A Mass Storage System Administrator Autonomic Assistant

Milton Halem
University of Maryland Baltimore County
halem@umbc.edu

Randy Schauer
Raytheon Company
rschauer@arl.army.mil

Abstract

System administrators of today's high performance computing systems are generally responsible for managing the large amounts of data traffic and archival querying that mass storage systems must provide to users who compute on hundreds or thousands of processors at once. The file management systems that have been engineered to handle this workload generally consist of a reliable compute server, high- and low-speed disks, robotic tape silos with thousands of cartridges, and various network interconnects. As a result of the wide variety of mechanical components used, storage system administrators maintain the functional aspects of system operations and troubleshoot the day-to-day lower level physical system and software failures under severe constraints. The Mass Storage System Administrator Autonomic Assistant (MSSAAA) is being developed to reduce the burden on storage system administrators and improve the reliability, availability and serviceability of mass storage systems.

1. Introduction

A mass storage system being utilized at a high performance computing (HPC) center typically can become a single point of failure for the entire suite of resources if a major problem renders it unavailable. Finding the root cause of a problem in a mass storage system is not trivial. In addition to the physical compute server, the current generation of mass storage solutions also includes high- and low- speed disks, a large robotic tape silo, and various network interconnects that service each component. This design is further complicated by including the software installed to drive each of these peripherals and the possibility of problems existing on the HPC resources that have mounted various file systems from the storage server.

The administration of these storage systems is growing with the complexity. At many HPC centers, dealing with petabytes of data to archive, backup and manage will be the norm in the next few years, if it is not already. Estimates have been made that the average mass storage installation has one system administrator per ten terabytes of data [2]. This number should be evolving with time since the density of storage media has been growing at an average rate of ten times every five years [5]. Regardless, this number is extremely alarming as HPC centers head toward Petaflop computing, where the online data storage needs will grow, and the need to archive these larger amounts of data for longer periods of time. The march toward Petaflop computing will drive the growth of mass storage technologies over the next decade [1].

Reliability, availability and serviceability (RAS) is of growing interest to HPC centers among all types of systems as Teraflop computing becomes common and Petaflop computing exists on the horizon. In [4], the National Coordination Office for Information Technology Research and Development has made RAS one of the broad areas of focus for the coming years. Mass storage systems must become more self-aware to achieve this goal. This self-awareness should be able to identify problems as they occur and implement corrective actions.

The solution and prototype described in the following sections are based on research done with storage system administrators at the U.S. Army Research Laboratory Major Shared Resource Center and the NASA Center for Computational Sciences at the Goddard Space Flight Center. Conversations and observations of their actions provide much insight into the benefits of an autonomic tool to monitor specific aspects of the system and automatically handle many issues that come up on a daily basis.

2. Proposed Solution

This growing complexity leads us to introduce the Mass Storage System Administrator Autonomic Assistant (MSSAAA). Its primary goal is to reduce the low-level tasks that system administrators deal with the most frequently. The tasks selected include those that require manual intervention and those that could be monitored to improve performance of the entire storage system. The assistant's success with these tasks will make the

Proceedings of the Second International Conference on Autonomic Computing (ICAC'05)
0-7695-2276-9/05 $ 20.00 **IEEE**

administrators more efficient and allow them to better handle higher-level issues that arise. The implementation makes use of the IBM Autonomic Toolkit, in particular the Generic Log Adapter and Autonomic Monitoring Engine.

3. Prototype Design

The MSSAAA is being developed to provide a framework for automating low-level tasks. The prototype consists of six modules designed to improve the reliability, availability and serviceability of mass storage systems.

A module to measure and analyze file transfer performance over a period of time can be used to determine if conditions exist for users to experience problems with the storage system. This performance is measured by retrieving a specific file from tape across the network from a client Network File System (NFS) mount. The results can be used to make the administrators aware of potential performance issues as they transpire instead of waiting for a user to report the problem.

A module to perform an audit on a tape when it returns an error status can improve tape availability. Many tapes are placed into an error status when no problem actually exists. Administrators spend a good deal of time auditing these tapes just to find there is no problem and then place them back into production. By auditing tapes when they report an error status and filtering out the tapes with no errors versus the tapes with actual physical problems, administrators are sure every error they see is actually a problem.

A module that automates the assessment of damaged files on a tape can give system administrators the level of detail they need to determine the amount of data loss and media damage. If a problem is discovered with a tape during the staging process of a file, the other files on that tape should be staged automatically. The files that report damage during the staging of the tape should then be undamaged and restaged in an attempt to solve the problem without involving the administrators. Only those files that continue to be damaged should be reported. Over time, the amount of files reported for assessment would be a fraction of those initially reporting damage.

A module to automate the movement of data between Sun's Storage and Archive Management (SAM) instances can save time and work in the background while other tasks are executing on the storage systems. If the storage systems could gauge the loads on each system, they could keep a continual flow of data moving between them and increase or decrease the flow as their loads fluctuate. By keeping the systems aware of their environment and granting them the autonomy to handle this function, the administrators would only have to become involved during the process if major problems occur.

A module to develop autonomic migration policies will be able to judge when data should be moved to another archive medium to facilitate the best environment for decade or century long preservation. Data should be monitored in such a way to realize when it has met criteria for archive preservation and should be migrated to the assigned storage media. This will eliminate the need for administrators to perform periodic checks on existing data to see when it should be moved onto long-term tape storage.

A module that consolidates statistics from different storage system architectures will give administrators a high level assessment of system performance. When reviewing logs for potential problems and to analyze metrics, administrators would benefit from already having a consolidated log in place. This data could also be used to determine issues that occur regularly and may benefit from being added to the assistant.

4. Conclusions

The complexities of mass storage systems are growing at a rapid pace to meet the needs of users who require petabyte-scale data storage. Storage systems are required to provide optimal levels of stability and efficiency at all times. Having a self-management solution means that system administrators will not continue to be called on during nights and weekends to troubleshoot tasks that are now solved autonomically [3].

An effective MSSAAA has the potential to reduce system support costs while improving system reliability and enabling system administrators to focus on higher-level tasks. System administrators will be able to alter policies on the system in a fraction of the time it would take to modify a script or set of scripts to accomplish the same task. The looming crisis of managing data intensive distributed computational resources with constrained budgets for human resources is already a problem facing HPC centers, and autonomic assistants such as MSSAAA can contribute to a solution.

5. References

[1] Dongarra, J. J. and D. W. Walker. "The Quest for Petascale Computing." Computing in Science and Engineering, 32-39, May/June 2001.

[2] Gray, J. A conversation with Jim Gray. ACM Queue, 1(4). ACM, June 2003.

[3] Kephart, J. O. and D. M. Chess. "The Vision of Autonomic Computing." Computer. 36. 1, 41-50, 2003.

[4] National Coordination Office for Information Technology Research and Development, "Guide to the NITRD Program FY 2004 – FY 2005," 2005.

[5] Thompson, D. A. and J. S. Best, "The future of magnetic data storage technology," IBM Journal of Research and Development, 44, 311-322, May 2000.