.

For these reasons, RAID storage systems have become central to storage planning in corporate and application environments. RAID has become the workhorse within Direct Attached Storage (DAS), Network Attached Storage (NAS) and Storage Area Network (SAN) environments supporting a critical role of storage as a network utility.

In response to this growth and the critical role it plays, RAID has evolved and RAID options have increased, continuing to offer storage managers new flexibility and control over their storage environment.

However, with close to ten RAID levels to choose from, matching the right RAID system and RAID levels to your applications and environment is important in assuring data availability and performance.

RAID Primer

RAID systems provide large amounts of storage by making the data on many small disks readily available to file servers, host computer(s) or the network as a single array. Applications present a storage system with data types ranging from large sequential file transfers as in streaming media to high volumes access to small files as found in many web applications.

Data throughput and availability are dependent upon the architecture of the storage system, how the data is stored within the array (RAID level) and application utilization of the data.

RAID System Architecture

RAID storage systems contain three main components, as illustrated in the accompanying figure:

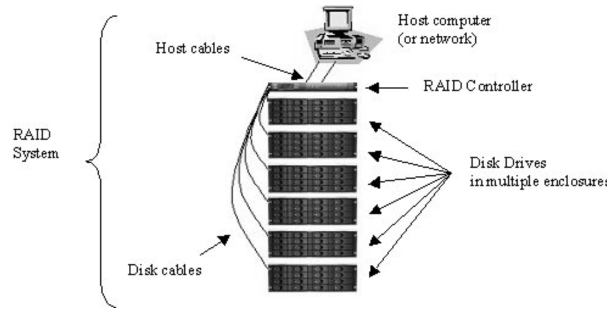

1. <u>The RAID Controller</u>
   The RAID controller is an intelligent electronic device that routes, buffers and manages data flow between the host computer(s) or network and the array of disks. Regardless of whether the RAID storage system is attached to a Storage Area Network (SAN), a server in a Network Attached Storage (NAS) system, or directly attached to one or more servers, the RAID controller makes the array of disks look like one or more virtual disks that are very fast, very large, and very reliable.

   The RAID controller connects to the host computers – either directly to server(s) directly or the network – using a high-speed interface such as Fibre Channel or SCSI. A high performance RAID controller uses two or more parallel data paths to allow host computers to communicate with the RAID system. This enables demanding applications, such as high-speed transaction processing, that would otherwise be impossible or extraordinarily expensive. These connectivity options continue to expand with choices such as iSCSI and Infiniband.

   The controller then connects the drives using similar connections. Most RAID systems today connect to drives using Fibre Channel, SCSI, or ATA interfaces.

   A high performance RAID controller uses parallel data paths to read and write information to the disks in a RAID array. Thus, throughput is increased as it performs the operations of reading and writing information to several disks simultaneously.

   A RAID controller that uses six or more disk channels can address up to several hundred disks and provides scalability to more than 100 terabytes (100,000 gigabytes) in a fully configured RAID system.

Advanced RAID controllers will reconstruct the lost data from a failed disk onto a spare disk, so that the system can survive another disk failure. The most sophisticated RAID systems have 6 or more disk channels, so the RAID controller can stripe data across disks on separate channels, a technique called orthogonal RAID.
These systems provide the ultimate in data availability, because they survive not only failure of a disk but also failure of a disk channel that may contain several disk enclosures.

What is important is that modern RAID storage controllers have made the connectivity between the disk array and the servers or network seamless and transparent; the servers or network do not require changes to software or special drivers to deploy and manage even complex RAID storage systems.

2. <u>The Array of Multiple Disks</u>
The design of the array of disks is an important determinant of performance and data availability in a RAID system. This includes both the disk enclosure and the disks. The disk enclosure provides high-speed access ports to disks, the ability to lock disks for security, and visual feedback on operation. Sophisticated disk enclosures include load sharing power supplies and hot swap components such as fans and power supplies, offer greater expandability through multiple access ports, report status to the RAID controller, and can support from 12 to 16 disks in a cabinet 19 inches wide and 5¼ inches high. The disks in corporate environments typically are Fibre Channel, SCSI, or ATA – each with a variety of capacities, rotation speeds, read/write performance characteristics, data cache, Mean Time Between Failure (MTBF) characteristics and cost per unit of storage.

Limitations in scale and performance are highly dependent upon channel technology. For instance, Fibre Channel can address 125 devices per channel, SCSI can address up to 15 devices per channel, while the lower performing ATA technology can address only a few devices per channel.

3. <u>Cabling</u>
Cables that connect the RAID controller to the array of disks and to host computers or network devices provide critical data and status communication links. Depending on whether the disk array is to be connected to a Storage Area Network, to a server in a Network Attached Storage System, or directly attached to servers, the cabling can range from fiber optic to SCSI.
Fibre Channel has out grown SCSI in popularity because of its high-speed data transfer (2 gigabits per second) over long distances (tens and hundreds of meters) and scalability (can support 125 drives per Fibre Channel), while SCSI cables are limited in capacity (up to 15 drives per SCSI channel) and cable length. In addition, fiber optic cables have high immunity to electrical noise and environmental factors, yet are lightweight, robust, and easy to install.

RAID Levels

RAID controllers can organize data on the disks in several ways to optimize the performance and reliability of the system for different types of applications. In fact, advanced controllers offer the ability to format multiple logical units (LUNs) in a RAID array independently with different RAID levels to accommodate the varying applications and needs of an enterprise.

RAID levels include mirroring and parity checking alternatives to improve data availability. In mirroring, the RAID controller replicates data on multiple disks to create multiple copies of all information. Although this technique is fast and very reliable, it does require twice as many disks for all information. With parity checking, when the RAID controller writes information onto the disks, it also writes redundant information called parity bits. Should a disk fail, this parity information enables the RAID controller to recreate the lost information as it is requested. Since the parity information is spread across multiple disks, only a percentage of disks is used for parity information, thus improving the efficiency of available storage space.

There are several popular RAID levels from which to choose, each with unique capabilities and limitations. The chart below provides a look at each of these popular RAID levels and its intended use.

Common RAID Levels (note that RAID level numbers indicate different types of data layouts, not higher performance or availability)					
RAID Level	**Description**	**Data Reliability**	**Performance**	**Application Strength**	**Cost**
RAID 0 Striping	Data striped across multiple disks	Not true RAID – Any disk failure causes loss of data.	Faster than single disk for reads because data striped across several disks can be read simultaneously. Similar to single disk for small writes.	General	Low, because there is no redundancy
RAID 00* Two Dimensional Striping	Data striped across disks on separate channels <u>and</u> across disks on the same channel	Not true RAID – Any disk failure causes loss of data.	Faster than RAID 0 for reads. Similar to RAID 0 for writes.	General	Moderate. Designed for large arrays
RAID 1 Mirroring	All data copied onto 2 separate disks	Very high. Can withstand selective multiple disk failures	Faster than single disk for reads because data mirrored on 2 disks can be read simultaneously. Similar to single disk for small writes.	General	High. Requires twice as many disks for redundancy
RAID 10 Stripped Mirroring	Two copies of data striped across disks	Very high. Can withstand selective multiple disk failures	Very high. For reads, access is very fast because data is both mirrored and striped (i.e., there are 2 disks from which to read any piece of data, and striping spreads data across more disks)	High data reliability and performance applications such as ERP and image processing.	High. Requires twice as many disks for redundancy
RAID 100* Mirroring with Two Dimensional Striping	Two copies of data striped across disks on separate channels <u>and</u> across disks on the same channel	Very high. Can withstand selective multiple disk failures	Highest of all RAID levels	Communication services, supply chain or enterprise applications which are customer facing.	High. Requires twice as many disks for redundancy
RAID 2	Data bit-striped across disks with error correcting codes on additional disks	Very high Can withstand multiple disk failures	Slowest of all the RAID levels due to bit striping	Slow speed makes RAID 2 unattractive. It is not used commercially.	

RAID 3	Data striped across disks on separate channels with dedicated parity disk	Much higher than single disk. Can withstand single disk failure	Faster than a single disk, owing to parallel disk accesses	Video, prepress, medical imaging, and other large file applications	Low. Requires only one disk per RAID group for redundancy
RAID 30	Data striped across disks on separate channels with dedicated parity disk and across disks on the same channel	Much higher than single disk. Can withstand single disk failure	Faster than RAID 3	Video, prepress, medical imaging, and other large file applications	Moderate. Designed for large arrays
RAID 4	Data block-striped across disks on separate channels with dedicated parity disk	Much higher than single disk. Can withstand single disk failure	Faster than a single disk, owing to parallel disk accesses	Video, prepress, medical imaging, and other large file applications	Low. Requires only one disk per RAID group for redundancy
RAID 5	Data and parity striped across multiple disks	Much higher than single disk. Can withstand single disk failure	High compared to single disk for reads but lower than single disk for writes	OLTP, email, ERP, web, CRM.	Low. Requires only one disk per RAID group for redundancy
RAID 50	Data and parity striped across multiple disks and across disks on the same channel	Much higher than single disk. Can withstand single disk failure	Faster than RAID 5	Transaction processing with high read to write ratio	Moderate. Designed for large arrays

* *Designed by Digi-Data*

Summary

With increasing demands for mass storage performance, capacity, and reliability, many firms are adopting RAID technology to make their servers and networks more flexible and responsive to changing environments, while supporting a wide variety of applications such as digital imaging, prepress, on-line transaction processing, data warehousing, and file servers.

To optimize your environment using RAID, start by listing and prioritizing your current and anticipated needs for performance, capacity, reliability and scalability based on the applications and number of servers and users you intend to support. Determine whether your RAID storage system will be directly attached to one or more servers, attached to the network via NAS, or attached using a SAN. Then carefully match these needs with the right RAID levels and capacity for your environment. Working with storage integrators and resellers who understand storage capacity planning and networking are the best people to help you validate these estimates.

You also want to select and deploy a RAID storage system that has the flexibility and versatility you need to respond to unforeseen needs as they arise. This not only means selecting a system that can accommodate capacity expansion and offers multiple RAID levels, but one that can scale performance and capacity proportionally, and offer multiple RAID levels across different LUNs within the storage system to optimize the use of all the storage you own.