Availability Impact of Hard and Soft Failures in Enterprise Storage Systems

Hossein Asadi, Alireza Haghdoost, Reza Eftekhary
Department of Computer Engineering
Sharif University of Technology, Tehran, Iran
asadi@sharif.edu, haghdoost@ce.sharif.edu, eftekhary@ce.sharif.edu

Abstract

Enterprise storage systems have been an integral part of any high-end computing systems. Availability of storage systems is essential to timely provide requested data to servers. Any service disruption, caused by either data loss or data unavailability, can be very costly for enterprise applications such as e-banking or e-shopping. Component failure is one of major causes of data loss or data unavailability. Although data loss can be prevented by data protection techniques such as remote mirroring, snapshots, and backups, these techniques have been less effective to avoid data unavailability. In this work, we investigate the impact of hard and soft failures of processor cores on the overall availability of enterprise storage systems. We use an analytical technique based on Markov model to estimate number of system downtime events.

1 Introduction

Information systems or data storage systems have become an important part of any high-end *Information Technology* (IT) computing system. Data storage systems are now widely being used in IT infrastructure of enterprise applications such as online financial transactions, e-shopping, e-banking, and online travel booking. Data unavailability or data loss in these systems can have expensive, or even catastrophic, consequences [1, 2]. Although data loss can be prevented by data protection techniques such as local snapshots, incremental archives, and *Redundant Array of Inexpensive Disks* (RAID) arrays, these techniques are inefficient to avoid data unavailability. Recent system failures or outage of some e-commerce and communications companies such as *eBay* or *T-mobile* show that a system down time can cost as much as few million dollers per hour or even more [3].

In this work, we investigate the impact of hard and soft failures of processor cores running in both front-end and backend logic of a high-end storage system. We utilize an analytical technique based on Markov model to predict number of data unavailability events during life-time of a high-end storage system.

2 A High-End Data Storage System Model

To satisfy increasing demand for high-speed data transactions, enterprise storage systems have extensively become more complex and have included ever growing integrated circuits. A typical enterprise storage system includes few hundred disk drives in the disk subsystem, a Tera-scale memory subsystem , and more than hundred processor cores running to transfer customer data from the disk subsystem to the servers connected directly or indirectly to the storage system. If customer data not exist in the memory subsystem, it is transferred from the disk subsystem to the memory subsystem, and from the memory subsystem to the storage front-end logic, and from the front-end logic to the servers. Any disruption in the disk subsystem, the memory subsystem, or the front-end logic can delay data transfer and cause data unavailability for host applications running on the servers. A block diagram of a high-end storage system is illustrated in Fig. 1.

As the number of components and computing resources within a storage system increases, the frequency of failure of components within storage systems increases as well. Although most components and computing resources within a high-end storage system use mirrored configuration, a recent study shows increasing number of hard and soft failures in both midrange and high-end storage systems [4].

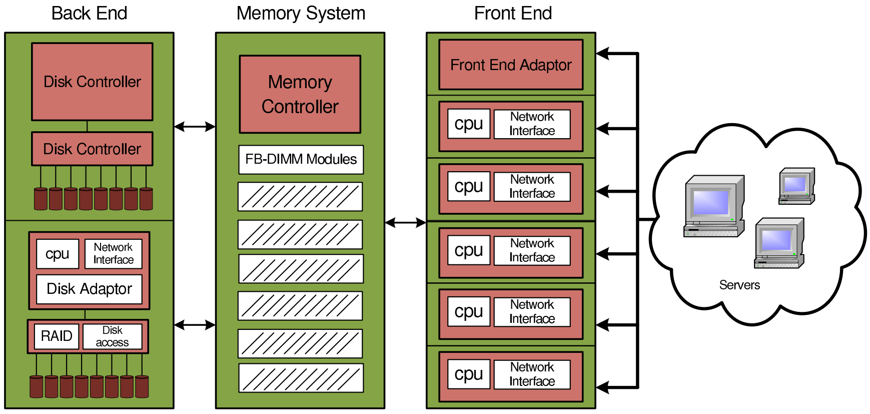

Figure 1. Block diagram of a high-end storage system

3 Root Cause of System Downtime and Proposed Approach

Recent literatures on failure analysis of a supercomputer discuss various causes of increased downtime of supercomputers [5, 6]. In this work, we further investigate failures causes in high-end storage systems. In particular, possible failures in storage systems can be caused by following factors: (a) transient hardware faults due to single event upsets, (b) permanent hardware faults due to chip wearout, (c) hardware failures due to manufacturing or design defects, (d) software failure due to a buggy code, (e) cooling system failure due to mechanical or electonic component failure (f) environmental conditions such humidity or power outage, or (g) human error during system maintenance.

In addition to the above-mentioned failure types, there are other types of failure occurring outside of storage system that can lead to overall service downtime. Examples of such failures are network failure and application failure. We will exclude such failure in our analysis as it is not related to the storage system itlsef.

In particluar, we are working to investigate the impact of hard and soft failures on the overall availability of high-end sotrage systems. It has been shown in [4] that in high-end storage systems 32% of system down-time is caused by hardware components failures. Additionally, it is reported that 15% of system down-time is due to soft errors caused by energetic particles such as neutrons or alpha particles [4]. As device minimum feature sizes continue to shrink, CMOS processing devices and storage devices such as microprocessors and SRAM will experience an increasing number of soft errors [4]. These bit flips can temporarily corrupt the data being stored or processed. In the past, soft errors were regarded as a major concern only for space applications. Yet, for designs manufactured at advanced technology nodes, such as 65 nm and smaller, system-level soft errors are much more frequent than in the previous generations.

References

[1] S. Gaonkar, K. Keeton, A. Merchant, and W.H. Sanders. Designing dependable storage solutions for shared application environments. In *IEEE International Conference on Dependable Systems and Networks*, pages 371–382, June 2006.

[2] S. Gaonkar, K. Keeton, A. Merchant, and W.H. Sanders. Designing dependable storage solutions for shared application environments. *IEEE Transactions on Dependable and Secure Computing*, 2009.

[3] Debbie Lewis. Partner techtip: The cost of downtime. Online page, http://www.mcpressonline.com/tips-techniques/high-availability-disaster-recovery/partner-techtip-the-cost-of-downtime.html, November 2006.

[4] S.Z. Shazli, M. Abdul-Aziz, M.B. Tahoori, and D.R. Kaeli. A field analysis of system-level effects of soft errors occurring in microprocessors used in information systems. In *IEEE International Test Conference*, pages 1–10, October 2008.

[5] Y. Liang, Y. Zhang, A. Sivasubramaniam, M. Jette, and R. Sahoo. Bluegene/l failure analysis and prediction models. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 425–434, June 2006.

[6] S. Gaonkar, E. Rozier, A. Tong, and W.H. Sanders. Scaling file systems to support petascale clusters: A dependability analysis to support informed design choices. In *IEEE International Conference on Dependable Systems and Networks*, pages 386–391, June 2008.