MTTF Computation for RAID Architectures

Boudewijn R. Haverkort
LuFG Informatik—Verteilte Systeme
RWTH–Aachen, D-52056 AACHEN

January 30, 1997

This short note describes how to compute first-order approximations for the MTTF of RAID architectures. We start with defining, for every seperate disk:

- MTTF: the mean time to failure,
- MTTR: the mean time to repair.

Under normal circumstances the MTTF of a disk will be much larger than the MTTR. For the former, think of years, for the latter of hours. Related to the mean failure and repair times are the failure and repair rates:

$$\lambda = \frac{1}{\text{MTTF}}, \quad \text{and} \quad \mu = \frac{1}{\text{MTTR}}.$$

These rate express the mean number of failures or repairs that will occur per unit of time (typically chosen to be an hour). Since the MTTF is much larger than the MTTR, the failure rate λ is much smaller than the repair rate μ. Due to this fact, we may assume that the probability that a disk fails in a period of time equal to the mean repair period is approximately equal to

$$\lambda \times \text{MTTR} = \frac{\text{MTTR}}{\text{MTTF}}.$$

Further notice that if we have N disks operatingn independently, than the overall failure rate of these disks is $N\lambda$, and, consequently, the MTTF for those N disks as a whole is MTTF$/N$.

Now consider the case of RAID-0. Here, we have N independent disks without any redundancy. When 1 fails, the whole system is considered non-operational. Thus, as per the above, we have

$$\text{MTTF}_{\mathit{Raid-0}} = \text{MTTF}/N.$$

Let us now consider RAID-1. Here, all N disks are duplicated. Thus, we have $2N$ disks now, and therefore an overall failure rate $2N\lambda$. However, a single failure does not really matter, there is always the mirrored disk still available. A problem arises when during the repair of the first failed disk, the mirror disk also fails. This happens, approximately, with probability $\lambda \times$ MTTR. As a consequence, the effective rate leading to a system failure equals

$$(2N \times \lambda) \times (\lambda \times \text{MTTR}) = \frac{2N \times \text{MTTR}}{\text{MTTF} \times \text{MTTF}}.$$

Consequently, the reciprocal value of this rate equals the MTTF for RAID-1:

$$\text{MTTF}_{\mathit{Raid-1}} = \frac{\text{MTTF} \times \text{MTTF}}{2N \times \text{MTTR}}.$$

As can be observed, the MTTF increases dramatically (recall that MTTF is normally very large).

Let us finally consider RAID-5. Here, next to the N data disks, there is 1 parity disk (the actual parity information is distributed over all disks, but that does not matter for our evaluation here). Thus, we have $N+1$ disks now, and therefore an overall failure rate $(N+1)\lambda$. However, a single failure does not really matter, there is always the parity disk to correct this failure. A problem arises when during the repair of the first failed disk, any of the other disks also fails. This happens, approximately, with probability $N\lambda \times \text{MTTR}$. As a consequence, the effective rate leading to a system failure equals

$$((N+1) \times \lambda) \times (N\lambda \times \text{MTTR}) = \frac{N(N+1) \times \text{MTTR}}{\text{MTTF} \times \text{MTTF}}.$$

Consequently, the reciprocal value of this rate equals the MTTF for RAID-5:

$$\text{MTTF}_{Raid-5} = \frac{\text{MTTF} \times \text{MTTF}}{N(N+1) \times \text{MTTR}}.$$

As can be observed, the MTTF also increases in comparison to RAID-0, but not as dramatically as in RAID-1. Note, however, that RAID-5 is much cheaper than RAID-1.

Finally, as an example, consider the case where we have $N = 5$ data disks and a MTTF of 10000 hours and a MTTR of 5 hours. For RAID-0 we compute the MTTF as $10000/5 = 2000$ hours $= 0.228$ years. For RAID-1 we find $10000^2/(10 \times 5) = 2000000$ hours ≈ 228 years. Finally, for RAID-5 we compute the MTTF as $10000^2/(6 \times 5 \times 10) = 333333$ hours ≈ 38 years.