Effects of Data Scrubbing to Reliability in Storage Systems

Junkil Ryu
lancer@postech.ac.kr
Pohang University of Science and Technology

Chanik Park
cipark@postech.ac.kr
Pohang University of Science and Technology

Abstract

Today, even bytes stored in expensive, professionally managed data centers, are vulnerable to loss and damage. The vulnerability increases when large volume of data must be stored indefinitely into the future, as required by emerging data services. The economic viability of these services depends on their keeping data unaltered. As disks outpace tapes in capacity growth, large scale storage systems based on disks are becoming increasingly attractive. Such systems will encompass a very large number of disks and store petabyte of data. In large systems encompassing thousands or even tens of thousands of active disks, disk failure will become frequent, necessitating redundant data storage. Until now, the reliability measure in disk drive-based storage systems focus predominantly on protecting against complete disk failures. However, very little is known about the incidence of latent sector errors i.e., errors that go undetected until the corresponding disks sectors are accessed. Hence, periodically accessing the entire storage (data scrubbing) is necessary to protect data loss due to latent sector errors. In this research, we made a standard Markov failure and repair model with data scrubbing mechanism to measure the reliability of the storage system. And we determined the replicas and the period of the data scrubbing for the demanded reliability of the disk-based storage system by using the proposed Markov model.

1. Introduction

Most storage systems today make certain disk drive reliability assumptions and build redundancy mechanisms such as RAID to compensate for complete disk failures. However, the inability to either temporarily or permanently access data from certain sectors can also affect the MTTDL (Mean Time To Data Loss). Such incidents are often referred to as latent sector errors because the disk drive does not report any error until the particular sector is accessed. The impact of latent sector errors on the MTTDL in RAID systems is well known. Hafner et al. pointed out that a single latent sector error can lead to data loss during RAID group reconstruction after a disk failure. Similarly, Baker et al. developed new RAID equations that account for latent sector errors when calculating the MTTDL.

2. Data Scrubbing

To obtain the data replication degree and the data scrubbing interval satisfying the demanded reliability, we make a simple Markov model. This model can be applicable in the distributed storage system and the direct-attached storage system. Data replication ensures that the system is able to guard against individual disk drive failures but does not ensure that the system is able to guard against data losses due to the latent sector errors. So, the system need to confirm untouched data periodically. To maintain data's long-term availability, the system must constantly repair replicas lost due to disk failures and latent sector errors. Thus, the three parameters, which the systems uses to ensure long-term availability of replicas, are the number of replicas, the repair rate, and data scrubbing interval.

For the reliability of the storage system, let's consider a simple storage system, which is comprised of 6 SATA disk drives (each drive's capacity is 1TB, one drive is spare). We assumed that average read performance of each drive is 100MB per second, in fact less than 80 MB per second in the case of sequential read operation. If a disk drive failure occurs, a repair process will be done. In this case, data can be lost when another disk drive fails or latent sector errors in another disk drive have been occurred. The time to repair failed data and to find latent sector errors becomes longer as the drive's capacity increases rapidly because the drive's bandwidth increases very slow. Reconstruction of a failed 1TB SATA drive in a 5 drive stripe can have up to 40% chance of an unrecoverable error ($5 \cdot 1TB \cdot 10^{-14}$ = 0.4, where a typical SATA drive specification for an unrecoverable error is 10^{-14}).

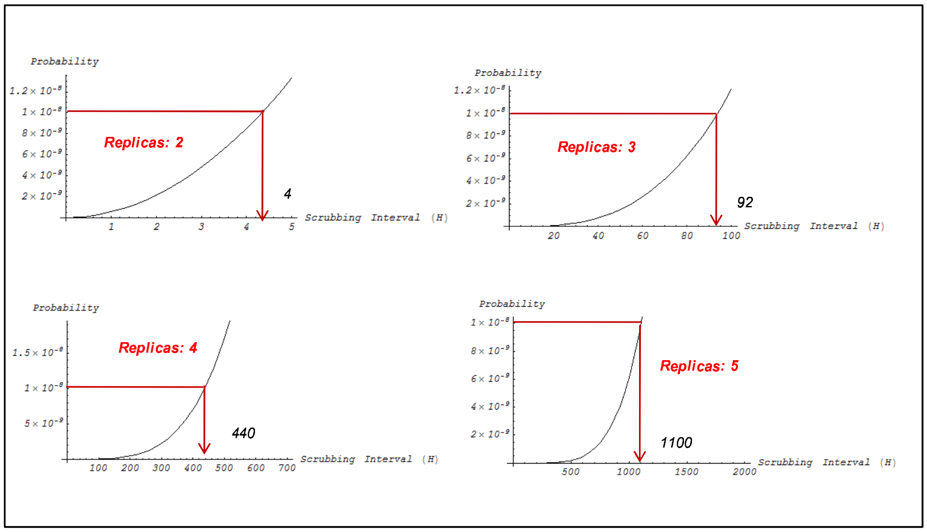

Figure 2. Data Loss Probability vs. Scrubbing Interval

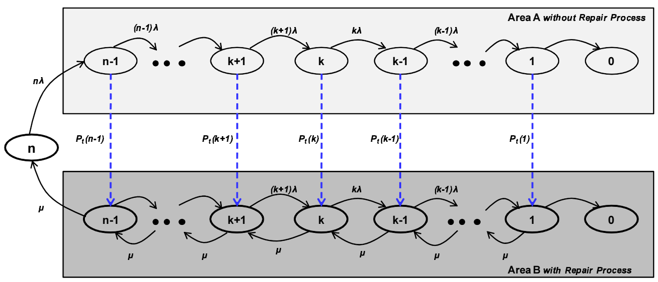

Figure 1. Markov model for reliability analysis

2.1. Analytical Model

We consider a system with N disks (or nodes with a disk drive) and the replication degrees of the data is n. We first assume that the latent sector error rates are larger than the disk drive failure rates, which are deduced from the Bairavasundaram et al. and etc. Hence, a failure rate in the proposed Markov model is obtained from the latent sector errors. And each latent sector errors and each disk drive failures are independent. Failure rates (latent sector errors) are assumed to follow an exponential distribution with MTTF (Mean Time To Failure) as its mean. To analyze the reliability of the storage system especially in the view of latent sector errors, we introduce a Markov model as Figure 1. In the model, state i represents the number of the live data replicas in the storage system. State i of area A in Figure 1 does not have the repair process but state i of area B has the repair process. After latent sector errors are reported through the data scrubbing process, the transition from area A to area B will occur.

The probability found in the k state of area B, $P_{state}(K)$, is ${}_n\mathrm{C}_k \cdot (e^{-\lambda t})^k \cdot (1 - e^{-\lambda t})^{n-k}$, where λ is the failure rate. The probability returning state n from state k, $P_{rel}(t)$ is $\{1 - \sum_{k=0}^{n-1} Q_k \times P_{state}(k)\}$, where $P_{rel}(t)$ represents the probability for at least a replica to survive after t time and the probability to return state n from state k. Q_k is the probability that the system reach state n before state 0 starting in state k.

We apply this model in the following storage system. λ (latent sector failure rate) is 1/43,800, which is induced from the Bairavasundaram et al. μ (repair rate) is assumed to 36, where replica size is 100MB. Repair bandwidth (or data scrubbing bandwidth) in this system is 1MB per second. In Figure 2, data scrubbing interval can increase as the number of data replicas increases. If the number of data replicas is 2, data scrubbing must be done within about 4 hours periodically for the seven nines' reliability (the chance to a data loss is 1×10^{-8}). According to the types of the storage system, the different parameters (the replication degree, the data scrubbing interval, the repair (or scrubbing) bandwidth) can be managed for the reliability. Changing the number of the replicas is available in the distributed storage system for the reliability. Changing the number of the replicas is unavailable in the disk array system due to given array configuration.