

Scaling File Systems to Support Petascale Clusters: A Dependability Analysis to Support Informed Design Choices

Shravan Gaonkar, Eric Rozier, Anthony Tong, and William H. Sanders
University of Illinois at Urbana-Champaign
Email: {gaonkar, erozier2, whs}@uiuc.edu, atong@ncsa.uiuc.edu

Abstract

Petascale computing requires I/O subsystems that can keep up with the dramatic computing power demanded by such systems. TOP500.org ranks top computers based on their peak compute performance, but there has not been adequate investigation of the current state-of-the-art and future requirements of storage area networks that support petascale computers. Dependable scaling of an I/O subsystem to support petascale computing is not as simple as adding more storage servers. In this paper, we present a stochastic activity network model that uses failure rates computed from real logs to predict the reliability and availability of the storage architecture of the Abe cluster at the National Center for Supercomputing Applications (NCSA). We then use the model to evaluate the challenges encountered as one scales the number of storage servers to support petascale computing. The results present new insights regarding the dependability challenges that will be encountered when building next-generation petabyte storage. Furthermore, we provide insight into a new design approach that will enable system designers to integrate the trace-based analysis of parameter values from real system data into their stochastic models to allow informed design choices.

Keywords: Simulation, data analysis, modeling techniques, reliability and availability, storage systems.

1. Introduction

Historically, scientific computing has driven large-scale computing resources to their limits. Towards the end of the current decade, we are likely to achieve petascale computing. While supercomputer performance has improved by over two orders of magnitude every decade, the performance gap between the individual nodes and the overall processing ability of an entire system has widened drastically [13]. This has led to a shift in the paradigm of supercomputer design, from a centralized approach to a distributed one that supports heterogeneity. While most high-performance computing environments require parallel file systems, there have been several file systems, such as GPFS [10], PVFS2 [16], and Lustre [3], that have been specifically proposed to support very large-scale scientific computing environments.

As the number of individual computing resources and components becomes very large, the frequency of failure of components within these clusters and the propagation of these failures to other resources are important concerns to high-performance computing applications. Failures can be caused by many factors: (a) transient hardware faults due to increased chip density, (b) software error propagation due to a large buggy legacy code base, or (c) manufacturing defects and environmental factors such as temperature or humidity.

Recent literature on failure analysis of BlueGene/L discusses

various causes of increased downtime of supercomputers [7]. It has been well-established that elimination of failures is impossible; it is only feasible to circumvent failures and to mitigate their effects on a system's performance. The standard approach to the mitigation of a failure is to checkpoint the application at regular intervals. Long et al., however, showed that checkpointing has a large impact on the performance of very large computer clusters with large numbers of nodes [15].

Increasing the number of compute servers in a cluster almost always increases the size of the desired subsystem. Depending on the type of parallel file system, that means an increase in the number of file servers that could accept requests from the compute servers to keep up with I/O requests. Compute servers and file servers have very different characteristics. First, a failure in a file server needs more attention than a failure in a compute node. A compute server might just be marked as unavailable until it is repaired, but a failed file server might have to be reconstructed, or its state might need to be transferred to another file server, depending on the replication strategy. Second, file servers are inherently slower than compute servers due to their I/O characteristics. This generally makes file servers the bottleneck for the reliability and performance of a cluster. Unfortunately, there has been a trend towards increasing failure rates for I/O subsystems that is similar to that for overall petascale clusters. This increase in failures can be attributed to the increase in the number of individual components that are used to build the whole I/O subsystem. Recent studies have shown that workload intensity is highly correlated to the failure rates [11, 14]. That emphasizes the need for thorough analysis to understand the impact of the I/O subsystems and their failures on petascale computers.

To address the research challenge of providing realistic prediction of petascale file system availability, we took a two-pronged approach. First, we have obtained the failure event log of the Abe cluster from the National Center for Supercomputing Applications (NCSA). The log contains the failures of individual nodes, file server nodes, and the storage area network (SAN). We preprocessed the event logs to determine various reward measures of interest corresponding to the file system, such as the availability of the file system over the lifetime of the log and the failure rate of jobs due to I/O failures and other transient failures. Then, we built and refined stochastic models of the file system used by these clusters that abstracts much of the operations, while generating reward measures that are comparable to the real log events. We then scaled the models to reflect the scale and magnitude of a future petascale computer and estimated the impact of current I/O and file system designs on a petascale computer. Furthermore, we evaluated strategies that could be used to mitigate the bottlenecks due to scaling of I/O file system and cluster designs from current supercomputers to petascale computers. Our analysis will give storage architects

support to make informed design choices as they build larger cluster file systems.

The rest of the paper is organized as follows. Section 2 outlines related work. Section 3 discusses the file system architecture of the Abe cluster at NCSA with the analysis of collected failure log files. Section 4 presents the conceptual stochastic activity network model of the Abe cluster. Section 5 covers results and analysis. Section 6 offers our conclusions and plans for future work.

2. Related work

The estimation and prediction of the failure of file systems are crucial to understanding the overall performance of petascale computers. Past literature describes several attempts to model and analyze different aspects of large-scale supercomputing systems.

Log/trace-based analysis: Recent literature using trace-based system analysis has shown that storage subsystems are prone to higher failure rates than their makers estimate because of underrepresented disk infant mortality rates [12]. In addition to disk failures, the analysis by Jiang et al. shows that interconnects and protocol stacks play a significant role in storage subsystem failure [6]. Furthermore, Liang et al. investigated the failure events from the event logs from BlueGene/L to develop failure prediction models to anticipate future fatal failures [7]. In general, trace-based analysis of logs provide good metrics for evaluating and understanding working systems, but is limited to the scope of events represented by these traces, making it difficult to study trends or behaviors not witnessed in the traces.

Model-based analysis: Wang et al. looked at the impact on system performance in the presence of correlated failures as the systems are scaled to several hundred thousand processors [15]. Rosti et al. presented a formal model to capture CPU and I/O interactions in scientific applications, to characterize system performance [9].

The use of simulation for evaluating the model provides the ability to predict behavior and bottlenecks of future designs, but the accuracy of predictions may often be compromised by assumptions of parameter values. Our approach focuses on integrating trace based analysis with model-based evaluation to form a combined approach, providing guidance to make informed choices for system design. Using failure data from the logs of the cluster as parameter values, we verify our models against the real system. We then analyze the impact of current design choices when the system is scaled. Our approach reduces the burden of sensitivity analysis, reducing the design space to a moderate size providing the opportunity to perform a robust analysis of the system.

3. Abe Cluster: System Configuration and Log File Analysis

The Abe cluster architecture is the current state-of-the-art. Abe consists of 1200 blade compute nodes, i.e., 9600 core CPU Intel 64 (2.33 GHz dual-socket quad-core) processors, 8/16 GB shared RAM per node, and an InfiniBand (IB) interface. The cluster runs Red Hat Enterprise Linux 4 (Linux 2.6.9) as its operating system. The cluster can provide a peak compute performance of 89.47 PFLOPS. The Lustre file system supports a 100TB parallel cluster file system for the Abe cluster’s compute nodes [2].

Lustre-FS outage time			
Cause of Failure	Start time	End time	Hours
I/O hardware	07/21/07 23:03	07/22/07 12:00	12.95
I/O hardware	07/31/07 01:49	07/31/07 20:01	18.18
I/O hardware	08/22/07 18:08	08/23/07 02:15	08.12
I/O hardware	08/28/07 16:20	08/29/07 18:01	01.67
I/O hardware	09/25/07 18:00	09/26/07 09:30	15.50
I/O hardware	10/04/07 09:30	10/04/07 21:55	12.42
Batch system	10/16/07 17:56	10/16/07 21:24	03.47
Network	10/29/07 11:53	10/29/07 15:15	03.36
File system	11/16/07 09:30	11/16/07 10:00	00.40
File system	11/19/07 09:04	11/19/07 11:00	01.93

Table 1. User notification of outage of the Lustre-FS

3.1. General Cluster File System (CFS) Architecture

A typical storage architecture for a cluster file system consists of a *metadata server*, multiple *file servers*, and *clients* [3]. The metadata server maintains the file system’s metadata, which includes the access control information, mapping of files and directory names to their locations, and mapping of allocated and free space. The metadata server serves the metadata to the clients. The file servers maintain the actual data and information about the file blocks stored on the connected I/O disks and serve these file blocks to the clients. For reliability/performance, the file blocks can be replicated/striped over multiple disks. The client communicates first with the metadata server and then with the appropriate file server to perform the required read and write operation. The readers are referred to [3] for further details.

3.2. Abe CFS Server Hardware

The Abe Lustre-FS is currently supported by 24 Dell dual Xeon servers that provide 12 fail-over pairs¹. One OSS serves the metadata of the Lustre-FS, 8 OSSes serve the /cfs/scratch OSS, and the remaining 6 servers handle the remaining partitions of the shared file systems (home, local, usr, etc.) of the cluster. Each server self-monitors its file system’s health. The 2 metadata OSSes are connected to the storage I/O through a dual 2Gb fiber channel (FC).

3.3. Abe CFS Storage Hardware

Scratch partition: 2 S2A9550 storage units from DataDirect Networks Systems provide the storage hardware for the CFS’s scratch partition. Each S2A9550 supports 8 4Gb FC ports. Each port connects to 3 tiers of SATA disks. Each tier has (8+2) disks in a RAID6 configuration. Therefore, there are 480 disks, each with a 250GB capacity, that form the scratch partition providing 96TB of usable space.

Metadata: DDN EF2800 provides the I/O hardware to support the metadata of the Lustre-FS. It is connected to the 2 metadata OSSes through a dual 2Gb fiber channel. The EF2800 has one tier of 10 disks in RAID10 configuration.

¹We refer to a fail-over pair as an OSS in the remainder of the paper.

Date	#	Date	#	Date	#
07/03/07	102	07/19/07	258	08/16/07	375
08/20/07	591	09/05/07	005	09/17/07	002
09/18/07	004	09/19/07	003	09/28/07	463
09/29/07	477	10/01/07	051	10/02/07	035

Table 2. Lustre mount failure notification by compute nodes from 07/01/07 to 10/02/07; column with “#” represents the number of compute nodes that experienced mount failure

Total jobs submitted between 05/13/07 and 10/02/07	44085
Total failures due to transient network errors	1234
Total failures due to other/file system errors	0184

Table 3. Job execution statistics for the Abe cluster

Other partitions: 10 IBM DS4500s serve an approximate total of 40TB of usable space over a SAN via a 2Gb FC.

Lustre settings: Lustre version 1.4.10.X runs on all of the OSS's hardware. Most of the reliability is provided by the SAN hardware; therefore, the Lustre reliability features are switched off.

3.4. Abe Log Failure Analysis

All NCSA clusters have elaborate logging and monitoring services built into them. The log data set used in this study was collected from 05/03/2007 to 10/02/2007 for compute nodes (compute-logs) and from 09/05/2007 to 11/30/2007 for the SAN (SAN-logs). The compute-logs and SAN-logs are monitored precisely, and the logs provide details about the events taking place in the cluster. Events are reported with the node IP addresses and the event times appended to the log information. To extract accurate failure event information, we filter failure logs based on temporal and causal relationships between events.

Table 1 provides the availability of the Abe cluster based on the notifications provided by the SAN administrators to the users [1]. The availability of Abe's SAN can be estimated to be between 0.97 and 0.98 depending on the dates one chooses as the start and end times for the measure computation. Table 2 shows Lustre-FS mount failures experienced by individual compute nodes aggregated on a per-day basis. Lustre-FS mount failures do not always imply the failure of the CFS, as these errors could be caused by intermittent network unavailability. Nevertheless, those errors are perceived as failures from the cluster's perspective.

Table 3 presents the job failure/completion statistics obtained by analyzing the compute-log. The analysis shows that the transient errors causing network unavailability (between the compute nodes and the CFS or between the compute nodes and the login nodes) are 5 times more likely to cause job failures than other errors are (such as software errors or CFS failures). Earlier clusters had dedicated backplanes connected to compute nodes to provide communication. Current communication in Abe is through COTS network ports and switches. The change in the design choice was motivated chiefly by a desire to lower costs and increase flexibility in maintaining the system.

Table 4 provides the disk failure and replacement log from 09/05/2007 to 11/28/2007 for disks that support the scratch partition of the Abe's cluster. The authors of [12] estimated the disks' hazard rate function to be statistically equivalent to a Weibull distribution. We performed similar survival analysis on the disk failure data and found that Weibull with $\beta = 0.7$ was a good fit for Abe's disk drive failure logs. The key insights we gained from analyzing failure data and from discussions with cluster system administrators are as follows:

- The disk replication redundancy and replacement have been so well-streamlined that they almost never cause catastrophic failure of the CFS. On average, 0-2 disks are replaced on the Abe cluster per week.
- The Abe cluster's S2A9550 RAID6 (8+2) technology combines the virtues of RAID3, RAID5, and RAID0 to

Dates in September 2007	05	06	09	13	23
Number of failed disks	2	1	1	1	1
Dates in October 2007	08	17	24		
Number of failed disks	2	1	1		
Dates in November 2007	08	17			
Number of failed disks	1	1			

Survival analysis of the disk failures ($n = 480$) using Weibull regression (in log relative-hazard form) gives the shape parameter as 0.6963571 with standard deviation of 0.1923109 (95% confidence interval) [5]

Table 4. Disk failure log from 09/05/2007 to 11/28/2007 for disks supporting the Abe's scratch partition

provide both reliability and performance [8]. RAID6 prevents a second drive failure from occurring during disk re-mirroring. The *Blue Waters* petascale computer, which will be built at the University of Illinois, will likely have an (8+3) RAID configuration. That would make the failure of the file system due to multiple individual disk failures highly unlikely.

- Most file system failures are due to software errors, configuration errors, and other transient errors. The software errors take, on average, 2-4 hours to resolve. Most often, the fix is to bring the disks to a consistent state using a file system check (*fsck*). A hardware failure due to a network component or a RAID controller might take up to 24 hours to resolve, as these components need to be procured from a vendor.

4. Stochastic Activity Network Model: Cluster File System

The failure data analysis and the insights provide the details necessary to build a stochastic model of the Abe's cluster file system. Here, we describe the details of the stochastic activity network models using Möbius [4].

4.1. Overall Model

Figure 1 shows the *composed model* of the Abe cluster using replicate/join composition in Möbius. The leaf nodes in the replicate/join tree are stochastic activity network models that implement the functionalities. Space limitations do not permit detailed descriptions of these submodels. The CLUSTER model has two main submodels connected using a join, where the models share states on error propagation from their CLIENT to the CFS. The CLIENT represents the behavior and interaction of the compute nodes and the communication network between the compute nodes and the CFS. The CFS_UNIT emulates the Abe's cluster file system. It is composed of the OSS, OSS_SAN_NW, SAN, and the DDN_UNITS. The OSS implements the availability and operational model of the metadata server and the file server. The OSS_SAN_NW implements the failure model of the network ports and switches that connect OSS to the DDN_UNITS. The SAN emulates the operations provided by the network to communicate between OSS and the DDN_UNITS. The OSS, OSS_SAN_NW, SAN, and the DDN_UNITS communicate by sharing information about their current state of operation and availability. The DDN_UNITS composes multiple RAID6_UNITS with RAID_CONTROLLER. The failure of disks in RAID6_UNITS is assumed to follow a Weibull distribution. RAID_CONTROLLER emulates the failure and operation

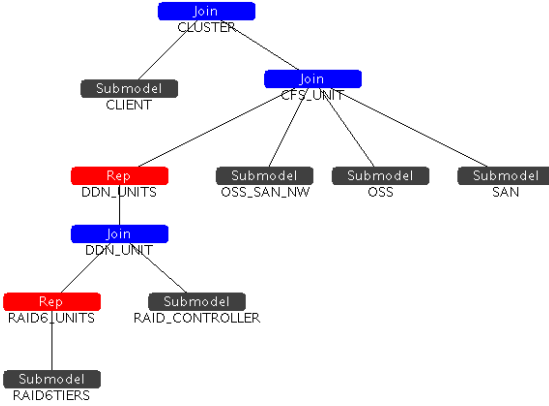


Figure 1. Compositional model of the CFS.

of a typical RAID6 architecture. The DDN_UNITS is replicated to emulate multiple S2A9550 units.

Since the goal is to investigate the impact of availability of file systems to petascale computers, the stochastic activity network models do not consider hardware failure in compute nodes. Our model incorporates only the behavior of the scratch partition and the metadata servers of the CFS, because a cluster’s utility depends mainly on its scratch partition’s availability. Finally, hardware and software misconfiguration errors occur in the early deployment phase of the system; therefore, we exclude them from the models. In the following subsections, we describe the reward measures and the failure model used to represent the Abe’s CFS.

4.2. Reward Measures

The **availability** of the cluster file system is defined as the ability of the CFS to serve the client nodes. More precisely, it is defined as the fraction of time when all the file server nodes (OSSes), the DDN, and the network interconnect between the OSSes and the DDN are in the *working* state.

The **disk replacement rate** is defined as the number of disks that need to be replaced per unit of time to sustain the maximum availability of the CFS.

The **cluster utility, CU**, is an availability metric from the cluster’s perspective. To be precise, it is defined as $CU = \left(1 - \frac{\text{Compute cycles lost due to unavailable file system}}{\text{Total available compute cycles}}\right)$. CU is a metric that is different from the availability metric of the CFS². The cluster users and SAN administrators tend to notice different levels of availability. The reasons are failures in network communication between the compute nodes and the CFS as well as failures due to intermittent transient errors that make CFS appear unavailable even though it has not failed.

4.3. Failure Model for Abe’s CFS

The Abe’s cluster suffers from failures mainly because of 3 types of errors: hardware errors, software errors, and transient errors. Each kind of error affects all the CFS’s components.

The **hardware errors** in the metadata/file servers (OSSes) occur in the hardware components that are built to operate the system. These errors include processor, memory, and network errors. Hardware errors are assumed to be less frequent than

²CU does not distinguish between compute cycles used to perform checkpointing and those used for actual computation.

Model parameter	Values (range)
Disk MTBF ²	100000-3000000
Annualized Failure Rate (AFR)	0.40%-8.6%
Weibull distribution’s shape parameter ¹	0.6-1.0
Number of DDN ¹	2-20
Number of compute nodes ¹	1200-32000
Average time to replace disks ³	1-12 hours
Average time to replace hardware ³	12-36 hours
Average time to fix software ³	2-6 hours
Job request per hour ¹	12-15 per hour
Hardware failure rate ¹	1-2 per 720 hours
Software failure rate	1-2 per 720 hours
Annual growth rate of disk capacity ²	33%
DDN_Units ¹	2-20
OSS Units ¹	8-80

Parameter values obtained from: log file analysis¹, data specification from literature and hardware white papers², discussions with NCSA cluster administrators³

Table 5. Abe cluster’s simulation model parameters

disk failures, occurring at the rate of 1-2 per month. The RAID controllers in the DDN or network ports/switches that connect DDN to OSS show similar failure rates. The repairs of these components take 12-36 hours depending upon the severity of the failure (as reported by SAN administrators), as the needed replacement parts have to be shipped from the vendors. Most of the hardware is replicated with fail-over mechanisms. Failure of both members of the fail-over pair causes the unavailability of the CFS system. The replacement of failed disks is modeled as a deterministic event. The repair time is varied from 1 to 12 hours across simulation experiments.

The **software errors** that cause failure of the cluster file systems are mainly due to the corrupted supercomputing applications running on the compute nodes (implemented in the CLIENT submodel) or the Lustre-FS (implemented in the OSS submodel). Since we do not have accurate estimates on software corruption errors, we assume that the rates are similar in the orders of magnitude to hardware error rates. The repair times for software errors are modeled as deterministic events. The repair time is varied from 2 to 6 hours across simulation experiments.

Transient errors occur in most components of the cluster model, but mainly in the network components. The error rates are obtained from the failure-log analysis as shown in Table 3. Transient errors are temporary, but hard to diagnose. Our model assume that one of these errors causes a few minutes of unavailability of components under transient failure. The jobs depending on those components fail due to the temporary unavailability.

Past literature has emphasized the importance of modeling **correlated failures** [14]. Most correlated errors occur because of shared resources. Correlated errors propagate to components that have causal or spatial proximity. In the CFS model, hardware errors propagate because other hardware components are connected to each other. Software errors propagate from compute nodes to OSS or from OSS to disk, leading to data corruption. Transient errors propagate errors into software. All failures except disk failures are modeled as exponential distributions. To model correlated failures, we model jobs and requests submitted to the CFS, and estimate the probability p that the job requires a resource that is inaccessible due to failure, causing errors to propagate through the system.




Figure 2. Availability of storage with respect to disk failures; Label with values (0.7,2.92,8+2,4) represents a tuple =(Weibull shape parameter β , AFR in %, RAID configuration, average disk replacement time in hours)

5. Experimental Results and Analysis

We evaluate the design of the Abe cluster’s availability using simulation in the Möbius tool. Table 5 summarizes the parameters collected through failure log analysis, hardware reliability specifications, and discussions with cluster administrators. In order to reflect the size and scale of a petascale computer and to determine the factors that impede the high availability of the CFS, we scale the number of individual components in the composed model. By implementing scaling through the addition of components (each of which has its own individual failure models) rather than by changing failure parameter values themselves, we ensure that the failure rates observed in the overall system model accurately reflect the new system size. All the simulation results are reported at a 95% confidence level.

5.1. Impact of Disk Failures on CFS

To evaluate the baseline effect of failures of disks on availability of the CFS, we evaluate the DDN_UNITS models associated with the RAID6 tiers and the RAID controllers in isolation from failures of other components of the SAN. Figure 2 shows the availability of the storage hardware as one scales the file system from the current 96TB (Abe’s file system) to 12PB (the Blue-Waters file system). The key observation is that the RAID6 architecture provides sufficient redundancy and recovery mechanisms to mitigate the impact of high disk failure rates to a very large extent. First, note that all configurations of failure and recovery rates for an Abe-sized cluster file system have nearly 100% availability (refer to the first data point in Figure 2). However, as the experiments are scaled from Abe’s system to a petascale system, our simulation results show that the RAID6 architecture cannot provide the same level of storage availability for some of the failure model configurations. The SAN architect’s plan to use (8+3) RAID in Blue Waters is important; it provides better reliability than the (8+2) RAID on petascale systems. While RAID6 provides a larger margin for disk failure rates, i.e., up to 8.6% AFR, it is very important that these rates be contained to lower thresholds by disk manufacturers and vendors to provide the adequate level of availability. If one makes a pessimistic assumption of a higher infant mortality rate in disks (Weibull shape parameter = 0.6), the availability

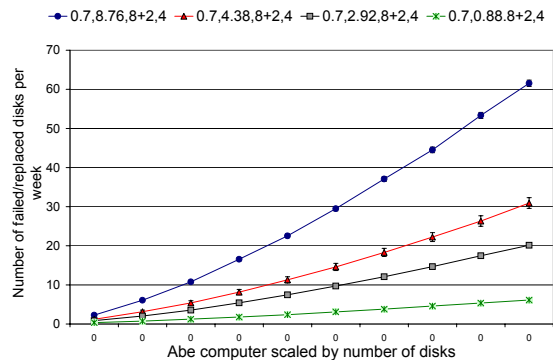


Figure 3. Average number of disks that need to be replaced per week to sustain availability

falls below 99.9% for petascale storage.

To better understand the cost of disk replacement, we compute the expected number of disks that need to be replaced per week for the RAID6 tiers. Figure 3 depicts the average number of disks that need to be replaced per week to sustain the availability so that the CFS does not suffer failure due to RAID6 failure. The configuration (0.7,2.92,8+2,4) corresponds to the Abe cluster with 0 to 2 disk replacements per week. Each time a disk fails, there is an operational cost (in dollars) that is borne by the SAN vendors as they provide extended support to their SANs. As the CFS system is scaled to support petascale computers, the number of disks that need to be replaced increases, increasing the labor cost and the replacement cost. Therefore, the SAN vendors have an incentive to increase the disk MTBF to reduce their overall support cost.

Survival analysis of the disk failure data provided a good estimate of the Weibull distribution’s shape parameter β , but the estimate for the scale parameter (MTBF) was insignificant [5]. Using simulations, we estimated an MTBF that matched the average disk failures per week for the scratch partition and determined that an MTFF=300,000 hours or an annualized failure rate (AFR) = 2.92% to be a good fit.

5.2. CFS Availability and CU

To analyze the impact of all components that determine the availability of the CFS, we evaluated the availability and CU of the Abe system. The experiments are scaled to the size of a petascale computer to allow understanding of the impact of failures on those measures. Figure 4 shows that the CFS availability decreases as one scales the system to support a petascale computer. Since most of the parameter values were obtained through the log data analysis and times reported by SAN administrators, our measures for CFS availability matched with Abe’s availability as shown in Table 1. Therefore, we have higher confidence in the measures of availability and CU as we scaled the models to represent a petascale computer with a petabyte storage system. The storage availability in Figure 4 refers to configuration (0.7,2.92,8+2,4), which models the Abe cluster’s current environment. We find that the RAID6 subsystem in this configuration continues to provide an availability of 1, but the CFS availability is reduced from 0.972 to 0.909 as one scales the design to support the petascale system. The reduction is mainly due to correlated failures in OSS and hardware. Improving upon Abe cluster’s design, the architect could provide




Figure 4. Availability and utility of the Abe cluster when scaled to petaflop-petabyte system

an additional standby or spare OSS to replace the failed OSS quickly. Our evaluation shows that this approach can improve the availability by 3%. To improve the availability further, the architects have to develop solutions to mitigate correlated errors. For example, improving the robustness of the Lustre-FS can reduce the software-correlated errors. The CU in Figure 4 shows that the cluster's network architecture between the compute nodes and the CFS has a profound impact on the cluster utility available to the users. The trend to move away from customized backplanes to COTS network hardware (with its complicated software stacks) has decreased the CU. The transient errors seen in the network can be mitigated by providing multiple network paths between the compute nodes and the CFS.

6. Conclusion and Future Work

Many researchers have focused on developing and understanding dependability of clusters for supercomputing applications. In our paper, we have taken steps to understand the dependability and availability of the ABE cluster through failure data analysis and discussions with administrators at multiple levels of the cluster operation, from the lowest level of the SAN's availability, to the cluster's availability, and to user perception of cluster utility at the top level. Our key findings through analysis and simulation showed that the RAID6 design for a disk's dependability has limited the impact of disk failures on the CFS, even when the model is scaled to evaluate the support for petascale systems. On the other hand, transient errors, hardware errors, and software errors contribute significantly to failures, and these components are the limiting factors for the high availability of the CFS. We believe that petascale architects will have to focus on those issues to develop solutions to improve the overall availability of the CFS.

Future work: Our work has mainly focused on evaluating the availability of the ABE's CFS through data collection, analysis, and system modeling. We showed that system modeling, combined with data analysis from real systems, provides better insight to design future systems. NCSA has other operational clusters, such as Mercury and Tungsten, with different architectures. We may evaluate and analyze these systems to determine how the ABE architecture was designed based on the lessons learned from the older clusters. That analysis could provide insight into future CFS architectures on petascale systems. The

models and data analysis (on larger failure data sets) can be extended to evaluate performance metrics that would complement the dependability measures conducted in this research.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant Number CNS-0406351 and a generous gift from HP Labs. The authors would like to thank the security, storage and other affiliated groups at NCSA for their inputs, log files data and time that enabled us to consolidate this practical report. We would also like to thank Jenny Applequist for her editorial comments.

References

- [1] Notifications to teragrid users about CFS unavailability <http://news.teragrid.org/user.php?cat=nscsa>, november 2007.
- [2] NCSA intel 64 linux cluster abe technical summary <http://www.nscsa.uiuc.edu/UserInfo/Resources/Hardware/Intel64Cluster/TechSummary/> july 9, 2007.
- [3] R. L. Braby, J. E. Garlick, and R. J. Goldstone. Achieving order through CHAOS: The LLNL HPC Linux Cluster Experience. In *The 4th International Conference on Linux Clusters: The HPC Revolution 2003*, San Jose, CA, USA, 2003.
- [4] G. Clark, T. Courtney, D. Daly, D. Deavours, S. Derisavi, J. M. Doyle, W. H. Sanders, and P. Webster. The Möbius modeling tool. In *Proceedings of the 9th International Workshop on Petri Nets and Performance Models*, pages 241–250, 2001.
- [5] D. W. Hosmer and S. Lemeshow. *Applied Survival Analysis: Regression Modeling of Time to Event Data*. John Wiley & Sons, Inc., New York, NY, USA, 1999.
- [6] 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. In *File And Storage Technologies, FAST*, pages 111–125, 2008.
- [7] Y. Liang, Y. Zhang, A. Sivasubramaniam, M. Jette, and R. Sahoo. BlueGene/L Failure Analysis and Prediction Models. In *Dependable Systems and Networks, DSN*, pages 425–434, 2006.
- [8] L. McBryde, G. Manning, D. Illar, R. Williams, and M. Piszczek. Data Management Architecture. US Patent number 7127668, Oct 24 2006.
- [9] E. Rosti, G. Serazzi, E. Smirni, and M. S. Squillante. Models of parallel applications with large computation and I/O requirements. *IEEE Transaction in Software Engineering*, 28(3):286–307, 2002.
- [10] F. Schmuck and R. Haskin. GPFS: A shared-disk file system for large computing clusters. In *File And Storage Technologies, FAST*, pages 231–244, 2002.
- [11] B. Schroeder and G. A. Gibson. A large-scale study of failures in high-performance computing systems. In *Dependable Systems and Networks, DSN*, pages 249–258, 2006.
- [12] 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 *File And Storage Technologies, FAST*, pages 1–16, 2007.
- [13] E. Strohmaier, J. J. Dongarra, H. W. Meuer, and H. D. Simon. The marketplace of high performance computing. *Parallel Computing*, 25(13–14):1517–1544, 1999.
- [14] D. Tang and R. K. Iyer. Dependability measurement and modeling of a multicomputer system. *IEEE Transactions on Computers*, 42(1):62–75, 1993.
- [15] L. Wang, K. Pattabiraman, Z. Kalbarczyk, R. K. Iyer, L. Votta, C. Vick, and A. Wood. Modeling coordinated checkpointing for large-scale supercomputers. In *Dependable Systems and Networks, DSN*, pages 812–821, 2005.
- [16] W. Yu, S. Liang, and D. K. Panda. High performance support of parallel virtual file system (pvfs2) over quadrics. In *International Conference on Supercomputing, ICS*, pages 323–331, 2005.