Beyond MTTDL: A closed-form RAID 6 reliability equation

Jon Elerath and Jiri Schindler

NetApp Inc.

1 Introduction

Several researchers [1], [4] have noted that the original RAID reliability equation formulated by Gibson and Patterson in 1993 and expressed as mean-time-to-data loss (MTTDL) [3] has outlived its useful life. Since then, HDD technology has evolved, the design of RAID systems has grown in complexity, including processes for proactive scanning and repair of media defects. Finally, we have now a much better understanding of HDD failure modes and non-constant time-to-failure distributions.

Our work aims to derive a closed-form equation for assessing the reliability of *D*+2 parity RAID groups – the most prevalent parity RAID configuration in deployed systems today in the form of EVENODD or RAID-DP. The equation builds upon previous work on formulating equation for single *D*+1 parity RAID [1]. It models many advanced features of modern disk arrays, allows for expression of time-variant HDD operational failure rates and *accurately* models the effects of latent sector errors and media scrub processes have on overall RAID system reliability.

The original MTTDL reliability equation generates unrealistically high numbers for RAID reliability, especially for RAID-6 [1] [4], leading to large discrepancies between the predicted reliability and that experienced in the field. For example, previous work [1] showed that the MTTDL equation for RAID-4 can be off by 3000 times or more and concluded that the MTTDL metric is too inaccurate to be useful or practical when considering the design of a RAID system. It thus developed an alternate model with errors less than 10%. Our work achieves similar level of accuracy for a more complex *D*+2 RAID-6 configuration.

Poor design choices based on inaccurate models can have large impact on performance or lead to unexpected data loss in deployed systems. Previous research attempted to formulate more accurate models. However, they are impractical to use as they require complex formulations with Markov chains or time-consuming Monte-Carlo simulations to obtain one result for a single design point [1]. Additionally Markov chains require constant transition rates that do not adequately capture observed events.

Our goal is to provide an easy-to-use equation to provide system designers, who are not reliability experts, with the ability to readily explore the design space for their next-generation RAID solution and understand its performance and reliability tradeoffs. Therefore, it should be easily understood and applied by reliability neophytes and, most importantly allow for speedy execution in order to allow designers to explore a variety of design points fast. We believe that our formulation, as well as the web-based Javascript implementation we are developing is meeting these goals and that it can put to rest the MTTDL RAID reliability equation that has been shown to be woefully inadequate for 21st century HDD and RAID technology [4].

2 Formulation

There is no exact statistical basis for a closed form equation that models *m*-out-of-*n* failures with repair. By that we mean it is not possible to begin with fundamentals such as PDFs, and derive a closed form equation. Thus, we first model and analyze a RAID-6 system using a Monte Carlo (MC) simulator that has already been shown to be accurate, albeit very time consuming, in these situations [1]. It takes between 14s and 18 hours to produce a result for a single set of inputs. Next, we formulate our equation and compare its results to the field data gathered from many enterprise-class storage systems.

The closed-form equation we developed estimates reliability as the expected number of data-loss events for a redundant array of inexpensive disks in a RAID-6 configuration. The equation expresses HDD operational failures, their restorations, latent (sector) defects, and disk media scrubbing by non-constant distributions that can represent non-homogeneous Poisson processes.

We use a two-parameter Weibull distribution, which can take on many different shapes, modeling increasing-, decreasing-, or constant-over time occurrence rates for each of the modeled events. In formulating our equation, we derived the parameters of these distributions from real-world data collected from thousands of enterprise-class storage systems deployed in the field and containing 100,000s of HDDs arranged into RAID-DP groups – a variant of RAID-6 with 16 disks (*D*=14).

3 The RAID-6 Equation Results

RAID system is continuously succumbing to HDD failures and latent defects, and having these restored and scrubbed. Each time a failure or defect occurs, the RAID group enters into a degraded mode. This process continues throughout time until an operational failure occurs when the RAID group is already in a degraded mode by having an operational failure and a latent defect, or two operational failures.

Our equation computes the expected number of data loss events, $N_{TDF}(t)$, as a function of time and expressed as the number of simultaneously occurring triple disk failures – partial or whole

$$N_{TDF}(t) = (DM_1 + DM_2) \cdot D \cdot H(t) \qquad (1)$$

where, DM_1 is the probability of being in "degraded mode 1", DM_2 is the probability of being in "degraded mode 2", D is the number of data disks in the RAID group, and $H(t)$ is the time-dependent hazard rate for the HDDs. Due to space constrains, we cannot provide the derivation of the terms here.

Figure 1 compares the accuracy of our equation. It shows the results from our closed-form equation fitted with parameters to the four input distributions that we obtained from a data analysis of the field data, including individual run times of scrubbing and RAID repair processes for different disk types. We show results for one particular disk model.

We compare our closed-form equation to the results obtained by a sequential Monte Carlo simulation model originally developed by Elerath for RAID-4 [1]. We adopted it to RAID-6 and used the same parameters for the four input distributions. The graph shows that our closed-form equation yields similar accuracy. More importantly, it yields significant time savings over the MC simulation. Our Javascript implementation calculates results with no human-perceivable delay (approximately 0.1ms). A single MC simulation, on the other hand, ran between 14 seconds and 18 hours. The advantage of the short run time becomes apparent as system designers routinely consider 100s of different parameter settings and distributions for bit error rates, the aggressiveness of the scrubbing and RAID reconstruct background processes.

Finally, we explore the question whether our equation is "better" than MTTDL for predicting the expected number of data loss events, expressed as triple-disk failures (TDF). Using the inverse of the MTTDL as a rate of occurrence of failure for the RAID group and multiplying by the time in the field, which is erroneous but done regularly, results in a linearly increasing function, as shown in Figure 1. The MTTDL equation-

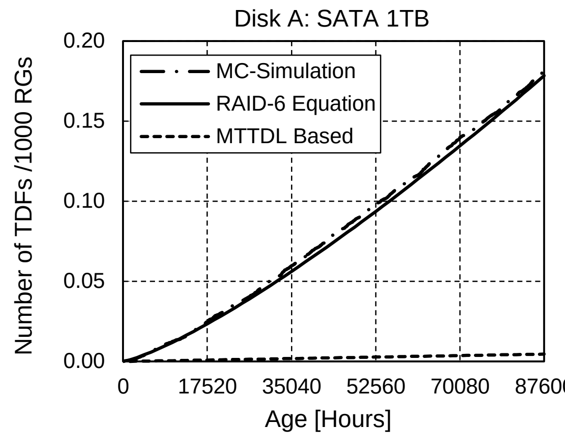

Figure 1: Comparison of the Monte Carlo simulation, the MTTDL equation, and the closed-form RAID-6 equation results.

predicted number of TDFs in 10 years is so small for it appears to be on the horizontal axis.

To the best of our knowledge, this work the first formulation of a closed-form RAID-6 reliability equation that is simple enough to be plugged into a spreadsheet yet accounts for the various advanced features built into modern systems, including the availability of spare drives and rate-limited continuous media scrubbing. Our hope is that it will replace the original MTTDL reliability formula whose time has passed. We also believe that, thanks to the Javascript implementation that we made available, it will be used by RAID architects to explore design tradeoffs as well as system dependability experts who wish to accurately capture the reliability of a storage system used in practice.

References

[1] J. Elerath and M. Pecht. A highly accurate method for assessing reliability of redundant arrays of inexpensive disks. IEEE Trans. on Computers, vol. 58, no. 3, 2009.

[2] Y. Gao, D. Meister, and A.Binkmann. Reliability analysis of declustered-parity RAID 6 with disk scrubbing and considering irrecoverable read errors. IEEE Conf. on Networking, Architecture, and Storage, 2010.

[3] G. Gibson, and D. Patterson, Designing disk arrays for high data reliability. Journal of Parallel and Distributed Computation, vol. 17, January 1993.

[4] K. Greenan, J. Plank, and J. Wylie. Mean time to meaningless: MTTDL, markov models, and storage system reliability. In *Proc. of the 2nd USENIX conference on Hot Topics in Storage and File Systems*, 2010.

[5] D. Kececioglu. Reliability & Life Testing Handbook, Volumes 1 & 2. Prentice Hall, 1993.

[6] D. A. Patterson, G. Gibson, R. Katz, A case for redundant arrays of inexpensive disks (RAID). In *Proc. of the ACM SIGMOD Conference*, 1988.