North River Solutions
7 Colby Court
Bedford, NH 03110
Phone: (888) 761-4411
Fax: (646) 349-3652

# Calculating the True Reliability of RAID Systems

The Brookings Institute estimates that 85% of the digital assets of the average corporation reside on servers, desktops and laptops deployed outside of the data center. This shift in the location of vital records and critical applications is forcing a change in data protection strategies everywhere.

Disk arrays offer many desirable performance and capacity characteristics. They also have a reputation as highly reliable technology, but are they as good as advertised? This article explores the reliability claims of disk array technology.

There are several types of disk arrays, each with its own set of performance, reliability and capacity characteristics. Referred to as RAID *Levels*, these configurations differ in the way in which they store data on magnetic disks. Some RAID Levels, such as 2 and 4, do not offer any advantages and have never been mass-produced. RAID Level 3 has effectively been rendered obsolete by price/performance enhancements to RAID Level 5. For all practical purposes, there are four types of RAID systems in use today:

- RAID Level 0, simple disk stripping which yields a high performance/high capacity system.
- RAID Level 1, a mirroring strategy where two copies of information are kept on separate but logically connected disk drives.
- RAID Level 5, a configuration where mathematical techniques are used to create a software representation of stored data. This strategy allows the reconstruction of missing data in the event of a disk drive failure.
- RAID Level 5+, enhancements to RAID 5 that improve the capacity, performance, reliability or a combination of these factors.

Some combinations of these levels are in use such as RAID Level 10, which is a combination of RAID Levels 1 and 0.

The important thing to remember is that no single RAID system meets the requirements of all applications. Some analysis is required to match applications to the appropriate type of disk array and it is common to have multiple types of array groups co-existing on the same server since the needs of the resident applications are different.

### *Price Breakthroughs*

Historically, cost has been the over-riding consideration in selecting storage. In light of rapidly dropping prices this issue in no longer paramount and other factors such as reliability and performance now take top consideration.

But are the claims of dramatically improved reliability for RAID systems true? To understand this issue it is necessary to become acquainted with the terms shown in Table 1.

**Table 1 – Terms Used to Define Reliability**

1) *Mean Time Between Failure* (MTBF). Perhaps the best known of the reliability statistics, most people believe this figure indicates the number of hours a drive will run before it fails. In reality, MTBF is a statistical calculation that projects the *average* life expectancy of a *typical disk* in a *large population* of drives. MTBF does not indicate how long a disk drive will run before it fails, just the probability of failure.

2) *Annualized Failure Rate* (AFR). Although AFR is not as well-known a measure of reliability, it does have its uses. AFR is calculated from the observed failure rate of a device over time. Disk resellers are particularly sensitive to this issue since this statistic influences their spare parts strategy.

3) *Mean Time Between Data Loss* (MTBDL). MTBDL is a statistical figure that attempts to predict how long an array group will operate before suffering a catastrophic failure and data loss. Data loss occurs when information is written to the failed disk array subsequent to its last backup. RAID technology allows some disk array groups to survive the failure of one or more disks making MTBF less significant. MTBDL is a more meaningful measure of reliability than MTBF or AFR since it considers the performance of the array, not just a sub-component.

4) *Mean Time To Repair* (MTTR). MTTR is the time it takes to repair a failed part. MTTR figures assume that no time elapses between the failure and the beginning of repair operations. That is to say, the appropriate person becomes aware of a bad component the instant it fails, has the necessary replacement parts available, and begins the repair process *immediately*. Unfortunately, this is not usually a valid assumption. As disk capacity increases, the time to complete the rebuild process grows in importance since it extends the time that the total array is vulnerable to a catastrophic failure.

5) Redundant Arrays of inexpensive Disks (RAID). The term RAID is applied to a collection of techniques that store data on a group of logically bound disks in order to provide high throughput and, in most cases, parity protection for high reliability.

## *Calculating the MTBF of a Single Disk*

How does a disk manufacturer determine that the MTBF for a brand new drive should be 100,000 hours (11.4 years)?

Initially the manufacturer does a statistical calculation that projects an MTBF number. The method used to calculate MTBF is conceptually quite simple. During the development cycle, the manufacturer sets up a testing facility (commonly called a "test bed") and collects data on the new device. For example, the vendor might run 1,000 new disks in parallel for 13 weeks. At the end of this three-month period, the vendor has accumulated 2,184,000 hours of "disk in service" time (1,000 disks x 13 weeks x 7 days per week x 24 hours per day). If the vendor experienced 21 failures out of the 1,000 disks tested, the calculated MTBR for the drive in question would be 2,184,000 hours divided by 21, or a Mean Time Between Failure of 104,000 hours.

An alternative approach, known as the Mil-Spec method, is sometimes used. To calculate reliability via the Mil-Spec method, the MTBR figures of various components are added and then averaged according to a well-defined process. This approach is not commonly used by disk drive manufacturers, but is widely used by printed circuit board makers and chip designers. Historically, the Mil-Spec method yields very conservative MTBF numbers.

The third method of computing MTBF relies on customer experience. Once a device has been installed for a meaningful period of time (usually one year), the vendor has enough data to measure the observed reliability of a particular device. The method used to compute this observed MTBF number is also statistical in nature. In this case, MTBF is calculated by dividing the total number of operating hours in a year (8,760) by the Annualized Failure Rate (AFR). This method provides a more realistic MTBF estimate since it is based on demonstrated rather than projected failure rates.

### *Disk Stripping: Higher Performance - Lower Reliability*

Disk arrays make multiple individual disks look like one large drive. Each RAID level differs in the way it organizes data and performs under various workloads. One technique is known as disk striping. Disk striping divides a data file into several component segments and writes this data in parallel across several logically bound disks forming a RAID Level 0 array.

Disk striping offers the same capacity and significantly higher performance than that delivered by a group of individual disks. The combined performance and capacity benefits of RAID 0 explains the interest in this technique. However, since the information is spread evenly over the entire family of drives, the failure of one disk causes the entire array to fail with no way to restore the files that were written to the array after its last back-up.

Since all the data stored on a RAID 0 array is lost if a member disk fails, it is logical to ask how the reliability of a disk array compares with that of a single disk drive? Is the MTBF rating of a RAID 0 group greater or less than that of its component disk drives?

To calculate the reliability of a RAID 0 array, the MTBF of any one disk is divided by the total number of drives in the group, since each drive in the array has an equal probability of failure.

$$\textbf{MTBF of the array} = \frac{\textbf{MTBF individual disk}}{\textbf{N}}$$

**(where N = the number of drives in the array)**

For example, assume that there are 5 disk drives, each with an MTBF of 100,000 hours (11.4 years). Based on this figure, the MTBF of a five-drive disk array is:

$$\textbf{MTBF of the array} = \frac{\textbf{100.000 hrs}}{\textbf{5}} = \textbf{20,000 hrs}$$

As more drives are added to this array the overall MTBF drops. A 20-drive array would have an MTBF of about 5,000 hours (100,000 hours divided by 20 drives). This is a far cry from the 100,000-hour MTBF figure cited earlier for each of the array's component disk drives.

What can be done to increase the reliability of disk array, and is MTBF the most appropriate measure of array reliability?

## *The Reliability of RAID 5*

RAID levels 1-5 use various parity checking techniques to overcome the low reliability of disk arrays. If a disk drive fails, the parity information is used to calculate the lost information. This is similar to the way error-correcting memory uses parity bits to correct corrupted entries. RAID 5 distributes both data and parity information evenly among its drive group. If a drive in a RAID 5 array fails, the parity enables the array group to continue operating and simultaneously perform an on-line rebuild of the failed drive once it is replaced. All this occurs without any loss of data or downtime although there is some storage overhead due to the need to store the parity data.

How is the reliability of such parity-protected arrays measured? Since the failure of a single drive is compensated for by the operation of the array, MTBF ratings may not provide useful information.

In parity-based arrays, Mean Time Between Data Loss is a more appropriate measurement since MTBDL takes into account the probability that the array group will suffer the loss of a second drive before the first drive can be repaired. MTBDL is calculated as follows:

$$\textbf{MTBDL of the array} = \frac{\textbf{MTBF of the array}}{\textbf{Probability of a 2}^{\textbf{nd}}\textbf{ failure}}$$

Using a five-drive array with an MTBF of 20,000 hours, calculate the MTBDL as follows:

$$\textbf{MTBDL of the array} = \frac{\textbf{20,000 hrs}}{\textbf{Probability of a 2}^{\textbf{nd}}\textbf{ failure}}$$

To calculate the probability of a second disk failing before the first can be restored to full operation take the disk drive's MTTR and dividing it by the MTB of the disabled array.

$$\textbf{Probability of a 2}^{\textbf{nd}}\textbf{ failure} = \frac{\textbf{MTTR of a disk drive}}{\textbf{MTBF of the disabled array}}$$

To calculate the MTBF of the disabled array, divide the MTBF of a single disk by one less drive.

$$\textbf{MTBF of the disabled array} = \frac{\textbf{MTBF individual disk}}{\textbf{(N-1) drives}} = \frac{\textbf{100,000 hrs}}{\textbf{4 drives}} = \textbf{25,000 hrs}$$

In choosing an MTTR figure it is prudent to consider a worst case scenario. If a failure occurs at the beginning of a holiday weekend, it might be 48 hours before the problem is detected, 3 hours until a technician can respond with the appropriate parts, and 30 minutes or more before the actual repair is completed. If this was an unprotected array (ex. RAID 0 array), additional time would be required to locate, mount and restore data from the last backup tape onto the array. From a manager's perspective, the MTTR would equal the sum of all these activities, or 51.5 hours.

In this example, the probability of the $2^{nd}$ failure would be:

$$\textbf{Probability of 2}^{\textbf{nd}}\textbf{ failure} = \frac{\textbf{MTTR of the damaged drive}}{\textbf{MTBF of the disabled array}} = \frac{\textbf{1.5 hrs}}{\textbf{25,000 hrs}} = \textbf{0.0021}$$

This means that for every 10,000 *single drive failures*, the likelihood is that you would simultaneously experience the loss of a second drive only 21 times. An event that is not very likely!

Having determined the probability of a $2^{nd}$ failure, the MTBDL of the five-drive RAID 5 array can be calculated as:

$$\textbf{MTBDL of the array} = \frac{\textbf{20,000 hrs}}{\textbf{0.0021}} = \textbf{9,523,810 hrs or 1,087 yrs}$$[1]

As this calculation shows, there is a considerable difference between the reliability of a parity-protected RAID 5 configuration and that of a RAID 0 array. The RAID 5 array is also significantly more reliable than any individual disk drive.

---

[1] The formula for calculating MTBDL can also be expressed as:

$$\textbf{MTBDL of the array} = \frac{\textbf{(MTBDL of the disk)}^2}{\textbf{(N)x(N)x(MTTR)}}$$ (where N = the # of drives in the original array)

## *Enhanced RAID 5*

If there was a way of reducing the time it took to return a damaged disk to service, the MTTR of a failed disk would drop, and the MTBDL of the array would increase proportionately. One way of reducing MTTR would be to minimize the waiting period between drive failure and its replacement.

Some disk array products achieve this by supporting a global spare option which can eliminate most of the waiting period. A global spare is an extra disk (or several disks) mounted in the array, but not assigned to any of the array groups. Since the spare disk is not used for data storage, it is available to support any array. If a disk fails, the array automatically begins to rebuild data and parity on the global spare, speeding the return of the array to fully protected status.

To determine the effect that a global spare would have on the MTBDL of the array, the probability a second failure is recalculated using a revised MTTR figure of 1.25 hours. In this example, the MTTR consists of only the time it takes to do an online rebuild. Since the spare is already formatted and mounted in the unit, no time is required for detection, response, and repair. This 1.25 hour figure is in sharp contrast to the 51.5 hour figure used earlier. Thus:

$$\textbf{Probability of 2}^{\textbf{nd}}\textbf{ failure} = \frac{\textbf{MTTR of the damaged drive}}{\textbf{MTBF of the disabled array}} = \frac{\textbf{1.25 hrs}}{\textbf{25,000 hrs}} = \textbf{0.00005}$$

Using the previous figure of 20,000 hours, the MTBDL of the array can be calculated.

$$\textbf{MTBDL of the array} = \frac{\textbf{20,000 hrs}}{\textbf{0.00005}} = \textbf{4000,000,000 hrs or 45,662 yrs!}$$

By narrowing the time in which the array could experience a second drive failure, we can dramatically increase the overall reliability of the entire array.[2]

## *Conclusion*

Disk arrays represent an effective way of delivering increased performance to many applications and represent a cost effective way of adding online storage capacity. However, reliability must always be a consideration when formulating storage strategy.

Even though the MTBF ratings of disk drives are improving, some failures are inevitable. Multiple disk drives configured as RAID disk array groups can provide levels of data protection and reliability previously unavailable to MIS managers at an affordable price. Parity-protected disk array systems can be constructed that are

[2] It is important to realize that these calculations only deal with disk drive reliability. A similar set of calculations are required for the power systems, bus channels, system processor, etc., in order to determine the MTBDL of the entire array subsystem.

slightly more expensive than individual disks, but are nearly one hundred times more reliable.[3]

TABLE 2: Comparative Reliability Rating of Various Storage Configurations

| **Configuration** | **#Drives** | **Reliability Criteria** | **Reliability Rating** |
|---|---|---|---|
| Individual disks | 1 | MTBF | 11.4 years |
| RAID 0 | 5 | MTBF | 2.2 years |
| RAID 5 | 5 | MTBDL | 1,087 years |
| RAID 5 w/global spare | 5 | MTBDL | 45,662 years |

Table 2 shows that a 5-drive RAID 5 system, equipped with a global spare is 4,000 times more reliable than a single disk drive and 20,000 times more reliable than a RAID 0 array. Even more reliable systems can be constructed if ECC Hamming codes or dual parity-drive schemes are used.

RAID 5 arrays offer several key benefits but the most important is that a RAID 5 array continues to operate despite the failure of a disk drive. This *continuous operation* capability is particularly important in mission-critical situations where the cost of computer "downtime" is enormous. And, since rebuilds are done online from the data on the remaining set of active disks in the group, no time is lost searching for the last backup which might exist offsite on magnetic tapes. Tapes that could be lost or stolen and most likely have a reliability issue all their own. Finally, even in the event of a drive failure, *all* information on the array is restored, including any new data generated while the failed drive was still in place.

No storage system, not even an advance RAID unit, can completely eliminate the danger of data loss. User error and accidental file detection still account for most problems. For these and other reasons, frequent backups and periodic archiving are essential. By combining conservative operational procedures with disk array technology, a storage strategy can be built that offers the best defense against catastrophic disk failures.

***About the Author***

Don Byrne has spent over 30 years in the High Tech industry with concentrations in storage technology, business continuity, and enterprise software. He has worked as a venture capitalist, a corporate development executive handling mergers and acquisitions, a *"turn around"* CEO, and is the cofounder of five companies. His latest, North River Solutions, offers consulting and business development services. HE holds degrees in Mathematics, Philosophy and International Marketing. He is completing his PhD in Organizational Behavior and is the author of a forthcoming book on business continuity due out in winter of 2007.

[3] Hardware-based disk arrays range in price from between six to eleven dollars per megabyte. Individual disks on the other hand, cost between 2 and 4 dollars per megabyte.