

Scalable Open Storage for Active Digital Archives

Sun StorageTek™ 5800 System with DSpace



As the trend toward digital content continues to accelerate, many organizations are struggling to find efficient and effective methods to store, manage, and preserve their digital content. Today's digital libraries and research archives are often challenged by the complexity of storing and indexing research data in a way that enables users to easily find and quickly retrieve what they need. Many organizations are also becoming aware of the high cost of achieving persistence for data that must be preserved for decades into the future. Sun has partnered with DSpace to provide an open storage solution that can help organizations store, manipulate, and access their digital resources in a flexible and cost-effective manner. The solution combines the DSpace open source software platform with the Sun StorageTek™ 5800 system to provide a petabyte-scale object store for managing and preserving active data archives in an open storage format.

Highlights

- Reduced cost and complexity of managing and administering active digital repositories
- Massive scalability with the capacity to manage millions of digital objects in DSpace and seamlessly scale to hundreds of terabytes with the Sun StorageTek 5800 system
- Fast, easy deployment with out-of-the-box functionality and a proven interface between DSpace and the StorageTek 5800 system
- High data integrity through automatic block and file level checksums and distribution of file fragments across nodes to help ensure no single point of failure
- Extreme data protection with redundant design that allows multiple nodes and drives to fail without any resultant data loss

Managing the digital data boom

Libraries and research institutions that were once the domain of stacks of books and journals are now handling a vast and rapidly growing array of digital objects ranging from digital video and audio archives to ebooks and research data sets.

With both the quantity and size of digital objects increasing rapidly, today's archives can grow at unpredictable rates. Some of the primary challenges for owners of these digital repositories include:

- Determining the appropriate storage capacity to purchase when the number and size of content submissions to be ingested is difficult to predict
- Finding efficient ways to ingest and categorize data so that researchers can easily submit content
- Enabling users to easily search and find the data they want after it has been stored in the digital archive
- Managing the data from ingest through archiving in a manner consistent with the policies of the organization
- Managing an end-to-end solution without significant investment in IT staff resources


- Preserving data persistency and data integrity in a cost-effective manner
- Maintaining data currency through data transformations that help users avoid retrieving data in a format that can no longer be used

For many types of digital data, it is necessary to preserve both the information content and its authenticity over a period of several decades or even multiple centuries. The data must survive transitions in staff or user communities as well as replacement of hardware or software systems that contain the data. Preserving digital content thus requires active management to avoid potential pitfalls such as bit rot or storage formats becoming obsolete.

Owners of institutional repositories are looking for more efficient ways to manage their digital archives and to reduce the risk of damaged or lost data as the repositories age. With the digital data boom continuing to accelerate, it is becoming clear that traditional methods of managing and storing data are not sufficient for today's needs.

Organizations are looking for large-scale repository solutions that can offer:

- Efficient methods to preserve data integrity and prevent bit rot
- Reduced complexity in executing backups and administering storage systems
- High performance search and retrieval with multiterabyte archives
- An open data format that enables data to survive technology transitions
- Seamless scalability for easy and cost-effective capacity expansion
- Simplified discovery of data and the ability to search or sort based on wide ranging criteria
- Simplified capture of metadata when data is first stored in the repository
- Ability to capture and preserve the relationships between content items



DSpace users can transparently access data stored in the Sun StorageTek 5800 system through the integrated interface.

Sun and DSpace open storage solution for active archiving

Sun and DSpace have teamed up to provide an open storage solution that can provide an end-to-end solution for managing and preserving large scale repositories. The solution provides a petabyte-scale object store with an open storage format and built-in technology to help improve data persistence and reduce the cost of preserving data over long periods of time.

The Sun StorageTek 5800 system is a third-generation object storage system that has been designed to address today's large-scale repository needs. It can be used in combination with the DSpace software platform to provide an intelligent data repository for digital libraries and other applications that generate and preserve large volumes of data and associated metadata.

DSpace is an open source digital repository system that allows researchers to capture, store, index, preserve and redistribute digital data in virtually any format. More than 300

institutions worldwide use DSpace as their digital repository. DSpace provides organizations with an easy to use end-to-end solution for managing and providing permanent access to their digital works. DSpace was originally developed as a joint effort between MIT Libraries and Hewlett-Packard (HP). It is freely available to all commercial and non-commercial organizations under the BSD open source license.

With the help of DSpace community developers, DSpace was recently integrated with the Sun StorageTek 5800 system, creating a robust open storage solution that further simplifies archival storage and helps reduce operational costs for managing active digital archives. The integration of the StorageTek 5800 system takes advantage of a simple, lightweight storage API that will become part of the standard release after DSpace version 1.5.

Earlier versions of DSpace can be retrofitted to utilize this integration. The API enables seamless integration between DSpace and

the Sun StorageTek 5800 system.

Any user or application that utilizes DSpace software can now easily access the Sun StorageTek 5800 system through DSpace as shown in the figure.

Adding intelligent storage to Deep Blue at University of Michigan

Deep Blue is the University of Michigan's institutional repository for work that is produced or sponsored by University of Michigan faculty, staff, students, researchers or affiliates. The project began in 2005 and currently has more than 41,000 items stored in approximately 60 collections. The repository provides a permanent, safe, and accessible service that enables broad public access to the University's best scholarly and artistic work.

Deep Blue is based on DSpace and researchers at University of Michigan have been working with Sun Microsystems to augment their DSpace implementation by deploying the Sun StorageTek 5800 system as a storage module within Deep Blue. "The

“The Sun StorageTek™ 5800 system is compatible with our mission for Deep Blue. Because the system was designed specifically for active archival applications, it fits well with the technical architecture of DSpace, and we believe it will greatly simplify the task of managing data in the Deep Blue archive. As a result, we expect to spend less time and effort managing storage without compromising reliability or performance.”

Cory Snavelly

Manager of Library IT Core Services,
University of Michigan

Sun StorageTek 5800 system is compatible with our mission for Deep Blue,” said Cory Snavelly, manager of Library IT Core Services, and whose group is responsible for the operation and maintenance of servers and storage for Deep Blue. “Because the system was designed specifically for active archival applications, it fits well with the technical architecture of DSpace, and we believe it will greatly simplify the task of managing data in the Deep Blue archive,” continued Snavelly. “As a result, we expect to spend less time and effort managing storage without compromising reliability or performance.”

Integration of the StorageTek 5800 system with DSpace version 1.5 is in final testing by Sun and other DSpace contributors. Deep Blue will begin its on site integration when it upgrades to DSpace version 1.5 this summer. Researchers using Deep Blue will see no difference when they begin to use the StorageTek 5800 since the interface is hidden within DSpace.

Intelligent storage

The StorageTek 5800 system provides intelligent storage functionality at a very low level, reducing the manual effort required to maintain the data archive by moving the process of managing space and scalability down into the storage layer. There are no RAID sets to manage, no Storage Area Networks (SANs) to architect, and no file systems to mount. Administration thus becomes a very simple task, reducing the cost and complexity of maintaining the archive.

The StorageTek 5800 system is also self-healing, meaning that it detects failures and reassembles the data away from affected hardware. When failed components are later replaced, Sun software that runs across all nodes automatically detects new replacement hardware and spreads the data across the replacement hardware. All storage

nodes internal to the storage system have equal capacity utilization, including new or replaced hardware.

The StorageTek 5800 system also does automatic data integrity checking and correction, including block and file level checksums. Checksums are stored separate from the object and are validated upon object retrieval. The system can sustain multiple, simultaneous drive failures or a failure of up to two complete nodes or up to a total of eight drives without losing data access or sustaining data loss.

DSpace is also committed to going beyond reliable file preservation to offer functional preservation or currency of archive data. Future versions of DSpace may support data transformation, making it easy for institutions to migrate files to currently supported formats as their needs evolve over time.

Modular growth without disruption

Both the Sun StorageTek 5800 system and the DSpace platform are architected for scalability, allowing institutions to start small and expand as needed.

The scalability of the StorageTek 5800 system results from its modular architecture. It incorporates clustered servers for both processing and storage, allowing the system to be easily scaled. A half-cell is composed of 8 nodes, and can be easily upgraded to a full 16-node cell. The system can then be expanded to a multicell configuration as data needs increase. With each new cell, more compute power is also added to the system, making the system scalable to hundreds of terabytes.

Similarly, DSpace is based on a modular architecture that enables collections and communities to be added as needed and spread across multiple storage devices. This

architecture gives DSpace the capacity to support multidisciplinary repositories that extend across institutional boundaries.

Investment protection through open source

One of the most challenging aspects of long term archival is that data must survive multiple generations of technology and complete turnover in the staff who manage the data. Open source solutions such as the DSpace platform and the technology of the Sun StorageTek 5800 system provide greater assurance that the digital archives can be preserved regardless of technology upgrades or staffing issues.

All DSpace code is publicly and freely available on sourceforge.net. With more than 300 institutions currently using DSpace software and more than 100 developers around the world contributing code, DSpace has wide acceptance and active support. A 2007 census of Institutional Repositories in the United States conducted by the Council on Library and Information Resources (CLIR) also found that DSpace was the preferred Institutional Repository system.

The Sun StorageTek 5800 Open Edition software is freely available as an open source project. Since it is based on OpenSolaris™ and Sun's commitment to open interfaces, the StorageTek 5800 system helps protect against proprietary lock-in. Data is also kept in an open storage format so that future devices that may be based on new advances in technology will be able to read it. Customers can therefore be virtually assured of being able to retrieve their data well into the future. More details can be found at opensolaris.org/os/project/honeycomb/.

Flexibility to adapt to your needs

DSpace is designed to work out of the box and yet it also provides the flexibility to be easily customized to meet an institution's unique needs. DSpace Manakin provides a modular user interface layer, enabling institutions to design a unique look-and-feel that can be different for each community, collection and item across the repository. Manakin also allows the user interface to extend outside of DSpace into an existing Web presence.

DSpace supports multiple types of storage devices through a lightweight storage API. The storage layer currently provides support for local file systems, Storage Resource Broker (SRB), Amazon S3, or the Sun StorageTek 5800 system. New storage devices or approaches can be quickly integrated using the existing storage APIs.

Let Sun and DSpace help you reduce cost and complexity

The Sun StorageTek 5800 has been designed from its very inception to meet the needs of institutional repositories such as those created using DSpace. Institutions that require high performance archives and want to reduce the cost of preserving large amounts of research data can benefit greatly by extending their DSpace implementation with the Sun StorageTek 5800 system. The combined solution is designed to provide high data integrity, robustness, reliability, and more effective management of metadata to help reduce the cost and complexity of managing active digital archives.

Learn More

To learn more about DSpace and Sun solutions for digital content, visit dspace.org and sun.com/storagetek.

Get started now with free downloads

You can get started immediately with free download software from Sun and DSpace. The DSpace code is available for download at dspace.org and on sourceforge.net. Sun StorageTek 5800 Open Edition software is available for download at opensolaris.org/os/project/honeycomb and includes all of the tools you need to test your application without having a StorageTek 5800 system physically available.