

## A Report on Simulation-Driven Reliability and Failure Analysis of Large-Scale Storage Systems



Lipeng Wan  
Feiyi Wang  
Sarp Oral  
Sudharshan Vazhkudai  
Qing Cao

**Approved for public release:  
distribution is unlimited.**

**November 2014**

## DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

**Website** <http://www.osti.gov/scitech/>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service  
5285 Port Royal Road  
Springfield, VA 22161  
**Telephone** 703-605-6000 (1-800-553-6847)  
**TDD** 703-487-4639  
**Fax** 703-605-6900  
**E-mail** [info@ntis.gov](mailto:info@ntis.gov)  
**Website** <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information  
PO Box 62  
Oak Ridge, TN 37831  
**Telephone** 865-576-8401  
**Fax** 865-576-5728  
**E-mail** [reports@osti.gov](mailto:reports@osti.gov)  
**Website** <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

National Center for Computational Sciences

## **A Report on Simulation-Driven Reliability and Failure Analysis of Large-Scale Storage Systems**

---

Lipeng Wan\*   Feiyi Wang†   Sarp Oral†   Sudharshan Vazhkudai†   Qing Cao\*

\* University of Tennessee

† Oak Ridge National Laboratory

Prepared by  
OAK RIDGE NATIONAL LABORATORY  
Oak Ridge, Tennessee 37831-6283  
managed by  
UT-BATTELLE, LLC  
for the  
U.S. DEPARTMENT OF ENERGY

This research was supported by, and used the resources of, the Oak Ridge Leadership Computing Facility, located in the National Center for Computational Sciences at ORNL, which is managed by UT Battelle, LLC for the U.S. DOE (under the contract No. DE-AC05-00OR22725).

## Abstract

High-performance computing (HPC) storage systems provide data availability and reliability using various hardware and software fault tolerance techniques. Usually, reliability and availability are calculated at the subsystem or component level using limited metrics such as, mean time to failure (MTTF) or mean time to data loss (MTTDL). This often means settling on simple and disconnected failure models (such as exponential failure rate) to achieve tractable and close-formed solutions. However, such models have been shown to be insufficient in assessing end-to-end storage system reliability and availability.

We propose a generic simulation framework aimed at analyzing the reliability and availability of storage systems at scale, and investigating *what-if* scenarios. The framework is designed for an end-to-end storage system, accommodating the various components and subsystems, their interconnections, failure patterns and propagation, and performs dependency analysis to capture a wide-range of failure cases. We evaluate the framework against a large-scale storage system that is in production and analyze its failure projections toward and beyond the end of lifecycle. We also examine the potential operational impact by studying how different types of components affect the overall system reliability and availability, and present the preliminary results.

## 1 Introduction

The reliability and availability of data is at the core of any storage system and of critical concern for storage system architects. In order to provide reliable data services, a storage system must be able to tolerate different types of faults and failures. The complexity of this problem increases with scale. Large-scale storage systems are often built with *scalable units*. These basic building blocks are replicated to meet the required capacity and performance targets. Examples of such building blocks include a DDN SFA12K device with up to 1,680 disk drives [1], an IBM DS8870 system supporting 1,536 disk drives [2], a Panasas ActiveStor 16 system with 20 hard disk drives and 10 solid state disks per shelf with up to 100 shelves per system [3], and the NetApp FAS8080 EX series device which can support 12 HA pairs per system, with each HA pair containing up to 1,440 disk drives [4]. A perhaps dated example is the Google File System cluster, which is composed of more than 1,000 such building blocks [5].

Therefore, it would be logical to conduct reliability analysis at the scalable unit-level, assuming that the failures across scalable units are independent. However, many existing studies in this domain heavily concentrate on analytically estimating the reliability at either the disk population or the redundant disk array level. While they may be one of the most critical parts of the system, they are not sufficient to represent the entire system. Simple failure models coupled with imperfect assumptions allow for the use of close-form calculations, however, they also compromise the overall accuracy and usefulness of the reliability estimation [6]. As an example, one such common model is built with continuous Markov chains, which has an underlying assumption that the failure or repair rates are constant (time independent)

[7, 8, 9, 10]. However, it has been shown that these rates can vary with time, which makes obtaining an analytical solution very difficult [11].

In addition to the analytical modeling approach, there have been several research efforts on reliability estimation through simulation [11, 12, 13]. Compared to analytical modeling, simulations are more intuitive and easy to understand, and they can capture the temporal variations in failure or repair rates with relative ease. To the best of our knowledge, existing simulation approaches only focus on disk redundancy group failures. However, a failure of any given subsystem or component may cause a cascading effect and significantly impact the overall reliability of a storage system [14]. For instance, if a disk enclosure fails, all disk drives in that enclosure will be inaccessible until it is repaired. Therefore, failure characteristics of all hardware components and their dependencies must be evaluated systematically.

We propose a generic simulation framework for analyzing the reliability of storage systems at scale. The framework has been designed to accommodate a diverse set of components and subsystems, with different reliability characteristics. It provides flexible mechanisms to support temporal variations for each component. It is topology-aware, in the sense that the physical connection of components and their failure dependencies are tracked and analyzed. We validate our framework with real-world failure data, which was obtained from a large-scale production storage systems (Spider I) deployed at the Oak Ridge National Laboratory (ORNL) [15]. Our analysis shows that, forklift upgrades can be preferred when increasing number of data loss events are projected due to accelerated failure rate of aging hardware. The results further reveal surprising insights on the impact of spare components to the overall system reliability. This work also demonstrates the necessity of conducting a reliability evaluation at the system-level.

## 2 Storage System Reliability Simulator

This section aims to present both an overview on the architecture design as well as some implementation details of the simulation framework, with particular focus on mechanisms for achieving failure dependency analysis and synthesis of the two-phase simulation results to cover wide-ranging failure scenarios.

### 2.1 Design Considerations

A large-scale storage system is often complex and comprises of tens of thousands of disk drives and other hardware components, such as power supplies, RAID controllers, etc. Some of these components may consist of even smaller subcomponents. For example, a disk drive consists of spindle motor, and actuators. It is neither feasible nor necessary to simulate every single hardware component inside a disk drive or a RAID controller. Selecting the right abstraction with the right amount of details to simulate (and details to ignore) has always been the most challenging aspect of simulations.

System vendors and HPC operators are often, understandably, cagey about sharing detailed errors and failure data about their systems. Yet, having that data is crucial to validate and provide confidence

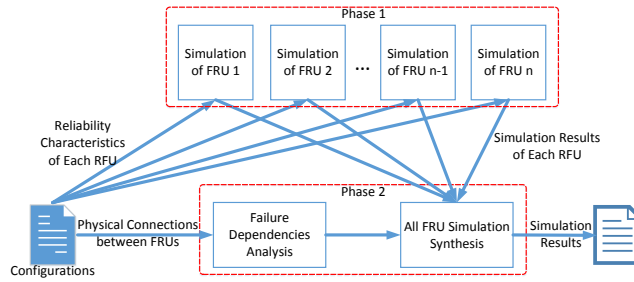


Figure 1: Two-phase Simulation Framework

to any simulation design of this nature. Our experience suggests that information such as field AFR (annual failure rate) of each FRU (field replaceable units) are more or less maintained by both vendors and operators alike. Therefore, it is not unreasonable to assume an FRU as a building block. With sufficient failure logs, we can establish reasonable failure models for each FRU, which will further drive the reliability estimate at the system level.

As we have alluded to earlier, simulating each FRU standalone is not enough to obtain the estimation on reliability of the entire storage system. One FRU's failure might also have a cascading effect on other FRUs, as there is often a correlation between failures of hardware components in close-coupled storage systems. For example, a RAID controller's failure might lead to the unavailability of hundreds of disk drives. Therefore, in order to estimate the reliability of the entire storage system, our simulation framework must be able to synthesize the simulation results of all FRUs by analyzing the failure dependencies between them.

To this end, we put forth a two-phase analysis approach as shown in Fig. 1, which is inspired by a conventional diagrammatic method for modeling the reliability of complex systems, namely the reliability block diagram (RBD) [16]. An RBD is composed of a set of blocks connected in parallel or serial, and each of which represents a component of the system. Essentially, RBD in this context reveals how failures of one component propagate and contribute to the failures of other components or the entire system.

The input to the simulation framework is a configuration file that contains both specifications of reliability characteristics for each FRU and layout information for physical topology. In phase 1, each FRU will be simulated independently based on its reliability characteristics, and the failure events are logged throughout its lifecycle. In phase 2, the framework extracts failure dependencies from all FRUs, and builds an RBD based on the topology of the storage system. Given all the extracted failure dependencies, the framework synthesizes the results across all components, and provides detailed information on the estimates of the various metrics of interest, such as average number of failed FRUs during the simulation window, events leading to data unavailability or data loss and for how long.

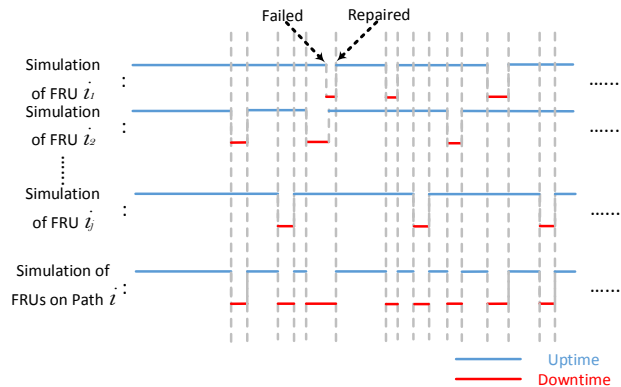


Figure 2: Uptime and Downtime of FRUs on Path  $i$

## 2.2 Implementation

### 2.2.1 Simulate Individual FRU

In prior work, both exponential and Weibull distributions have been widely used to model or simulate disk failures [11, 12, 13]. Exponential distribution has a constant failure rate, while Weibull distribution has a non-constant but monotonic failure rate. However, in reality, the failure rates of a component often follow a “bathtub” curve [17]. Therefore, to make our simulation more realistic, we need to build a probability distribution that has a non-monotonic failure rate. For this purpose and for greater flexibility, the framework provides users a mechanism to specify different failure rate characteristics for different time intervals.

The simulation of each FRU is based on discrete events. As illustrated in Fig. 2, for each FRU, the simulator creates an individual process which will randomly generate two types of discrete events (failure and recovery) based on the given reliability parameters. At the beginning of the simulation, the process generates a random number as the uptime. Then the process is suspended and will not resume until such amount of (simulated) time has passed to trigger the failure event. Once a failure occurs, the process will generate another random number as downtime and is suspended again to wait for the recovery event.

### 2.2.2 Analyze Failure Dependencies

In order to obtain the reliability estimation of the entire storage system, we build the RBD of the storage system and analyze the failure dependencies between different FRUs. For example, the RBD of a typical scalable storage unit that contains FRUs listed in Table 1 is shown in Fig. 3, in which each block is assigned a unique ID to represent an FRU.

Similar to a tree structure (but essentially not), in the RBD, each block can have several child blocks, which are below it in the diagram. A block that has a child is called the child’s parent block. Particularly, for the convenience of using graph algorithms, we create a dummy block as the root (block o) of all blocks

Table 1: FRUs in a typical scalable storage unit

Name	Number	IDs
Controller	2	15-16
Controller's Power Supply (House)	2	1-2
Controller's Power Supply (UPS)	2	8-9
Disk Enclosure	5	27-31
Disk Enclosure's Power Supply (House)	5	3-7
Disk Enclosure's Power Supply (UPS)	5	10-14
I/O Module	10	17-26
Disk Expansion Module (DEM)	40	32-71
Baseboard	20	72-91
Disk Drive	280	92-371

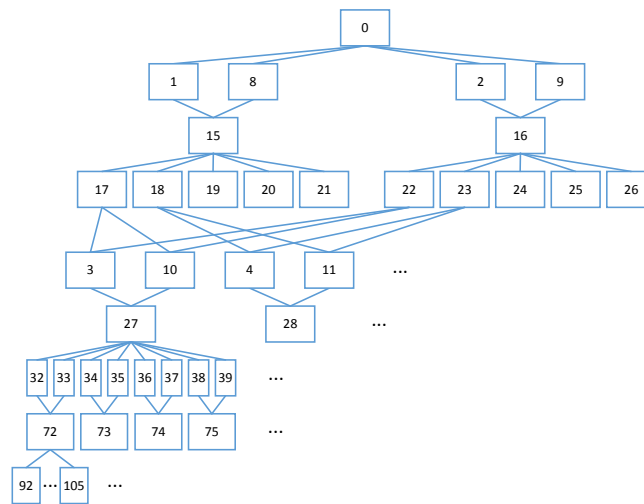


Figure 3: Reliability Block Diagram

in the RBD, which does not represent any real FRU. Different from the tree structure, a block in the RBD can have more than one parent. The reliability of each block depends on its parents, while determining that of its children. For instance, one of the controllers (block 15) has two parents, the house power supply (block 1) and UPS power supply (block 8). If and only if both of these two parents fail, the controller will fail. Meanwhile, the controller also has five children which are the I/O modules (block 17-21). If the controller fails, all five I/O modules will fail.

### 2.2.3 Synthesize Simulation Result at System Level

Once we obtain the simulation results of all FRUs, we need to synthesize them based on the RBD we built (failure dependencies between different FRUs), to derive the duration of temporary data unavailability and permanent data loss. The following are the concrete steps.



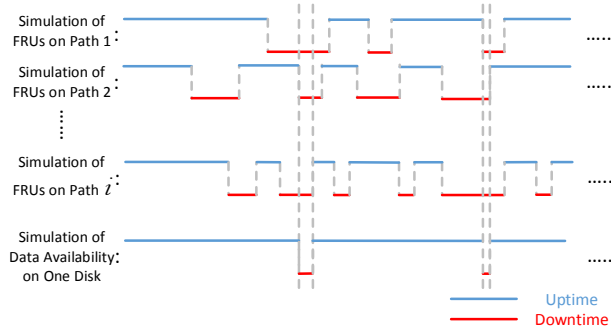


Figure 4: Data Unavailability of One Disk

First, we need to find all paths that start from each disk drive to the root in the RBD using a depth-first search (DFS). For instance, in the storage unit shown in Fig. 3, there are sixteen paths from each disk drive to the root. Second, we apply the “AND” operation to up/downtime sequences of FRUs on each path to calculate the simulation results of each path, as shown in Fig. 2. Third, we apply the “OR” operation to results of all paths belonging to each disk to obtain the duration data is unavailable, as shown in Fig. 4. Finally, with the knowledge of which disk drive belongs to which RAID group and the RAID level, we can derive the duration of temporary data unavailability and permanent data loss for each RAID group. For example, disks in Fig. 3 are organized as RAID level 6, and each RAID group can tolerate up to 2 disk failures. If more than 2 disks are unavailable at the same time, the RAID group will encounter a temporary data unavailability, while if more than 2 disks fail at the same time, a permanent data loss will occur.

### 3 Evaluation

#### 3.1 Disk Subsystem Reliability Evaluation

Following a bottom-to-top approach, we start the evaluation of our simulation with disks. We used the failure logs from the Spider I production storage system at the Oak Ridge National Laboratory. Spider I was the 10 PB production storage system for the Jaguar and Titan supercomputers, with approximately 13,400 disks. Spider I was decommissioned after six years of production.

In order to accurately simulate the failure characteristics of a disk population, we used a crafted distribution that has a non-monotonic failure rate. We used two independent Weibull distributions, with decreasing and increasing failure rates, respectively. Specifically, we used a method called inverse transform sampling [18] to build a probability distribution that has the “bathtub” failure rates.

The manufacturer’s specification often indicates that the disk drives have an annualized failure rate (AFR) less than 1%, or an MTTF of more than 1,000,000 hours. These are too optimistic to be accurate and real production disk subsystems typically observe 2-4% and up to 13% AFR in practice [19].

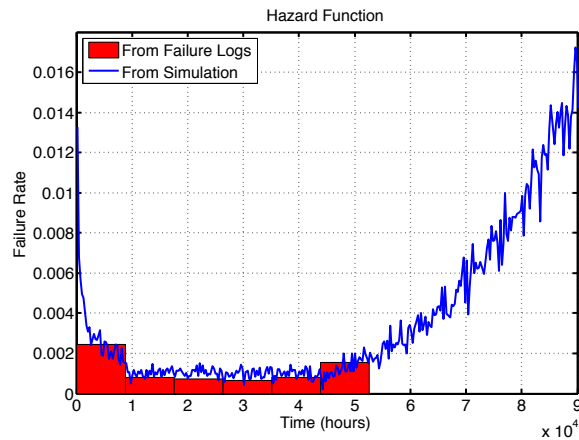


Figure 5: “Bathtub” Hazard Function

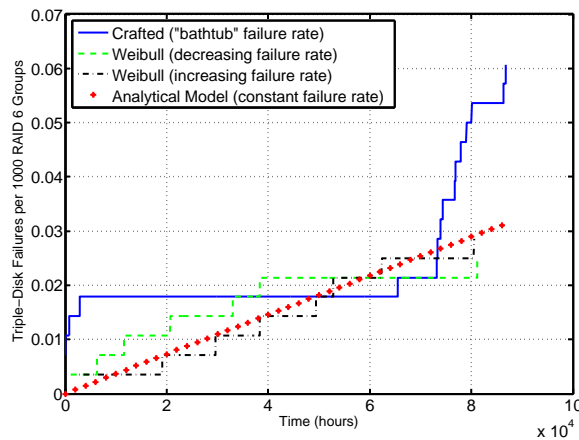


Figure 6: Cumulative Failure Function Comparison

Solid bars in Fig. 5 show the actual AFR for six years (deployment to decommission) for Spider I, while the curve represents the failure rates calculated by our simulation analysis, using the crafted failure distribution. The simulation was run 10,000 times. The mission time in simulation was set to 87,600 hours (10 years), since very few disk drives ever survive for more than 10 years of operation.

As can be seen from Fig. 5, our simulation using a crafted distribution is very accurate. Field obtained disk failure data forms a bathtub curve and the simulated AFR very closely tracks this curve.

The cumulative failure function is illustrated in Fig. 6, where the  $x$ -axis represents time and  $y$ -axis represents triple-disk failures per 1,000 disk redundancy groups (RAID 6). As observed in Fig. 6, our crafted distribution estimates an explosive increase in triple disk failure per 1000 RAID 6 groups after 6 years of operation. The analytical model with a constant failure rate falls short of capturing this.

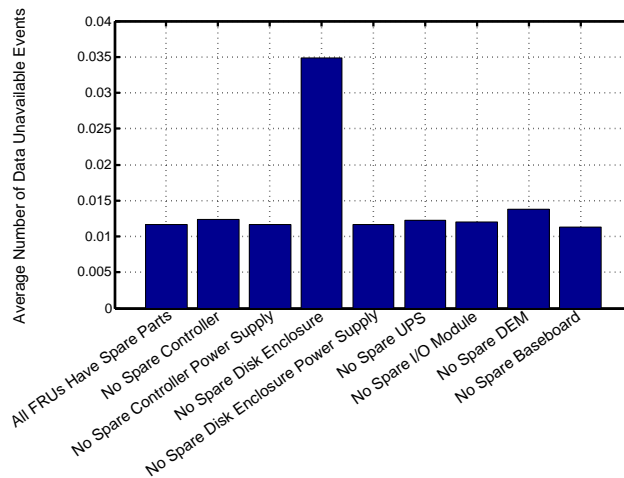


Figure 7: The impact of FRU repair times

### 3.2 Storage System Reliability Evaluation

Storage systems are composed of multiple subsystems and components. Our simulation framework classifies and captures these subsystems as different FRUs and establishes dependencies among them to provide an accurate reliability estimation of the end-to-end storage system.

Large-scale storage systems are often built with scalable units. A typical scalable unit consists of FRUs listed in Table 1. However, since there is almost no dependency between different scalable storage units, we can estimate the reliability of one unit and simply duplicate the results to obtain the estimation on the entire storage system. The simulation parameter settings of all FRUs are listed in Table 2. These values are selected accordingly with field failure data where available (Spider I), or from manufacturers' data sheets as needed. We ran the simulation of one scalable unit for 10,000 times and the simulated mission time is 87,600 hours (10 years). The results are shown in Table 3.

The repair time is directly correlated with the on-site availability of spare parts. In order to assess the overall system reliability with respect to the number of components in the spare pool, we varied the repair time in our simulation. As an example, an analysis of system reliability with respect to the number of spare RAID controllers can be found in Fig. 7. Here, the y-axis represents the average number of data unavailability incidents during a 10 year mission. Similar evaluations for other FRU types can also be found in 7.

From the results, we can observe that the availability of spares for all FRUs have a similar impact, except the disk enclosures. The simulation study suggests that HPC center should provision more budget for spare disk enclosures.

We also compare the average number of simulated data unavailability incidents with respect to varying disk enclosure MTTR values in Fig. 8. As can be seen, the average number of data unavailability

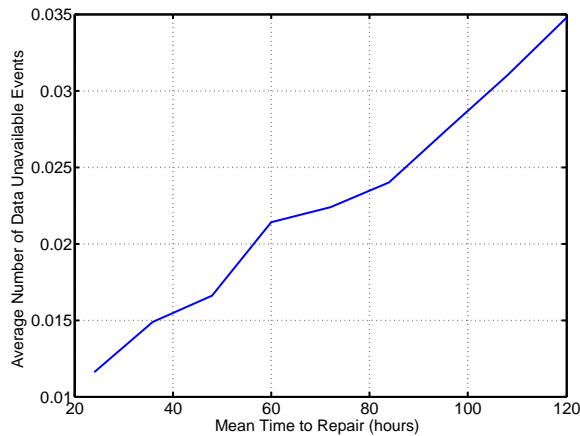


Figure 8: The impact of increasing disk enclosure MTTR

incidents linearly increase with disk enclosure MTTR. As an example, for a large-scale storage system with 48 scalable units and 5 disk enclosures per unit, if one wishes to limit the data unavailability incidents due to faulty disk enclosures, the average number of data unavailability incidents should be less than  $\frac{1}{48}$  which is  $\approx 0.0208$ . Therefore, the MTTR of disk enclosures should be less than 60 hours.

### 3.3 Discussion

Our evaluation sheds light on how to intelligently balance the depth of spare pools across various components under a fixed budget. The simulation results summarized in Figure 7 are surprising, as they show that the availability of spare disk enclosures had the biggest impact on the system reliability. For other systems the outcome might be different but we strongly believe that performing a similar evaluation would be beneficial towards optimizing spare pool resources with respect to system availability and reliability.

A component-level, close-form calculation can provide unrealistic MTTF values (*meantime to meaning-less* [6]). As our evaluation shows (Figure 7), all components have an impact on the overall system reliability. Therefore, we particularly recommend performing a system-level reliability simulation incorporating failure dependencies among subsystems.

Another insight is on upgrades. A system can be upgraded incrementally. However, performing a forklift upgrade can be technically more appropriate and cost effective under certain conditions. As shown in Figure 5, our simulation results are indicating a strong possibility of increasing number of disk failures after the first 6 years of operation. Also, Figure 6 projects a higher possibility of triple disk failures (i.e., data loss incidents) for the same time window. Given this information, a forklift upgrade is the right technical and cost-effective solution for an aging storage system, such as Spider I for providing continued and reliable operations. We highly recommend performing a similar evaluation before making any upgrade decisions.

## 4 Related Work

Reliability studies for storage system have been conducted on several fronts. Analytical modeling, coupled with field data analysis and fitting are among the most common approaches. A large body of existing work focused on building probability models for failures of disk drives and data losses of RAID groups [7, 8, 9, 10, 20, 21]. Almost all of these models assume that disk failures and recoveries fit the memory-less, exponential distribution, which means the failure or repair rates of disk drives are constant (time independent). In cases where these rates can vary over time [19, 22], such time dependency makes it difficult to derive an accurate analytical solution [11]. For example, by analyzing the disk failure logs collected from real storage systems, Schroeder and Gibson have shown that the disk failure better fits a Weibull distribution, which has a decreasing failure rate [19].

Besides analytical modeling, a few existing studies have also tried to estimate the reliability of a storage system through simulation [12, 11, 13]. In particular, Elerath and Pecht [12] have implemented a Monte Carlo simulation for RAID 4 groups to evaluate how time dependent failure and repair rates impact the average number of data loss events that could occur during a given mission time. Greenan developed a high-fidelity reliability simulator for erasure-coded storage systems which is capable of analyzing the reliability of arbitrary, linear multi-disk fault tolerant erasure codes [11]. However, most if not all of the simulation-based approaches focus at the component-level, i.e., disk or RAID group failures. While some of field data suggests that failures of other hardware components also contribute to considerable percentages of storage system failures, and therefore, an end-to-end system-level approach is better suited to capture such failures.

## 5 Conclusions

We have presented a simulation-driven study for analyzing the reliability and availability of large-scale storage systems. Our approach considers the end-to-end storage system as a series of connected components and subsystems. It calculates independent failure models for each and also tracks how these failures then propagate through out the entire storage system. As a case study, we evaluated our simulation framework against the field failure data obtained from a large-scale production storage system located at the Oak Ridge National Laboratory (ORNL), which allowed us to investigate, gain insights, and make projections on *what-if* scenarios of interests. We believe the framework and approach are general enough to have wide-ranging applications.

## Acknowledgements

This research used the resources of the Oak Ridge Leadership Computing Facility, located in the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725.

## References

- [1] DataDirect Networks, Inc. (2014) DDN SFA12K Family. [Online]. Available: <http://www.ddn.com/products/storage-platform-sfa12k/>
- [2] IBM. (2014) IBM DS8000 Series. [Online]. Available: <http://www-03.ibm.com/systems/storage/disk/ds8000/overview.html>
- [3] Panasas, Inc. (2014) Panasas ActiveStor 16. [Online]. Available: <http://www.panasas.com/products/activestor>
- [4] NetApp, Inc. (2014) NetApp FAS8080 EX. [Online]. Available: <http://www.netapp.com/us/products/storage-systems/fas8000/>
- [5] S. Ghemawat, H. Gobioff, and S.-T. Leung, "The Google File System," in *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles*, ser. SOSP '03. New York, NY, USA: ACM, 2003, pp. 29–43.
- [6] K. M. Greenan, J. S. Plank, and J. J. Wylie, "Mean time to meaningless: Mttddl, markov models, and storage system reliability," in *Proceedings of the 2nd USENIX conference on Hot topics in storage and file systems*. USENIX Association, 2010, pp. 5–10.
- [7] G. A. Gibson and D. A. Patterson, "Designing Disk Arrays for High Data Reliability," *Journal of Parallel and Distributed Computing*, vol. 17, no. 1-2, pp. 4–27, Jan. 1993.
- [8] P. M. Chen, E. K. Lee, G. A. Gibson, R. H. Katz, and D. A. Patterson, "RAID: High-performance, Reliable Secondary Storage," *ACM Computing Surveys*, vol. 26, no. 2, pp. 145–185, Jun. 1994.
- [9] M. Schulze, G. Gibson, R. Katz, and D. Patterson, "How Reliable Is A RAID," in *COMPCON Spring 89. Thirty-Fourth IEEE Computer Society International Conference: Intellectual Leverage, Digest of Papers*. IEEE, 1989, pp. 118–123.
- [10] D. A. Patterson, G. Gibson, and R. H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," in *Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '88. New York, NY, USA: ACM, 1988, pp. 109–116.

- [11] K. Greenan, "Reliability and Power-Efficiency in Erasure-Coded Storage Systems," University of California, Santa Cruz, Tech. Rep. UCSC-SSRC-09-08, Dec. 2009.
- [12] J. G. Elerath and M. Pecht, "Enhanced reliability modeling of raid storage systems," in *In Proceedings of the International Conference on Dependable Systems and Networks (DSN)*, 2007, pp. 175–184.
- [13] J. G. Elerath and J. Schindler, "Beyond MTDDL: A Closed-Form RAID 6 Reliability Equation," *Trans. Storage*, vol. 10, no. 2, pp. 7:1–7:21, Mar. 2014.
- [14] W. Jiang, C. Hu, Y. Zhou, and A. Kanevsky, "Are Disks the Dominant Contributor for Storage Failures?: A Comprehensive Study of Storage Subsystem Failure Characteristics," *Trans. Storage*, vol. 4, no. 3, pp. 7:1–7:25, Nov. 2008.
- [15] G. Shipman, D. Dillow, S. Oral, and F. Wang, "The Spider Center Wide File System: From Concept to Reality," in *Cray User Group (CUG) Conference*, Atlanta, May 2009.
- [16] M. Rausand and A. Høyland, *System Reliability Theory: Models, Statistical Methods and Applications*, 3rd ed. Wiley-IEEE, Nov. 2003.
- [17] B. L. Amstader, *Reliability Mathematics: Fundamentals, Practices, Procedures*. McGraw-Hill, 1971.
- [18] L. Devroye, "Sample-based Non-uniform Random Variate Generation," in *Proceedings of the 18th Conference on Winter Simulation*, ser. WSC '86. New York, NY, USA: ACM, 1986, pp. 260–265.
- [19] B. Schroeder and G. A. Gibson, "Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You?" in *Proceedings of the 5th USENIX Conference on File and Storage Technologies*, ser. FAST '07. Berkeley, CA, USA: USENIX Association, 2007.
- [20] Q. Xin, E. L. Miller, T. J. E. Schwarz, D. D. E. Long, S. A. Brandt, and W. Litwin, "Reliability Mechanisms for Very Large Storage Systems," in *IEEE Symposium on Mass Storage Systems*, 2003, pp. 146–156.
- [21] K. K. Rao, J. L. Hafner, and R. A. Golding, "Reliability for Networked Storage Nodes," in *International Conference on Dependable Systems and Networks (DSN)*. IEEE Computer Society, 2006, pp. 237–248.
- [22] E. Pinheiro, W.-D. Weber, and L. A. Barroso, "Failure Trends in a Large Disk Drive Population," in *Proceedings of the 5th USENIX Conference on File and Storage Technologies*, ser. FAST '07. Berkeley, CA, USA: USENIX Association, 2007, pp. 2–2.

Table 2: Parameter Settings of FRUs' Failure Simulation

FRU's Name	Time to Failure		Time to Repair	
	Distribution	Parameters	Distribution	Parameters
Controller	exponential	rate = $5.3 \times 10^{-6}$	weibull	shape = 1.0, scale = 24.0, location = 4.0
Controller's Power Supply (House)	exponential	rate = $9.5 \times 10^{-7}$	weibull	shape = 1.0, scale = 24.0, location = 2.17
Disk Enclosure	exponential	rate = $2.6 \times 10^{-7}$	weibull	shape = 1.0, scale = 24.0, location = 5.0
Disk Enclosure's Power Supply (House)	exponential	rate = $9.1 \times 10^{-8}$	weibull	shape = 1.0, scale = 24.0, location = 2.17
Power Supply (UPS)	exponential	rate = $4.4 \times 10^{-6}$	weibull	shape = 1.0, scale = 24.0, location = 3.0
I/O Module	exponential	rate = $4.3 \times 10^{-7}$	weibull	shape = 1.0, scale = 24.0, location = 2.17
Disk Expansion Module (DEM)	exponential	rate = $2.6 \times 10^{-7}$	weibull	shape = 1.0, scale = 24.0, location = 2.17
Base-board	exponential	rate = $2.6 \times 10^{-7}$	weibull	shape = 1.0, scale = 24.0, location = 5.0



Table 3: Reliability simulation for a scalable storage unit

Failure Type	Average Number of Failures
Controller Failure	0.9222
Controller's Power Supply (House) Failure	0.1665
Disk Enclosure Failure	0.1089
Disk Enclosure's Power Supply (House) Failure	0.0392
Power Supply (UPS) Failure	2.7017
I/O Module Failure	0.3876
Disk Expansion Module (DEM) Failure	0.906
Baseboard Failure	0.46
Disk Drive Failure	232.2646
Data Unavailability	0.0116
Data Loss	0.0017