

July 29, 2008

RE: Hard Drive reliability and RAID 0

By Thomas Hoops



Summary: Hard drive makers tend to obscure useful hard drive reliability information. This paper will attempt to clear the air as this information relates to RAID 0 drive arrays.

With the advent of the RAIDFrame, High-Rely entered a unique world of which there was little comparison. The RAIDFrame currently offers RAIDPacs in either a RAID 5 configuration (recommended) or RAID 0. High-Rely's FirstRAID G2 and Tandem DXR also provide RAID 0 functionality but are not offered in that configuration from the factory. A RAID 0 array complicates the issue of drive reliability because the MTBF (Mean Time Between Failure) rates of the component hard drives are divided.

For example, a common MTBF given from drive manufacturers is 250,000 hours (about 28.5 years). Few of us have lived to see a hard drive live that life. Either MTBF is meaningless or it's just computed in a way that is useless. In either case, it seems that another source of reliability data is needed.

We sought information from a company who is most likely one of the world's largest users of hard drives in their contemporary form – Google. Google was nice enough to openly publish a document called “Failure trends in large disk drive populations” by Eduardo Pinheiro, Wolf-Dietrich Weber and Luiz Andr’e Barroso of Google, Inc – an ideal source of baseline information for this document.

Because a RAID 0 array consists of individual drives, it is necessary to know what to expect of

the individual drive. Google's test case for this paper came from the experience of 100,000 hard drives in use from December of 2005 to August of 2006. They acknowledge the many variances between makes, models, hosts and use. Nonetheless, their data is extrapolated to a net failure rate – a baseline that should be very useful.

In their report, they state that a failure rate of 1.7% was experienced through the first year of operation. It is important to understand that the nature of these failures was diverse and was not always a hardware failure of the drive itself. Google accepts systemic compatibility issues and includes those in this figure. For example, a drive may have problems working reliably in one server and after three months is pulled and installed in another server for three years without trouble.

Similar systemic compatibility issues are seen in High-Rely to server configurations as well. So, this is a good realistic base number to work with. The data presented below applies to ALL RAID 0 systems, not just High-Rely. These statistics are inescapable for any manufacturer of RAID 0 systems by the RAID 0 definition mathematically.

The table (1) below condenses Google's results:

Years of Operation	Percentage Failure
One	1.7%
Two	8.6%
Three	17.3%

Table 1 – Google Hard Drive Failure Statistics

Thus, in three years of 24/7 server operation, there was an average probability of 17.3% that data would be lost or corrupted on any particular hard drive.

These are significant numbers. In computing this probably for a RAID 0 array, the following formula is necessary:

$$\%Probability\ of\ Data\ Loss = 100 \times (1 - (1 - \%SDF/100)^{ND})$$

Where, %SDF is the percentage of Single Drive Failure and ND is the Number of Drives in the RAID 0 array.

Hence, in a RAID 0 array of four drives, the probabilities for data loss/corruption is as follows for the first year:

$$100 \times (1 - (1 - 1.7/100)^4) = 6.6\%$$

Hence the following table (2) details three years of a four drive RAID 0 array:

Year of Operation	% of Failure
One	6.6%
Two	30.2%
Three	53.2%

Table 2 – 4 drive RAID 0 Array failure Statistics (FirstRAID G2)

By the end of the third year of 24/7 operation, there is roughly a 50/50 chance that your four drive RAID 0 array will have failed.

The High-Rely FirstRAID G2 is capable of providing a four drive RAID 0 array. The

RAIDFrame RAIDPacs are capable of providing a three drive RAID 0 array. Here is the breakdown (table 3) for probability of failure in this configuration:

Years of Operation	Percentage Failure
One	5.0%
Two	23.6%
Three	43.4%

Table 3 – 3 Drive RAID 0 array failure statistics (RAID 0 RAIDPacs)

The FirstRAID G2 and the Tandem DXR can provide a two drive RAID 0. Here are the probabilities of failure for that configuration (table 4):

Years of Operation	Percentage Failure
One	3.4%
Two	16.5%
Three	31.6%

Table 4 – 2 Drive RAID 0 array failure statistics (FirstRAID G2, Tandem DXR)

These probabilities do not include the factor of RAID controller failure, this is just the probability of data loss on the hard drives themselves so probabilities of RAID 0 system failures is higher.

Of course, High-Rely RAIDPacs are not intended for 24/7 operation over a 3 year period as it is presumed they will be swapped on a regular basis. However, 24/7 operation versus power on/off has a reliability factor difference itself that applies to reliability in electronics and is usually lesser for the power on/off side.

Also, these are data loss/corruption errors and if you experience such, it does not necessarily mean that your hardware has failed. Reinitializing the array, reformatting or moving the hardware around may cure the issue.

In summary, if you're using a RAID 0 array, you should be aware of these probabilities and have a strategy for their effects. High-Rely devices are intended for backup purposes and as such, having the multiple backups will usually provide relief against these failure statistics.