1 storage arrangement and failure risks

- *Assuming that the risk of HDD failure is 1/200000 per hour and that the disks are in constant operation, what are the chances of at least one drive failure per year in a storage array made of: a) 6 drives (˜lac3) ? b) 36 drives (˜spcnas) ?*

If r_h is the risk of disk failure per hour of operation, the chance for failure-less operation over a year is $(1-r_h)^{n_h}$ where $n_h = 365 \times 24$ is the number of hours in a year. Thus the risk of failure of a given disk in a year is

$$r_a = 1 - (1-r_h)^{n_h}$$

For $r_h = 1/200000$ we find $r_a = 4.3\%$. The chances for a N drives array to operate in a year without failure are $(1-r_a)^N = (1-r_h)^{n_h N}$so the risk of at least one drive failing is

$$r_{aN} = 1 - (1-r_h)^{n_h N}$$

We find that $r_{a6} = 23\%$ and $r_{a36} = 79\%$.

- *What do you think about the result and its consequences ?*

Building a RAID 0 array of mechanical drives is likely to fail.

- *If a RAID 5 member drive fails, replacing the disk initiates a parity rebuilt process during which each disk must be read entirely. In that case, assuming that the drive r/w speed is 100MB/s, what are the risk of unrecoverable array failure for : a) a 10TB RAID5 array made of 6x 2TB drives (˜lac3) ? b) a 630TB RAID5 array made of 36x18TB drives (˜spcnas) ?*

The full read time at speed s of a disk of capacity C is $t = C/s$. The probability of N drives beeing read without failure for this duration in thus

$$p_{N,C} = (1-r_h)^{NC/s} \simeq 1 - N\frac{C}{s}r_h$$

The risk of failure among $N-1$ drives in the RAID rebuilding process is thus

$$r^5_{fN,C} = 1 - p_{N-1,C} = 1 - (1-r_h)^{(N-1)C/s} \simeq (N-1)\frac{C}{s}r_h$$

where C/s expressed in hours is a) 5.8 hours and b) 52 hours. We find $r^5_{f6,2} = 0.01\%$ and $r^5_{f36,18} = 0.9\%$.

- *For drives bought from the same series, the rate of failure triples if one drive fails. How does the risk of unrecoverable array failure evolve?*

Multiplying r_h by three yields $r^5_{f6,2} = 0.04\%$ and $r^5_{f36,18} = 2.7\%$.

- *What would the risks of unrecoverable array failure be for a RAID6 array (assuming drives are from the same series)?*

For a RAID6 array, we consider the rebuilding process after one drive failed. The risk of two drives failing during the rebuilding process can be evaluated as the complement of the probability that either none or exactly one drive fails during that process. The chances that no drive fails among $N-1$ drives was evaluated before : $p_{N-1,C}$. The chances that exactly one drive fails is $(N-1)(1-p_{1,C})p_{N-2,C}$. Thus the risk of unrecoverable array failure is

$$r^6_{fN,C} = 1 - p_{N-1,C} - (N-1)(1-p_{1,C})p_{N-2,C}$$

which is explicitly

$$r^6_{fN,C} = 1 - (1-r_h)^{(N-1)C/s} - (N-1)[1-(1-r_h)^{C/s}](1-r_h)^{(N-2)C/s}$$

We find $r^6_{f6,2} = 0.0000077\%$ and $r^6_{f36,18} = 0.04\%$

- *What conclusions do you draw from this study?*

The risk of unrecoverable array failure is not negligible for a large RAID5 array. It is acceptable for a RAID6 array. Interestingly, the risks for a 6x2TB RAID5 array and a 36x18TB RAID6 array are similar.