




What is RAID--BASICS?

Mylex RAID Primer

A simple guide to understanding RAID

Let's look at a hard disk...

- ◆ Several platters stacked on top of each other with a little space in between.
- ◆ One to n platters with a magnetic head on each side all connected rigidly to one spindle.




- ◆ Hard disk drives store the programs and files with which you work.
- ◆ The most important program is the Operating System which resides on the hard disk drive.
- ◆ Every computer ships with at least one hard disk drive

A Single Drive - As RAID sees it



- ◆ **A single physical drive or Multiple physical drives**
 - JBOD (**J**ust A **B**unch **O**f **D**isks)
- ◆ **Used as a single system drive (one physical drive)**
- ◆ **Used as multiple system drives (more than one physical drive) attached through one or more SCSI or Fibre channels**
- ◆ **The RAID controller acts as an intelligent SCSI I/O port**
 - Benefits include ECC memory, battery backup and other fault tolerant features not available in non-intelligent SCSI I/O port controllers
- ◆ **RAID stands for:**
 - **R**edundant **A**rray of **I**nexpensive/**I**ndependent **D**isks

Single Disk (No RAID)




Data Stripes



Disk Drive

eXtremeRAID™ 1100



Example: Data is written to one disk as it is read from system memory.

RAID Levels



- ◆ **RAID 0**
 - Striping. Data is striped across multiple disks.
- ◆ **RAID 1**
 - Mirroring. Data is written identically to a multiple of two disks simultaneously.
- ◆ **RAID 3**
 - Striping with Parity. Data is striped across multiple disks and parity is written to one additional disk.
- ◆ **RAID 5**
 - Data and parity are written across multiple disks simultaneously.


RAID Applications



- ◆ **Transfer Rate Intensive Applications:** *(Typically RAID 0 environments)*
- ◆ **RAID striping is ideal for transfer rate-intensive environments**
- ◆ **A transfer rate-intensive environment consists of :**
 - Applications that require a large amount of data to be processed in a fixed amount of time
 - Video playback and video editing are typical transfer rate-intensive environments
 - Photo processing, manipulation and rendering
- ◆ **Request Rate Intensive Applications:** *(Typically RAID 5 environments)*
- ◆ **RAID is used in highly multi-tasking, request rate-intensive environments**
- ◆ **A request rate-intensive environment consists of:**
 - databases, file/web servers: -- high number of random smaller requests.
- ◆ **A RAID drive can be configured to process each request within a stripe, allowing multiple requests to be processed in parallel.**


RAID 0 - Striping

- ◆ RAID 0 uses disk striping.
- ◆ Multiple hard disks used to form a logical drive.
- ◆ Hard disks are striped using the stripe size specified during configuration.
- ◆ A request less than the stripe size is sent to one hard disk. A request greater than the stripe size is divided into multiple requests and sent to multiple disks, in parallel.
- ◆ Typically used in data rate intensive applications (video editing) three-drive RAID with a predefined stripe size of 32K. A 72K request



Example: Consider a three-drive RAID with a predefined stripe size of 32K. A 72K request would be broken up into three chunks of 32K, 32K, and 8K, with the first 32K going to one drive, the next 32K going to the next drive, and the last 8K going to the third drive. Each of these requests would be sent in parallel to the three drives to be serviced simultaneously.

RAID Level 0 - Disk Striping






Data Stripes



Disk Drive


eXtremeRAID™ 1100



Example: Data split in equal chunks and is written to multiple disks in parallel. The data transfer rate is the aggregate sum of all the drives available in the array minus any overhead.

RAID 1 - Mirroring


- ◆ RAID 1 provides the highest level of fault tolerance to a single drive
 - duplicates each I/O request to a secondary mirror hard drive.
- ◆ RAID 1 requires a multiple of two drives -- the primary set and the mirror set.
- ◆ Mirroring is suited for any application where protection from drive failure is required.



Mirroring allows two drives to mirror on the same channel or "duplex" across two channels/two host adapters.

- ◆ The mirror drive must be equal or greater in capacity than the primary drive
- ◆ In case of a drive failure, all I/O requests are re-directed only to the remaining good drive
- ◆ Used to store only the important data from applications
- ◆ Also used to protect the main drive where applications are installed
- ◆ Provides immediate disaster recovery when a drive fails.

RAID Level 1 - Disk Mirroring





Data Stripes



Disk Drive

eXtremeRAID™ 1100




Example: Data is copied and written to an identical drive in parallel. This provides a redundant configuration for disaster recovery.

RAID 3 and RAID 5




- ◆ RAID 3 and 5 stripes data across multiple drives and adds parity for fault tolerance.
- ◆ RAID 3 writes all parity information to one single drive.
- ◆ Mylex RAID 5 controllers implement “rotational parity” for increased drive reliability and reduced “single point of failure”.
- ◆ RAID 5 is the predominant highly available, fault tolerant RAID level used in servers today.
- ◆ RAID 5 applications include OLTP (online transaction processing), file server and web server applications requiring 24/7 up-time.



RAID 5 Array


Example: Consider a three-drive RAID 5 set. Each block or “stripe” of data is split across as many drives (n-1) as will fit. The remaining drive contains the parity information to re-construct the previously written striped data in case of a failure. The parity block is rotated amongst all the drives so that one single drive does not get accessed more than any other drive.

RAID 5 - Disk Striping w/ rotational Parity




eXtremeRAID™ 1100





Data Stripes



Parity Stripes




Disk Drive



Example: Data is split in equal chunks and is written to multiple disks in parallel. The remaining drive contains the parity information to reconstruct the previously written striped data in case of a failure. The parity block is rotated amongst all the drives so that one single drive does not get accessed more than any other drive.

RAID 00, 10, 30, 50

- ◆ Provides Multi-layer RAID configurations
- ◆ Ideal for TPCC configurations
- ◆ Multiple RAID Level 0, 1, 3, or 5 Arrays can be “striped” together to create RAID Levels 00, 10, 30, or 50
 - Provides more capacity
 - Often referred to as ‘Spanning’
- ◆ Provides better I/O load balancing



Example: Three RAID 5 arrays are “striped “ together for added capacity and as a means to add additional drives to a large array.

RAID Controllers



- ◆ Low Cost RAID for Desktops and Entry-level Servers
 - 0 or 1 channel PCI-to-Ultra2 SCSI RAID
 - examples: Mylex AcceleRAID Series
- ◆ High Performance RAID for Mid-level to High-end Servers
 - Up to 3 channel PCI-to-Ultra SCSI & PCI-to-Ultra2 SCSI RAID; up to 128MB cache; ECC (some models)
 - examples: Mylex eXtremeRAID Series
- ◆ External RAID for enterprise level servers
 - 2 channel Fibre to 4 Ultra2 SCSI RAID; 256MB mirrored cache
 - examples: Mylex DAC-FL Series

RAID Advantages!



- ◆ RAID provides an environment of highly reliable, fault tolerant disk drive sub-systems
- ◆ RAID adds life to disk drives as the controller manages the I/O load equally amongst all the drives in the Array reducing the risk of “single-point-of-failure”
- ◆ A “successful” RAID environment includes reliable robust disk drives
- ◆ Each RAID environment will include **multiple disk drives**
- ◆ Mylex PCI RAID controllers can attach as many as 45 drives per RAID controller, up to 16 controllers per system for an astonishing -- ***720 disk drives per*** system configuration
- ◆ Mylex external RAID controllers can attach as many as 90 drives per storage subsystem and can reside as part of a SAN (Storage Area Network)
- ◆ Capacities range from terabytes to petabytes!

Where can I use RAID?



- ◆ LANs/WANs
- ◆ SANs
- ◆ Clustering environments
- ◆ Mission critical installations
- ◆ News centers
- ◆ Internet News Servers
- ◆ Enterprise Servers
- ◆ Performance Desktop Systems
- ◆ PC Workstations
- ◆ Workgroup/File Servers
- ◆ E-Mail Servers
- ◆ Intranet/Web Servers
- ◆ Application Servers

MYLEX RAID



- ◆ Broadest line of RAID controllers in the world
 - Low Cost RAID for Desktops and Entry-level Servers
 - High Performance RAID for Mid-level to High-end Servers
 - External RAID for Enterprise-class Servers
- ◆ Strategic relationships with industry leaders
 - Intel, Microsoft, System vendors, Drive vendors & Software suppliers
- ◆ Mylex has the world's fastest RAID Controller
- ◆ Partner with 7 of the top 10 server OEMs
- ◆ Leading RAID supplier in the channel
- ◆ Over 1,000,000 Mylex RAID controllers in the field
- ◆ Mylex rated #1 in non-captive market share by IDC, DataQuest & Disktrend ---- 1998
- ◆ **Mylex is synonymous with RAID**

Mylex RAID for All Server Solutions



- ◆ **AcceleRAID™ Series: Cost Effective PCI RAID Solutions for:**
 - Entry-level Servers
 - “Win/tel”-based Workstations
 - High-end Desktop Systems

- ◆ **eXtremeRAID™ Series: High Performance PCI RAID Solutions for:**
 - Mid-level Servers
 - High-end Servers

- ◆ **DAC Series: External RAID Solutions for:**
 - Enterprise-class Servers
 - Clustering and High Data Availability Applications
 - Storage Area **Networks**

AcceleRAID™ Technology



RISC -i960
Ultra2 SCSI
Low Cost



- ◆ 0 or 1 Ultra2 SCSI Channels
- ◆ Up to 4MB Flash EEPROM
- ◆ Up to 128MB MB EDO RAM w/ECC
- ◆ Up to 1 int. + 1 ext. SCSI connectors
- ◆ PCI 2.1 compliance
- ◆ PCI Hot Plug
- ◆ RAID 0, 1, 0+1, 3, 5, 10, 30, 50, JBOD
- ◆ Also used in Systems with a PCI slot dedicated for add-in RAID controllers, providing 1 SCSI channel in addition to the number of channels on the Motherboard (AcceleRAID™ 250)

eXtremeRAID™ Technology



**233 MHz SA 110,
Ultra2 SCSI,
64-bit PCI**



- ◆ 2 or 3 Ultra2 SCSI Channels
- ◆ Up to 4MB Flash EEPROM
- ◆ 32 or 64 MB parity SDRAM + BBU module
- ◆ Up to 3 internal and 3 external SCSI connectors
- ◆ PCI 2.1 compliance
- ◆ Dual Address Cycle
- ◆ PCI Hot Plug
- ◆ > 6000 Random I/O's
- ◆ > 75 MB/second sustained data transfer rate

DAC-FL: Technology



**Two RISC Processors,
2 Fibre Channel loops,
Ultra2 SCSI**



- ◆ Up to 256 MB Data Cache
- ◆ Dual PCI Bus design
- ◆ 5 1/4" form factor
- ◆ 2 Copper FC-AL host channels
- ◆ Burst cache transfer rates: Up to 97 MB/s per host channel
- ◆ 4 Ultra2 SCSI LVD channels
- ◆ 15 devices per Ultra2 SCSI channel - 60 devices/controller
- ◆ Up to 1.08TB per controller
- ◆ Sustained disk transfer rates: over 52MB/s

DAC-FF: Technology



**Two RISC Processors,
2 Fibre Channel to host,
4 Fibre Channels to drives**



- ◆ Up to 256 MB Data Cache
- ◆ Dual PCI Bus design
- ◆ 5 1/4" form factor
- ◆ 2 Copper FC-AL host channels
- ◆ Burst cache transfer rates: Up to 189 MB/s
- ◆ 4 Fibre channels to drives
- ◆ 125 devices per Fibre channel - 500 devices/controller
- ◆ Up to 9.0TB per controller