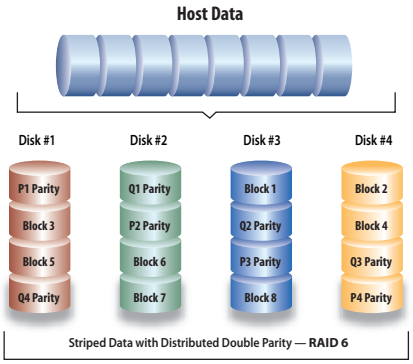

Figure 1: Two sets of parity data, P & Q, are striped across the disks. RAID 6 safeguards data against a second drive failure.

RAID 6: Greater Fault Tolerance

- **Higher data availability**
  Data is safeguarded against up to two consecutive drive failures
- **2-Drive Parity**
  Data from two failed drives can be rebuilt with assured data accessibility
- **RAID Protection in Degraded Mode**
  Data is protected against a single drive failure during rebuilds

RAID 6 Trade Offs

- **Reduced write performance**
  Second parity calculation causes system to work harder processing write transactions
- **Longer rebuild times**
  Twice the parity is used to reconstruct data
- **Minimum four drives required**
  Two of four drives exclusively dedicated to storing parity (N-2)
- **Higher system cost**
- **Lower available capacity**

AMCC STORAGE

RAID 6

Double-parity RAID, commonly known as RAID 6, safeguards against data loss during rebuild mode by allowing up to two consecutive drive failures

What is RAID 6

In a RAID 5 array, data is striped across all drives in the array. Parity information is rotated and stored across all the disks. If an individual drive fails, the surviving array operates in degraded mode until the failed drive is replaced and its data is rebuilt from the parity information retained on the surviving disks. RAID 5 arrays are vulnerable in degraded mode because all data will be lost in the unlikely event of a second drive failure during the rebuild. Rebuild times are getting increasingly longer due to today's increased hard disk capacity. Longer rebuild times widen the window of likelihood that a second drive will fail which will result in data loss.

RAID 6 eliminates the risk of data loss if a second hard disk drive fails while the RAID array is rebuilding. In a RAID 6 enabled system, a second set of parity is calculated, written, and rotated across all the drives. This second parity calculation provides significantly more robust fault tolerance and allows the array to survive up to two consecutive drive failures without losing data. A RAID 6 implementation is diagramed above (Figure 1).

RAID 6 Implementation Considerations

Performance

RAID 5 write performance is influenced by the number of disk accesses that are required during the write process. While there is no adverse effect on RAID 5 read performance, write performance drops by almost 50% between RAID 0 (data striping across multiple drives) and RAID 5 (data striping across multiple drives with rotating parity calculation)[1]. The effect on overall performance will depend on the ratio between reads and writes for a given application; more writes mean lower performance.

RAID 6 requires a second set of parity calculations to protect data against a second drive failure. This additional data-handling step adversely affects performance. Independent performance benchmarks show that a RAID controller can suffer a 20% drop in overall performance in RAID 6 compared to a RAID 5 implementation.[2] As with RAID 5, read performance is unaffected.

Capacity

RAID 5 implementations require a minimum of 3 drives and have the storage capacity of N-1 drives because the equivalent capacity of one drive is exclusively dedicated to holding parity data. For example in a 4 drive, 200 gigabyte per drive array, the total available storage capacity is 600 gigabytes out of 800 gigabytes.

RAID 6 implementations require a minimum of 4 drives and have the storage capacity of N-2 drives. The total available storage capacity, using the same example, is 400 gigabytes out of 800 gigabytes. The percentage of usable system capacity is greater in larger RAID 5 and RAID 6 configurations. In a typical 8-drive SATA RAID array, 25% of the total drive capacity will be used for RAID 6 parity, compared to 12.5% of a RAID 5 array (see figure 2).

Impact of Parity Calculation on Arrays

	Array capacity used for parity (%)		Storage efficiency (%)	
# Drives	RAID 5	RAID 6	RAID 5	RAID 6
3	33.3	N/A	66.7	N/A
4	25.0	50.0	75.0	50.0
8	12.5	25.0	87.5	75.0

Figure 2: Usable system capacity is greater in larger RAID 5 and RAID 6 systems. RAID 6 uses more capacity for additional parity storage.

Summary

RAID 6 provides higher fault tolerance compared to RAID 5 arrays. By assuring data availability following a second drive failure, RAID 6 provides additional protection during degraded mode. RAID 6 does not come without costs, however. Overall RAID 6 system performance can suffer a 20% drop compared to RAID 5; write performance is also adversely affected due to additional parity calculations on writes. Additionally, RAID 6 requires the equivalent capacity of two drives in the array to be dedicated to storing only parity information. At current market pricing, using 400 gigabyte drives, an 8 drive RAID 6 array would deliver 2.4 terabytes of actual data storage against a total array capacity of 3.2 terabytes, an additional cost to the system of approximately $300.

Avoiding a 2nd Drive Failure

RAID 5 provides robust redundancy during normal operation. RAID 6 further protects the RAID array against data loss during degraded mode by allowing up to two drives to fail during this vulnerable stage.

It is possible, however, to insure against the vulnerability of the system in degraded mode without incurring the costs associated with RAID 6. In general, the faster the rebuild is, the lower the risk of a second drive failure during rebuild. Building RAID 5 systems with reduced rebuild times in mind will minimize the chances of a second drive failure.

There are several ways of doing this:

1. Hot sparing with automatic rebuild. This does not speed up the rebuild, but does remove the time delay between drive failures and drive replacement. Multiple arrays on a single controller can share a single hot spare for automatic rebuild.
2. Set the rebuild priority to highest level. This will slow the application down during rebuilds but will minimize the exposure time.
3. Minimize the number of drives per array in line with the storage requirements. The greater the number of drives in a single array, the higher the probability of a second drive failure.
4. The higher the MTBF (Mean Time Between Failure) of the drive, the lower the probability of a drive failure to begin with. Always look for the highest rated drives for your RAID 5 array.
5. Use a higher number of smaller drives. The bigger the drive the longer the re-build time. Smaller drives will shorten the drive re-build time. In addition, smaller capacity drives tend to be significantly cheaper so the cost savings may cover the cost of a hot spare, as shown in the table below:

Storage Requirement	Capacity Drives	# Drives for RAID 5	Drive Cost	Controller Cost	Total Cost	Hot Spare Cost	Total Cost with Hot Spare
1.2 terabytes	400G	4	$312	$300	$1,548	$312	$1,860
1.2 terabytes	200G	7	$109	$490	$1,253	$109	$1,362

Prices based on **www.pricewatch.com** 2/16/05

[1] *Based on performance benchmarks completed on AMCC's 9000 series RAID controllers*

[2] *http://www.tomshardware.com/storage/20041227/areca-raid6-07.html*

AMCC

Sales Offices for 3ware Products:

USA: +1-877-883-9273
+1-408-523-1000
Europe: +00-800-3927-3000
Asia/Pacific: +65-6826-3381
Japan: +81-3-6717-4458

3waresales@amcc.com
www.3ware.com
www.amcc.com