

Hitachi Data Systems Object Storage Solutions



An Evolution in Storage: The Object Store

As unstructured data continues to grow faster than IT budgets, organizations are looking for ways to support growth while reducing complexity and easing the pressure on IT budgets. Hitachi Data Systems provides intelligent, object-based storage solutions that support diverse use cases, from a single cluster, simultaneously, to enable more efficient operations and help IT stay agile as organizations evolve.

Unstructured Data Challenges

The challenge with unstructured data is that it is unstructured, and many of the technologies to manage it were implemented when this data was a small fraction of the total compared to structured data.

As unstructured data began to grow more quickly, the fundamental differences between structured and unstructured data began to impact the IT environment in significant ways. In response, organizations deployed specialized technology to support the vast quantity of data being created. The technology of choice was network attached storage or NAS. Easy deployment and compelling cost led to storage sprawl, which created new challenges in managing, governing protecting and searching content.

Hitachi Data Systems provides solutions to these challenges through a single platform that can be divided into virtual storage systems, each configured for the desired level of service. The great scale and rich features of these solutions help organizations leverage a single storage investment for a variety of workloads.

Intelligent, Multipurpose Storage

Hitachi Data Systems object storage solutions avoid the limitations of traditional file systems by intelligently storing content in far larger quantities and in a much more efficient manner. These solutions provide for the new demands imposed by the explosion of unstructured data and its growing importance to organizations, their partners, their customers, their governments and their shareholders.

The Hitachi Data Systems object storage solutions treat file data, file metadata and custom metadata as a single object that is tracked and stored among a variety of storage tiers. With secure multitenancy and configurable attributes for each logical partition, the object store can be divided into a number of smaller virtual object stores that present configurable attributes to support different service levels. This allows the object store to support a wide range of workloads, such as content preservation, data protection, content distribution and even cloud from a single physical infrastructure. One infrastructure is far easier to manage than disparate silos of technology for each application or set of users.

Intelligent Objects

In most storage systems, the intelligence resides within the storage itself, which limits service to hundreds of millions to about a billion files. This volume, unheard of less than a decade ago, is now becoming more and more common.

To make the significant next jump in scale requires some intelligence to reside in the objects themselves. In such a model individual objects would have the “DNA” to know when to create clones of themselves and how to adjust to changes in environment. For example, in the case of a rush of read requests in a particular geography, objects would be cloned and migrate to the hot spot to service requests locally. Once read activity subsided, objects would know to die off, as there would no longer be a need for such a large population.

As a means of comparison, consider the human organism, which contains tens of trillions of cells. The human organism couldn't operate if it were solely governed by conscious control. Instead, the human organism is controlled by a set of

autonomic functions that operate independently of conscious thought and thus can perform the myriad functions necessary to keep such a complex of cells operating as a single unit. To achieve extreme scales in the tens of trillions of objects, intelligent object stores will likewise need to push down some of the intelligence to the objects themselves, thus creating “intelligent objects” capable of responding to changes in the environment.



Content Preservation

Many organizations want to ensure that digital content is preserved for the long term. Some of the reasons are regulatory, but others are to ensure content is preserved and protected for the future as an asset to the organization. Many times these assets can then provide a competitive advantage for an organization, driving value from the content assets.

Many organizations want to continue using their preferred software provider to interface the content source to the object storage infrastructure and remove their historical “islands of information.” This allows IT to shift its focus to implementing an enterprise-wide strategy with a common repository for long-term management, preservation, protection and search of content and its metadata. It allows IT to take the first steps toward “big data,” reduce the cost and risk associated with managing different “islands,” as well as optimize the return on investment and provide a long-term corporate repository. IT can also improve the cost-effectiveness of the organization's IT strategy and establish a solid platform for future compliance or information governance requirements.

These solutions provide an infrastructure that can be provisioned and configured to serve a wide range of use cases from a single infrastructure that provides key functionality, such as:

- “Write once, read many” (WORM) and content authenticity service for data integrity
- Encryption and access control for privacy and security
- Index and search for e-discovery
- Object tracking and event logging for audit support
- Metadata mining and full content search, which help gather metrics, look for trends and find relationships among data
- Spin-down disk support for reducing the cost for data protection copies and deep archives
- Multiple protocols, which can access advanced features to support multiple applications
- Retention and disposal management services to automatically govern how long content is kept and how it is deleted



Back Up Less...or Backup-less?

The growth in unstructured data stresses traditional, tape-based backup and restore operations. Numerous, disparate systems with large numbers of files and duplicate copies of data increase backup and restore times and impact the performance and availability of production systems. This drives up cost and complexity with the handling of increasing numbers of tapes, the management of offsite storage and the possibility of a compliance or legal action needing information stored in tape-based backups. Hitachi Data Systems distributed object storage solutions attack the problem in 4 ways that reduce the amount of data to be written to tape and streamline recovery processes.

First, the ability is provided to offload data from primary systems to the object store as an active archive. While archives used to be considered only the end of the line for content, the Hitachi Data Systems object storage solutions provide an environment that supports multiple versions of the same content. Multiple versions of less frequently used content can be in the object store

and be accessed directly by users and applications without requiring special tools or custom applications to view and access the archive. By moving less used and static content to an object store, IT vastly reduces the amount of data on expensive, heavily used primary systems. This reduces the amount of time spent backing up and, more importantly, restoring critical systems, and basically eliminates the hassle over less critical content.

Second, data deduplication and compression are used to control data size by eliminating unnecessary copies and shrinking the amount of storage used for a given piece of content. As new objects are written to the object store, the content is compared with similar objects and unnecessary, duplicate data is eliminated or compressed to save space. This capability combines with selective replication (where administrators can decide what data to replicate) to reduce the amount of data at replica sites and conserve precious replication bandwidth. Controlling the overall amount of storage consumed on the object store and any of its replica systems streamlines failover to secondary systems and recovery of primary systems once the failure is repaired.

Third, because of its content preservation capabilities, the object store already ensures data integrity with WORM, encryption and more. By adding services such as data protection levels, advanced replication, version awareness and the ability to browse the environment, the object store ensures objects are well protected and easily recoverable. As the data is onsite and on disk that can be easily browsed, content can be recovered quickly, on demand, at a particular point in time and in a self-service manner. This reduces help desk costs and avoids the hassle of finding the right tape, mounting it (assuming it is onsite), reading the catalog and spinning to the right point of the tape only to learn that another version is needed.

Fourth, the object store provides data retention and disposition services that automatically keep content for the prescribed duration. Barring a retention hold, it automatically deletes expired content so the capacity can be reclaimed and recycled back into available storage. These deletions can be logged and annotated to provide an audit trail of what content was removed,

when, by whom and why. These technologies are key as the traditional methods of keeping every file forever and backing all files up every week are too costly and risky in today's economic, regulatory and legal climates. By putting policies in place and adhering to them with automated tools that log important events, organizations can greatly reduce the risk of failing an audit or facing a fine due to rogue data in long forgotten tape.

By combining the capabilities of an object store with key attributes of data protection, Hitachi Data Systems object storage solutions give IT organizations the ability to deploy a single, intelligent, object-based storage infrastructure. This enables them to back up less data to tape without sacrificing recoverability or scrapping existing investments in backup infrastructure. In addition, Hitachi Data Systems object storage solutions position IT to pursue a backup-less strategy that provides greater protection and faster recovery, and is more reliable as well as easier to use and manage; and, by making use of spin-down disk, this approach rivals the cost of tape-based data protection.



Cloud Enabled

Consider these attributes: the security and integrity of an archive; the protection of RAID-6 erasure coding, advanced replication and failover capabilities; massive scale; support for thousands of tenants and namespaces; built-in chargeback capabilities; a management API; and a REST-based http interface. With these benefits, and more, Hitachi Data Systems object storage solutions compose an ideal platform from which to build the core of a private or public cloud.

Key to the economics of cloud is virtualization and secure sharing of a common set of physical resources. Hitachi Data Systems object storage solutions provide multitenancy that allows IT to securely provision a portion of the infrastructure and turn control of that storage and its capabilities to

HITACHI DATA SYSTEMS OBJECT STORAGE SOLUTIONS

- Address a variety of workloads and provide a variety of service levels to thousands of different tenants with a single storage infrastructure.
- Protect data with dynamic data protection levels, RAID-6 erasure coding, encryption and more.
- Control placement and distribution of content with advanced replication capabilities and intelligent objects.
- Monitor consumption of storage and bandwidth to report internally or charge for use.
- Ensure data integrity with “write once, read many” (WORM), encryption and access control.
- Deploy an edge device or “on-ramp” at distributed sites to connect to the core object store.
- Ensure proper governance of content and track objects in the event of an audit.
- Leverage metadata query tools and full-content search functionality for oversight, management, litigation e-discovery and compliance audits.
- Streamline backup and speed recovery by protecting less active content separately from new and changing data and storing data protection copies on spin-down disk.
- Reduce tape expenditures and limit the amount of data in offsite tapes.
- Reduce data growth with compression, duplicate elimination, and automated retention and disposal services.

the users of that storage. By imposing quotas on those tenants and charging based on their measurable usage, IT can better influence the behavior of users by showing them the cost of their storage practices.

Also important for cloud is the ability to easily adapt new storage models to current

user and application behavior. With an integrated “on-ramp” or “edge” device that connects applications and users at distributed sites to centralized object stores, the power of Hitachi is available to distributed consumers. This enables private organizations to reduce storage and data protection costs at remote or branch offices, and control the distribution of content to different geographies, lines of business and other appropriate audiences.

Cloud service providers can deliver an edge device that integrates directly with their core infrastructure, providing their customers with greater control and security for data in the cloud. In both cases, IT organizations can gain simplicity, focus on the business and speed return on investment.



Summary

Unstructured data has surpassed structured data in total volume and given rise to a new set of challenges for IT. Rather than continually deploying more capacity and suffering the effects of sprawl, the time has come for a change in how content is stored and managed. Hitachi Data Systems object storage solutions are the product of customer and partner input and are designed to address the challenges of fast growing, infrequently used and long-lived content. By integrating many key technologies in a single storage platform, Hitachi Data Systems object storage solutions provide a path to short-term return on investment and significant long-term efficiency improvements. They help IT evolve to meet new challenges and stay agile over the long term and to address future change and growth.

For More Information

To learn more about how Hitachi Data Systems can help you with your unstructured data and to read more about our solutions, please visit www.HDS.com/solutions, contact your local sale representative or solutions consultant, or call Hitachi Data Systems at 888-234-5601.



Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627 USA
www.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2012. All Rights Reserved. SP-012-C VA April 2012