Reliable Long Term Archiving Storage Architecture

Keith Rajecki; Sun Microsystems, Inc.; Menlo Park, CA/USA
Michael Selway; Sun Microsystems, Inc.; Middleton, MD/USA
Brian Parks; Sun Microsystems, Inc.; Calgary, Alberta/Canada

Abstract

This paper is intended to address the challenges of designing a reliable long term archiving storage architecture regardless of types of digital objects needing to be stored. While the digital object types and properties should be considered when determining the storage components of the architecture, the object characteristics are not as important as the access and preservation requirements.

The first step to building a long term archive storage architecture is the assessment of the business processes. Incorporating the business process into your architectural design is crucial to the overall success of the long term archive. The only way to ensure an appropriate architecture is to understand the business processes. Documenting your organizations policies and procedures including data types, length of archive, access methods, maintenance activities, and technical specifications will increase the probability your archive architecture will meet the business requirements.

Whether you are building your archive for historical preservation or to store data for business compliance, a tiered storage architecture can provide you with the most reliable and cost effective solution. Regulations often require that information be located and retrieved very quickly. If architected incorrectly, data searching and retrieval can be time consuming and costly. Traditional tape only archival methods simply can not meet the access requirements of many of today's repositories and long term archives. Likewise, storing all your data on disk requires greater administration and is more costly. The proposed tiered storage architecture provides a balance between disk and tape storage hardware to support long term archiving.

A reliable long term archive is also dependant on the software components being open and supporting interoperability. Storing, searching, and retrieving data is not sufficient criteria for a successful long term archive. A long term archive should incorporate open source standards based software to ensure future support.

The overall storage system architecture addresses the physical storage components and processes for long-term preservation. Key components to address when architecting your long-term archive are security, storage, and application interoperability. The security layer focuses on the data access in order to ensure integrity and privacy. Storage addresses the placement of the objects within the various hardware components based on retention policies. Application interoperability is the systems and applications ability to be backward compatible as well as the ability to support expanded system functionality.

Open Archive Information System

While the Open Archive Information System (OAIS) model has become the de facto standard for preservation archives, the design and implementation of a reliable long-term archive lacks adopted technology standards and design best practices. The proposed tiered storage architecture provides a viable, cost effective, and reliable long-term archiving storage system. This architecture is based on a combination of open source and commercially supported software and systems. An open source environment is likely to provide long-term viability.

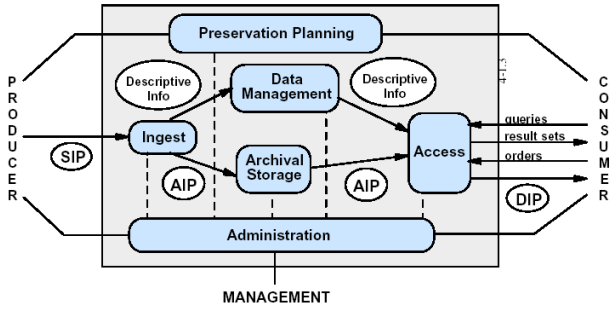

[1] Reference Model for an Open Archival Information System

Digital Repository Solutions

Fedora, Dspace, Eprints, Vital

The term digital repository is widely debated by some. For the purposes of this solution architecture, digital repository refers to the the system by which objects are stored for preservation archiving. There are a number of viable repository solutions available that provide the capability to store, manage, re-use and curate digital materials. Repository solutions support a multitude of functions and can be internally developed or extended. These repository solutions were highlighted for their ability to integrate into a tiered storage architecture and their interoperability. The repositories must be sustainable and supportable in order for the underlying storage system to operate.

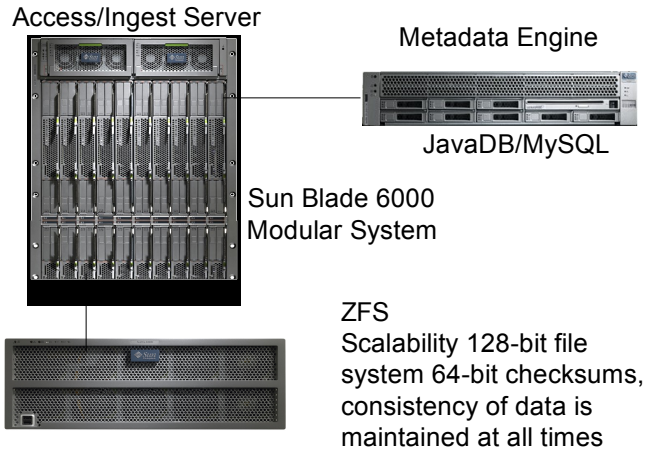

Operating Systems

Although several operating systems currently exist, the logical choice for an archive storage system is an open source operating system, of which there are two primary choices today: Linux and Solaris. There are many varieties of Linux available and supported by nearly all system manufacturers.

The Solaris Operating System is freely downloadable from Sun Microsystems and provides a number of technical advantages from file system support to security and supportability. The Solaris ZFS offers a dramatic advance in data management with an innovative approach to data integrity, tremendous performance improvements, and integration of file system and volume management capabilities. Solaris Dynamic Tracing (DTrace) allows you to analyze, debug, and optimize your systems and applications. Many variants of the Linux operating system and Solaris are available with support on a fee base.

Storage Solutions

There are fundamental differences between storage devices and a tiered storage solution for long-term archiving. In a long-term archive, the data needs to be stored on the least expensive, most reliable media form possible while still supporting access requirements. This is only achieved through tiered storage. You can not simply add more disk as your data increases. For efficiency you must assess the data access patterns to ensure data is being stored on the most appropriate media.

A disk based storage architecture alone remains the most expensive means of storing data. Magnetic tape should always be considered if the amount of data involved will exceed the capacity of a tier of disk. This solution architecture for a long-term storage archive includes storage software, file systems, and devices for rapid access to large amounts of data as well as slower less frequent and dark archives.

This tiered storage architecture also simplifies technology upgrades and the migration of data. The multiple distributed copies combined with the automated data integrity checks provides for a reliable system.

The Sun StorageTek Storage Archive Manager (SAM) software provides the core functionality of the recommended archive storage architecture. SAM provides policy based data classification and placement across a multitude of tiered storage devices from high speed disk, low cost disk, or tape. SAM also simplifies data management by providing centralized meta-data. SAM is a self-protecting file system with continuous file integrity checks.

The Hierarchical Storage System, or HSM, is a key software element of the archive. The HSM provides one of the key components that contributes to reliability through data integrity checks and automated file migration. The HSM provides the ability to automate making multiples copies of files, auditing files for errors based on checksum, rejecting bad copies of files and making new copies based on the results of those audits. The HSM also provides the ability to read in an older file format and write-out a new file format thus migrating the format and application information required to ensure archival integrity of the stored content. The automation of these functions provides for improved performance and reduced operating costs.

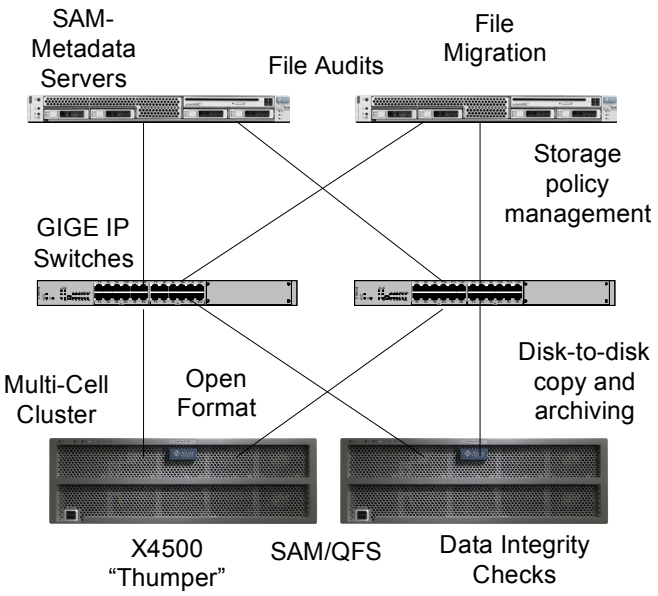

The Sun StorageTek 5800 provides the core hardware storage platform for the reliable archive storage architecture. The ST5800 object storage system provides data protection designed specifically for long-term preservation. The ST5800 provides simplified management of a highly scalable and flexible platform.

The Sun StorageTek 5800 is an object storage system optimized to store large amounts of digital content with some very unique capabilities including extensive metadata facilities that describe the object being stored, and an architecture that allows it to process locally the format of the object being retrieved or stored.

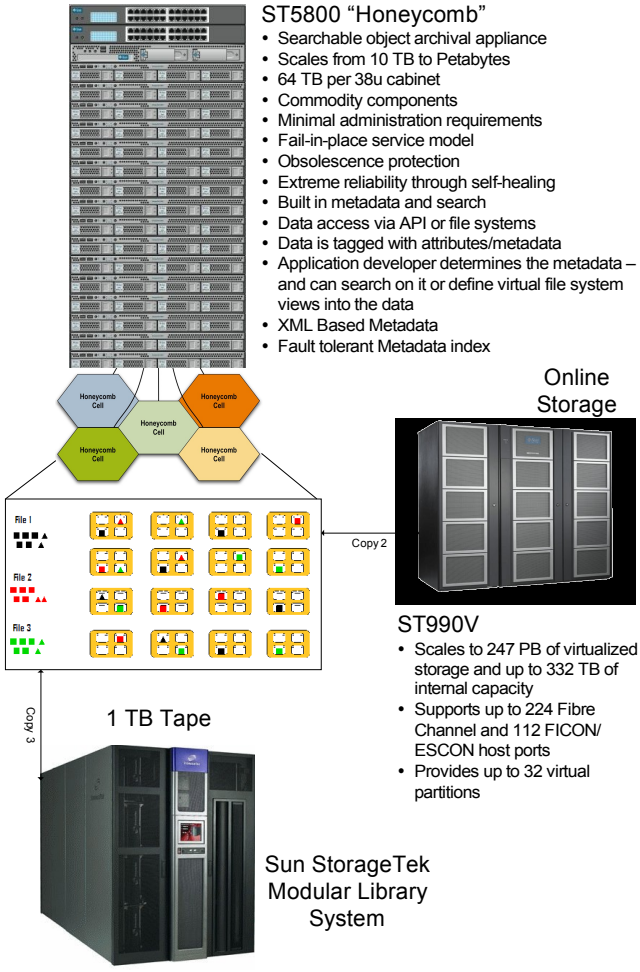

The ST5800 uses a multi-cell based symmetric, clustered architecture. Within a cell, all storage, control, data, and metadata path operations are distributed across the cluster to provide both reliability and performance scaling. Each storage node is independent of all other nodes, and there's complete symmetry in both the hardware and software on each storage node. The Sun StorageTek 5800 provides a comprehensive defensive strategy to corruption and data loss due to bit rot by integrating bit rot protection and real-time checksums. Reed-Solomon RAID 6 protection, advanced data placement algorithms, and self-healing.

Summary

A tiered storage architecture provides the most cost effective solution for object repositories and long-term archives while supporting scalability. The extent at which those storage tiers are deployed is dependant on the access patterns and archival policies. Although this architecture is not intended to cover all business requirements, it can be applied in a modular approach to address specific business requirements where one or more tiers may not be feasible due to business or technical requirements.

References

[1] Reference Model for an Open Archival Information System (OAIS) p.38, CCSDS 650.0-B-1 BLUE BOOK, Washington, DC 20546, National Aeronautics and Space Administration, January 2002

Author Biography

Keith Rajecki is a Solutions Architect with Global Government, Education, and Healthcare at Sun Microsystems. Keith possesses over 13 years of innovative experience in higher education where he has led datacenter consolidations and the standardization of university technology infrastructure. In his role as Solutions Architect, Keith develops strategic technology solutions to meet industry demands.