

November 2018

How to calculate RAID Reliability?

By Amir Segal

RAID (redundant array of independent disks) disk arrays are very common in critical data infrastructure. The basic principle of RAID is to add data redundancy such that no major data loss will occur when a bad sector appears, or even when a whole disk fails.

The question is: **how reliable are RAIDs?**

If you are designing an IT system that needs to be highly reliable, this paper is for you.

Following is an explanation and example regarding RAID reliability calculations:

A convenient way to calculate RAID reliability is by using Markov chains. BQRs RBD Markov module is ideal for such calculations.

In order to calculate the RAID reliability some parameters are required:

- Disk failure rate: usual values are between 0.5 and 2 failures per million hours, depending on HDD / SSD size and quality. For the following calculations 1 failure per million hours was assumed.
- Failure detection time: time until a bad block is detected. Bad blocks are detected in two cases:
 - The block is read due to user demand
 - The block is periodically read due to RAID scheduled test (scrubbing)

A detection time of 1 week is assumed for the following calculation.

- Rebuild time: The time to reconstruct the failed disk onto a spare or replacement disk. Reconstruction time depends on the amount of data in the failed disk, as well as on the load on the array during rebuild. Data reconstruction using parity calculations requires reading data from all the array disks, therefore the reconstruction time depends also on the number of disks in the array.

A rebuild time of 1 week is assumed for the following calculation. Usually rebuild times are faster, especially if there is low data demand in parallel to the rebuild process.

Note: Disks also have a rate for bad bit reads. For example: if the RAID disks have a rate of 1 bad bit read per 10^{15} bits, and a reconstruction involves reading 10^{14} bits, there is on average 0.1 incorrect reconstructed bits after the rebuild.

Example: RAID 5

Consider an array with 10 disks and a RAID 5 configuration. There are four possible states for the array:

- All the disks are good (this is the initial state)
- A disk failure occurred but is not yet detected
- A disk failure was detected and a rebuild takes place
- More than one disk has failed – **massive data loss!**

Following is a Markov chain diagram for the RAID5

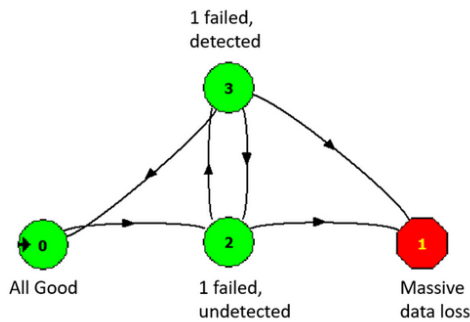


Figure 1: Markov chain diagram of RAID 5.

Transition rates for the Markov chain are:

From state	To state	Notes	Rate (per million hours)	Mean Time (days)
All Good	1 failed, undetected	Rate is 10 X single disk failure rate	10	4,166.666
1 failed, undetected	Massive data loss	Rate is 9 X single disk failure rate	9	4,629.629
1 failed, undetected	1 failed, detected	One week for detection	5,952.38	7
1 failed, detected	Massive data loss	Rate is 9 X single disk failure rate	9	4,629.629
1 failed, detected	1 failed, undetected	New disk failure during rebuild	1	41,666.666
1 failed, detected	All Good	One week for rebuild	5,952.38	7

Using the transition rates, the RAID 5 reliability was calculated for various mission times:

Years	Reliability
1	99.97%
5	99.86%
10	99.72%

Conclusions

RAID reliability depends on the specific RAID configuration, number of drives, the failure rate, rebuild time, and also on the detection time.

The reliability results described above may not be sufficient for critical systems. In many cases smaller arrays (less than 10 HDDs) are recommended.

BQR's RBD software can help you calculate the reliability of various RAID configurations as well as Availability of complex IT and computing system.

For more information please [contact us](#).