WHITE PAPER

on data

Hitachi Inspire the Next

How to Manage Unstructured Data with Hitachi Content Platform from OnData

By Hitachi Data Systems

April 2012

Hitachi Data Systems

Table of Contents

OnData

As a cloud managed service provider with vast experience in delivering and managing critical IT infrastructure, OnData excels in deploying scalable, high performance solutions.

Rather than promote a one-size-fits-all cloud solution, OnData works with clients to discover:

- How you can best utilise the cloud to ensure reliability and scalability while maximising the ROI
- Which systems you should locate and, most importantly, which you should not
- What are the most appropriate blue chip technology partners for you ensuring that fail-safe best-in-class infrastructure supports the applications you need to run

With all this in mind, OnData has highlighted Hitachi Content Platform as the best way to manage unstructured data. OnData can expertly deploy and manage this solution, enabling you to make the most of its key attributes.

To read about how Hitachi Content Platform can benefit you, please read on.

Executive Summary

Data growth, cost pressures and the complexity of supporting more users and applications pose significant challenges to IT departments. Not the least of these challenges is that many of the technologies to manage unstructured data were implemented when this data was a small fraction of the total compared to structured data.

As unstructured data grows at an explosive pace, the fundamental differences between structured and unstructured data begin to impact the IT environment in significant ways. In response, organizations deploy silos of specialized technology to support the vast quantity of data being created. The technology of choice was once network attached storage or NAS, with easy deployment and compelling cost. However, this lead to storage sprawl, which creates new challenges in managing, governing, protecting and searching content.

Particularly tricky is the management and proper handling of unstructured content across distributed IT environments, such as cloud service providers and organizations with branch or remote office sites. Distributed IT environments are often saddled with sprawling storage silos for a multitude of applications and workloads. They suffer from a buildup of backup infrastructure, and they have inadequate IT resources and safeguards at many locations outside the data center.

Hitachi Data Systems is focused on simplifying and accelerating cloud and other distributed IT environments to readily accommodate shifts in business, economic and regulatory demands.

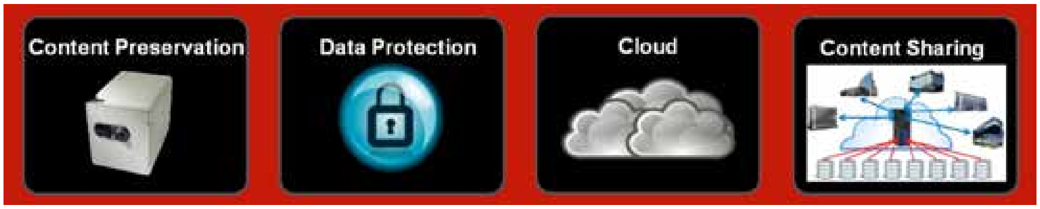

Hitachi Data Systems provides an alternative solution to these challenges through a single object storage platform that can be divided into virtual storage systems, each configured for the desired level of service. The great scale and rich features of these solutions help IT organizations in both private enterprises and cloud service providers leverage a single storage investment to address a variety of workloads. In addition, Hitachi uses a proven and integrated approach to manage distributed IT environments and the flood of storage requirements for unstructured content. With this approach, Hitachi helps cloud service providers and enterprise customers more easily provide storage services to geographically dispersed users.

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 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 and enables the cost to be spread across a large number of users, lines of business and customers.

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" and reduce the cost and risk associated with managing different "islands," as well as to optimize the return on investment (ROI) 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 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 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 (see Figure 1). 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.

Figure 1. Inherent Inefficiencies Associated with Traditional Full and Incremental Backup

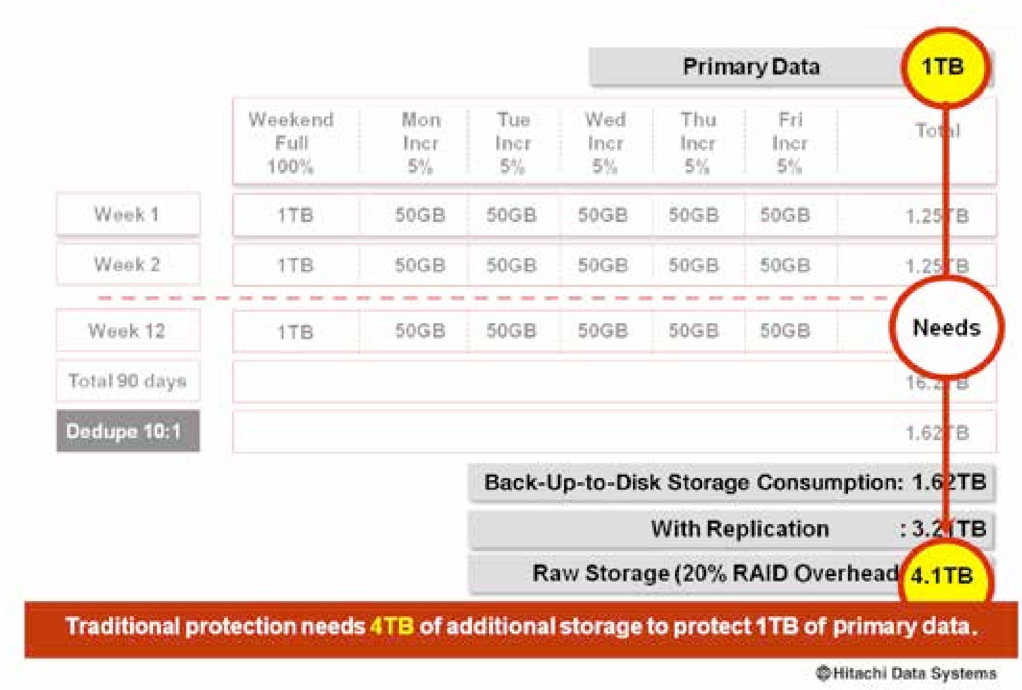

First, the ability is provided to offload data from primary systems to the object store as an active archive. While archives used to be considered only the end of the line for content, the Hitachi Data Systems object storage solutions provide an environment that supports multiple versions of the same 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 (see Figure 2). 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.

Figure 2. Data Deduplication and Compression, Which Are Used to Control Data Size

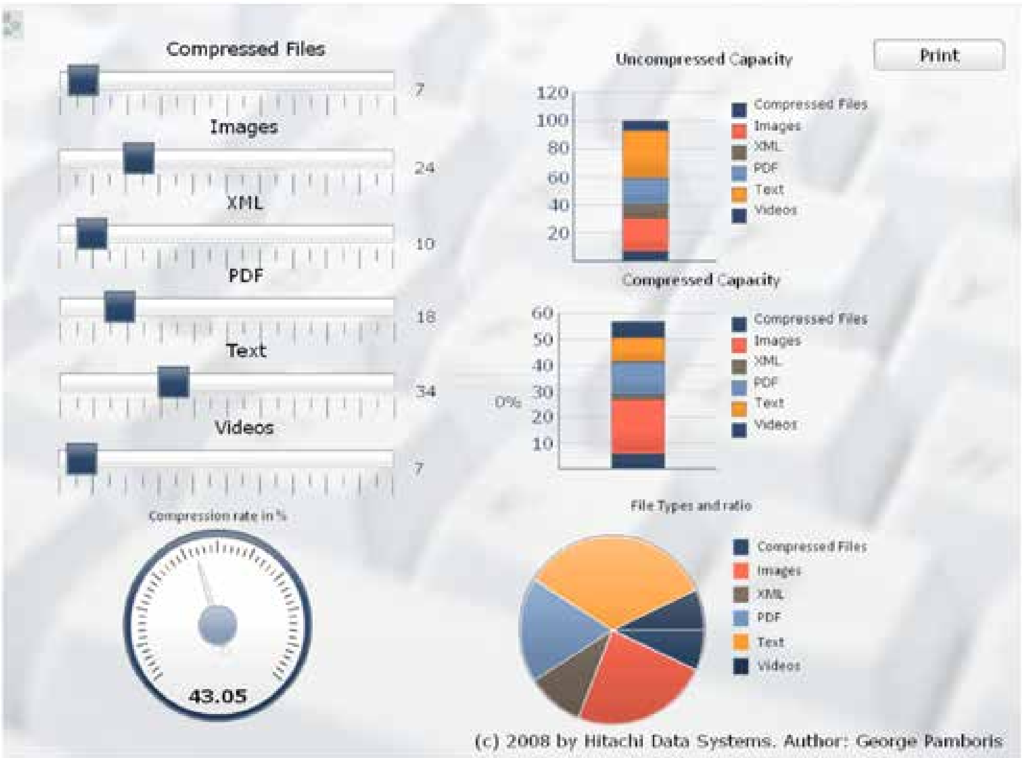

Third, because of its content preservation capabilities, the object store already ensures data integrity with WORM, encryption and more. By adding services such as data protection levels, advanced replication, version awareness and the ability to browse the environment, the object store ensures objects 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 (see Figure 3). 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, it rivals the cost of tape-based data protection.

Figure 3. Cost Reductions of the Hitachi Data Protection Suite and Hitachi Content Platform Solution (in US dollars)

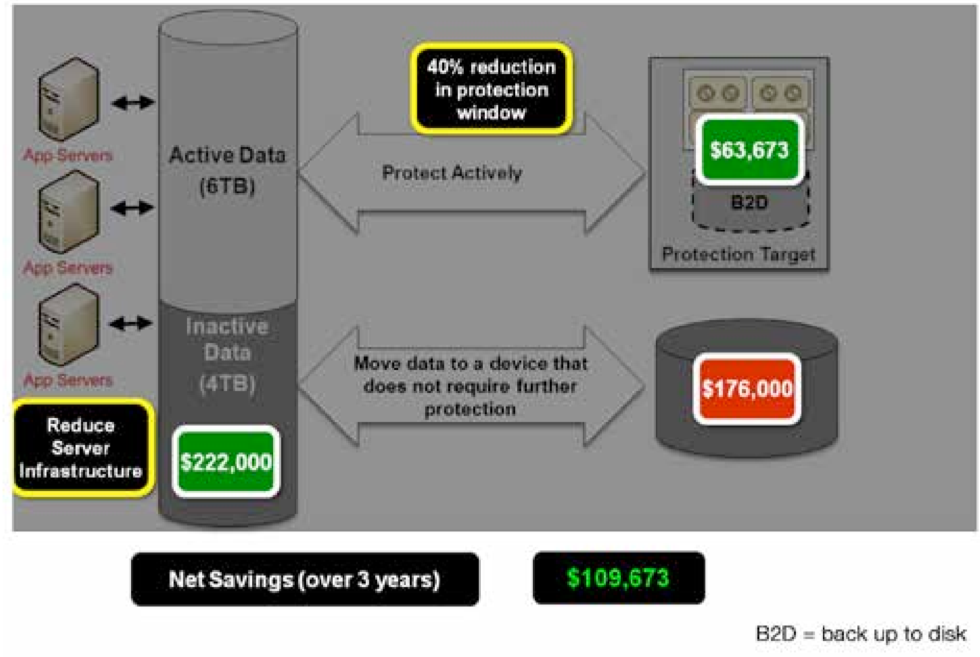

Cloud Enabled

Consider these attributes: the security and integrity of an archive, the protection of RAID erasure coding and advanced replication and failover capabilities, massive petabyte scale, support for thousands of tenants and namespaces, built-in chargeback capabilities, a management API and a REST 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 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 ROI.

Hitachi Data Systems takes the general concept of object storage a leap forward with the Hitachi edge-to-core storage solution. This solution provides distributed consumers of IT, including remote and branch offices and cloud storage patrons, with seamless, bottomless, and backup-free storage (see Figure 4). The 2 products that make up this offering are Hitachi Data Ingestor (Data Ingestor or HDI), a minimal-footprint or virtual appliance placed at the edge, and Hitachi Content Platform (Content Platform or HCP) as the core infrastructure. Hitachi Content Platform is a multipurpose object store with built-in intelligence, virtualization and massive scalability that eliminates the need for a siloed approach to storing unstructured content.

Figure 4. Hitachi Edge-to-core Storage Solution Providing Seamless, Bottomless and Backup-free Storage

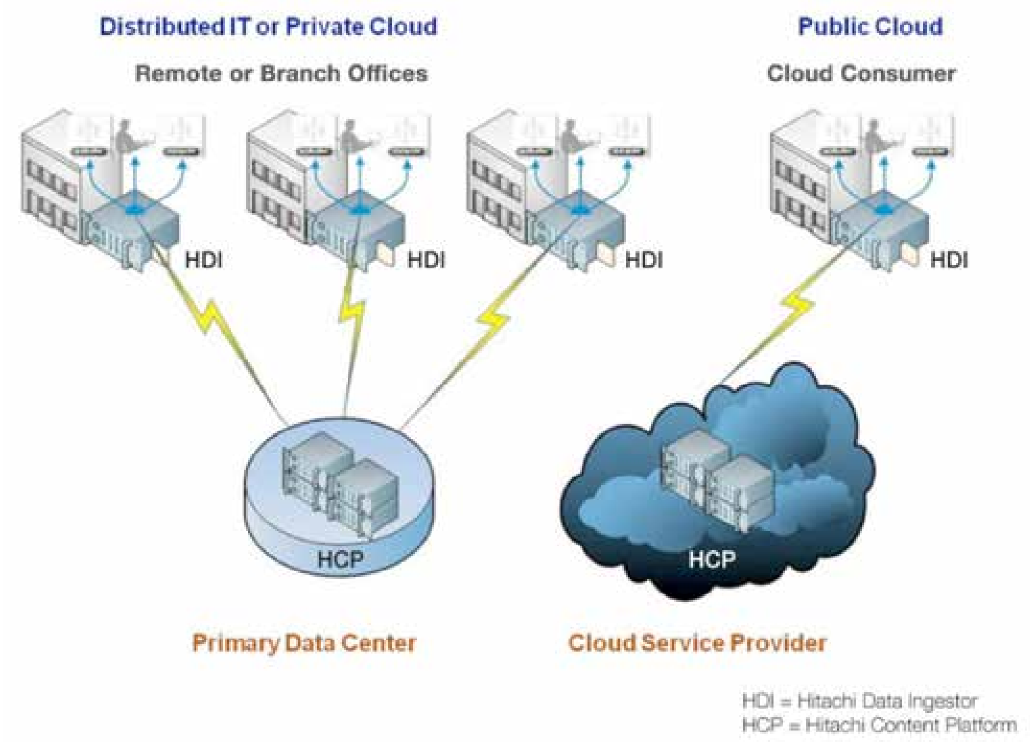

How it works for distributed work environments is simple. The Data Ingestor at each location is mapped to its designated core infrastructure tenant; each tenant can be shared by multiple Data Ingestors. Within the tenant are a number of namespaces. Each file system in that edge appliance is mapped to its designated namespace for proper segregation of data and end-to-end access control. This way, Data Ingestor systems and their clients can write only to their assigned tenants and namespaces.

Each namespace can be, if so desired, shared by multiple file systems on different Data Ingestors. However, only the designated one can write to the namespace while others can only read from it. This content sharing feature enables "edge dispersion," which allows one branch or remote office to access the content written by another. This is a great capability for any distributed or cloud environment that desires it.

Hitachi Data Ingestor and Hitachi Content Platform work together as a single solution to orchestrate the agility and robust characteristics necessary for rapid cloud adoption and cost-efficient distributed IT deployments. Complex data and storage management practices are cost-effectively centralized and automated and tape-based backup is eliminated. In addition, cloud service providers can deliver the benefits of cloud storage without having to build and maintain their own edge-to-core infrastructure.

Hitachi Content Platform Feature Overview

Hitachi Content Platform is an object storage system designed to support large, growing repositories of fixed-content data. An HCP system consists of both hardware and software that manage both data and the metadata that describes the data. It distributes these objects across the storage space but still presents them as files in a standard directory structure.

An HCP repository is partitioned into namespaces. Each namespace consists of a distinct logical grouping of objects with its own directory structure. Namespaces are owned and managed by tenants. HCP provides access to objects through a variety of industry-standard protocols, as well as through various HCP-specific interfaces.

Hitachi Content Platform is the distributed, fixed-content, data storage system from Hitachi Data Systems. It provides a cost-effective, scalable, easy-to-use repository that can accommodate all types of data, from simple text files to medical images to multigigabyte database images.

A fixed-content storage system is one in which the data cannot be modified. HCP uses WORM storage technology and a variety of policies and services to ensure the integrity of the stored data and the efficient use of storage capacity. It also provides easy access to the repository for adding, retrieving and deleting or shredding the stored data.

Object-based Storage

Hitachi Content Platform stores objects in a repository. Each object permanently associates data HCP receives (for example, a file, an image or a database) with information about that data, called metadata.

An object encapsulates:

- **Fixed-content data.** This is an exact digital reproduction of data as it existed before it was stored in HCP. Once it is in the repository, this fixed-content data cannot be modified.

- **System metadata.** These system-managed properties describe the fixed-content data (for example, its size and creation date). System metadata includes policies, such as retention and shred settings, that influence how transactions and services affect the object.
- **Custom metadata.** This is metadata that a user or application provides to further describe an object. Custom metadata is normally specified as XML.

Custom metadata can be used to create self-describing objects. Future users and applications can use this metadata to understand and repurpose object content.

HCP can store multiple versions of an object, thus providing a history of how the data has changed over time. Each version is an object in its own right, with system metadata and, optionally, custom metadata.

HCP also supports appendable objects. An appendable object is one to which data can be added after it has been successfully stored. Appending data to an object neither modifies the original fixed-content data, nor creates a new version of the object. Once the new data is added to the object, that data also cannot be modified. HCP stores directories and symbolic links as objects, too. These objects have system metadata but no fixed-content data or custom metadata.

Open Architecture

Hitachi Content Platform has an open architecture that insulates stored data from technology changes, as well as from changes in HCP itself due to product enhancements. This open architecture ensures that users will have access to the data long after it has been added to the repository.

HCP allows access to stored data by means of several industry-standard protocols. The HTTP, WebDAV, CIFS and NFS protocols support various operations: storing data, creating directories, viewing object data and metadata, viewing directories, modifying certain metadata, deleting objects and much more. You can use these protocols to access the data with a web browser, the HCP client tools, 3rd-party applications, Microsoft Windows Explorer, or native Windows or UNIX tools. HCP also allows special-purpose access to the repository through the SMTP protocol. This protocol is used only for storing email.

For backup, HCP supports the NDMP protocol, local replication to tape and spin-down disk. With NDMP, objects are backed up in OpenPGP format, which uses a tar file to package the files that represent an object. This standard format, which can be both signed and encrypted, allows backup objects to be restored to other storage systems. With local replication to tape, data protection object copies are made to a local tape library. When using spin-down disk, objects are written to disk and the disk is then powered down, saving on power and cooling costs and providing disk-based data protection copies for faster access in the event of failure or disaster.

Namespaces and Tenants

An HCP repository is partitioned into namespaces. A namespace is a logical grouping of objects such that the objects in one namespace are not visible in any other namespace. To the user of a namespace, the namespace is the repository. Namespaces provide a mechanism for separating the data stored for different applications, business units or customers. For example, you could have one namespace for accounts receivable and another for accounts payable.

Namespaces also enable operations to work against selected subsets of objects. For example, you could perform a query that targets the "accounts receivable" and "accounts payable" namespaces but not the "employees" namespace.

Namespaces share the same underlying physical storage. This, together with the multitenancy feature described under "Tenants" in Figure 5, enables HCP to provide support for cloud storage services.

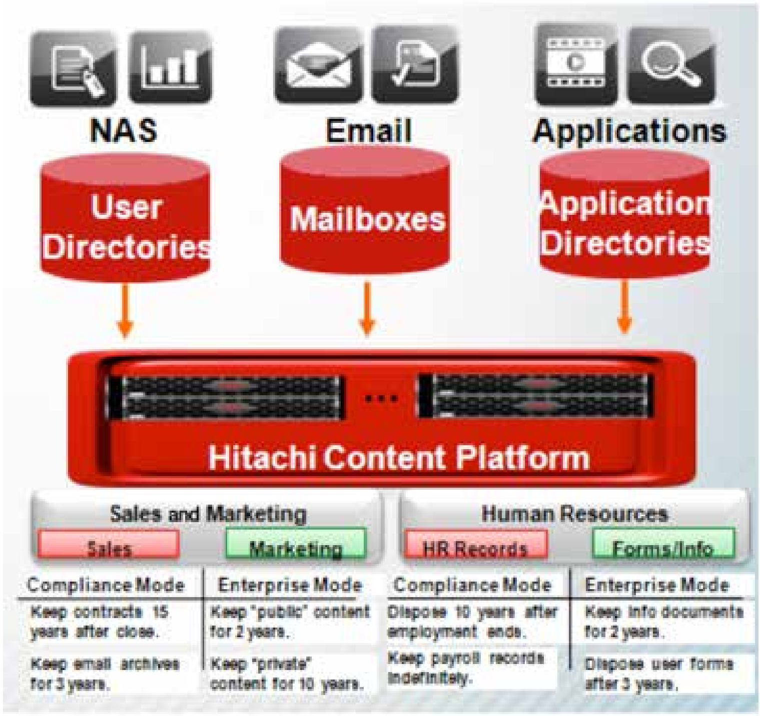

Figure 5. Hitachi Content Platform Tenants and Namespaces

Hitachi Content Platform and Default Namespaces

An HCP system can have up to 10,000 namespaces, including one special namespace called the default namespace. The default namespace is typically used for applications that do not use one of the aforementioned protocols.

Tenants

Namespaces are owned and managed by administrative entities called tenants. A tenant typically corresponds to an organization, such as a company, a division or department within a company, or a customer in cloud deployments.

HCP supports 2 types of tenants:

- *HCP tenants own HCP namespaces.* An HCP system can have multiple HCP tenants, each of which can own multiple namespaces. You can limit the number of namespaces each HCP tenant can own.

- *The default tenant owns the default namespace and only that namespace.* An HCP system can have at most one default tenant.

An HCP system can have 1000 tenants and up to 10,000 namespaces.

An HCP system has both system-level and tenant-level administrators:

- *System-level administrators* are concerned with monitoring the system hardware and software, monitoring overall repository usage, configuring features that apply across the HCP system and managing system-level users.
- *Tenant-level administrators* are concerned with monitoring namespace usage at the tenant and namespace level, configuring individual tenants and namespaces, managing tenant-level and namespace users and controlling access to namespaces.

Object Representation

For HCP namespaces, HCP presents each object as a URL. The root of the object path in the URL is always rest. Users and applications access objects through HTTP.

Here's an example of the URL for an object named wind.jpg in the images directory in a namespace named climate in a tenant named geo in an HCP system named hcp.example.com:

http://climate.geo.hcp.example.com/rest/images/wind.jpg

Other Data Access Methods

In addition to the namespace access protocols described earlier in this paper, HCP allows access to namespace content through:

- The namespace browser
- The data migrator feature
- A set of client command-line tools
- A metadata query engine and API
- The search console

Namespace Browser

The namespace browser for HCP lets you manage content in and view information about HCP namespaces. With the namespace browser, you can:

- List, view and retrieve objects, including previous versions of objects
- Create empty directories
- Store and delete objects
- Display namespace information, including:
  - The namespaces that you can access
  - Retention classes available for a given namespace
  - Permissions for namespace access

- Namespace statistics such as the number of objects in a given namespace or the total capacity of the namespace

The namespace browser is not available for the default namespace. However, you can use a web browser to view the contents of that namespace.

Data Migrator Feature

The data migrator feature for Hitachi Content Platform is a high-performance, multithreaded, client-side utility for viewing, copying and deleting data. With the data migrator feature, you can:

- Copy objects, files and directories between local file systems and HCP namespaces
- Delete objects, files and directories, including performing bulk delete operations
- Purge all versions of an object
- View the content of objects and files, including the content of old versions of objects
- Rename files and directories on the local file system
- View object, file and directory properties
- Create empty directories
- Add, replace or delete custom metadata for objects

The data migratory feature has both a graphical user interface (GUI) and a command-line interface (CLI).

Client Tools

Hitachi Content Platform comes with a set of command-line tools that let you copy or move data between a client and an HCP system. They let you search for the files you want to work with based on criteria you specify. Additionally, they let you create empty directories in a local or remote file system or HCP system.

The client tools support multiple namespace access protocols and multiple client platforms. The command syntax is the same for all supported configurations.

Metadata Query Engine and API

The metadata query engine and API features of Hitachi Content Platform let you search HCP for objects that meet specified criteria. This can search not only for objects currently in the repository but also for information about objects that have been deleted or purged. For namespaces that support versioning, queries can return both current and old versions of objects.

You can query for objects based on any combination of: name, type, size, hash, retention, ingest time, access time, update time, permissions, hold, shred, data protection level, indexed, replicated, namespace, owner and more to serve use cases, such as:

- Content Discovery
  - Return all files owned by "Bob Smith" created after "3/15/2011" with unexpired retention.
  - Locate all content changed within the last 90 days, or that has been modified since last Tuesday.
  - Return objects with a specific retention class defined.
  - Return objects that will expire within a given time range or objects already expired.

- Return objects which do or do not have custom metadata and/or access control lists (ACL)
- Return objects under retention hold, for litigation purposes.

- Intelligent Directory Services
  - Return all files with the extension "mp3 wav wma" in order to synthesize a "Music" folder.
- Bulk Operation Processing
  - Identify a subset of data for which to perform bulk updates.
  - Any user operation could be performed on each result.
- Metrics Gathering
  - What is the size distribution of files in the system?
  - What percentage of my files are Microsoft Word documents?
  - Which owners have generated the most content?
- Testing Metadata Queries Prior to Committing Them to API Code

Search Console

The search console for Hitachi Content Platform is an easy-to-use web application that lets you search for and manage objects based on specified criteria. For example, you can search for objects stored before a certain date or larger than a specified size and then delete them or prevent them from being deleted.

By offering a structured environment for performing searches, the Search Console facilitates e-discovery, namespace analysis, Google-like queries and other activities that look at the contents of namespaces. From the search console, you can open objects, delete objects, prevent objects from being deleted and export search results in standard file formats for use as input to other applications. You can also publish feeds to make search results available to web users.

For searches, you can also employ the Hitachi Data Discovery Suite (HDDS) search facility. This facility interacts with HDDS, which performs searches and returns results to the HCP search console. To use the HDDS search facility, you need to first install and configure HDDS, which is a separate product from HCP.

The search system maintains an index of data objects in each search-enabled namespace. This index is based on object content and metadata. The active search system uses the index for fast retrieval of search results. When data objects are added to or removed from the namespace or when object metadata changes, the search system automatically updates the index to keep it current.

Nodes and Storage

Hitachi Content Platform includes multiple servers (or blades in a blade server), called nodes, which are networked together. Each node can have multiple internal physical drives or connect to external Fibre Channel storage. At least 4 of the nodes in an HCP system must be storage nodes.

Storage Nodes

Storage nodes are the essential part of HCP. They manage the objects that reside in the system storage. To ensure data integrity and continuous availability in case of a hardware or software failure, HCP uses RAID technology and can also store the data and metadata for each object in multiple locations.

Each storage node runs the complete HCP software, with the exception of the HCP search software, if present. These nodes, therefore, act both as a repository manager and as a gateway that enables access to the data in the repository.

All runtime operations are distributed among the storage nodes, thereby ensuring reliability and performance as capacity grows. If a node fails, HCP adapts by redirecting processing to other nodes, so the stored data remains available to users.

Search Nodes

To support the HCP search facility, HCP must include one or more search nodes. The ratio of storage nodes to search nodes depends on the type of data you are storing, as well as the rate at which it is being added to the repository. For help determining how many search nodes your system needs, please contact your authorized HCP service provider.

Linear Scalability

A repository can accumulate a great deal of data over time. To accommodate more data, you can add nodes and storage to Hitachi Content Platform. HCP capacity can grow smoothly from hundreds of gigabytes to terabytes to petabytes of usable capacity. Because HCP uses a distributed processing scheme, the system can scale linearly as the repository grows in size and in the number of clients that have access to it. When you install HCP on new nodes, the system automatically integrates those nodes into the overall workflow, without manual intervention.

Hitachi Content Platform Architecture Overview

HCP runs on a networked redundant array of independent nodes (RAIN) or a SAN-attached array of independent nodes (SAIN). (SAN stands for storage area network.) HCP 300 systems use a RAIN configuration. HCP 500 systems use a SAIN configuration.

HCP hardware consists of:

- Nodes (servers)
- Internal, direct-attached or SAN-attached storage
- Networking components such as cables and switches
- Additional infrastructure items such as racks and power distribution units

HCP uses both back-end and front-end networks. The isolated back-end network connects the HCP nodes to each other through back-end switches. Each storage node is configured with 2 pairs of bonded Ethernet ports that allow external applications to access the system. The recommended setup includes 2 independent Ethernet switches that connect these ports to the front-end network.

SAIN System Architecture[1]

Figure 6 shows the architecture of an HCP system that uses external Fibre Channel storage. This system has 4 storage nodes and 2 modular storage trays.

Figure 6. SAIN Architecture and Interconnect Table

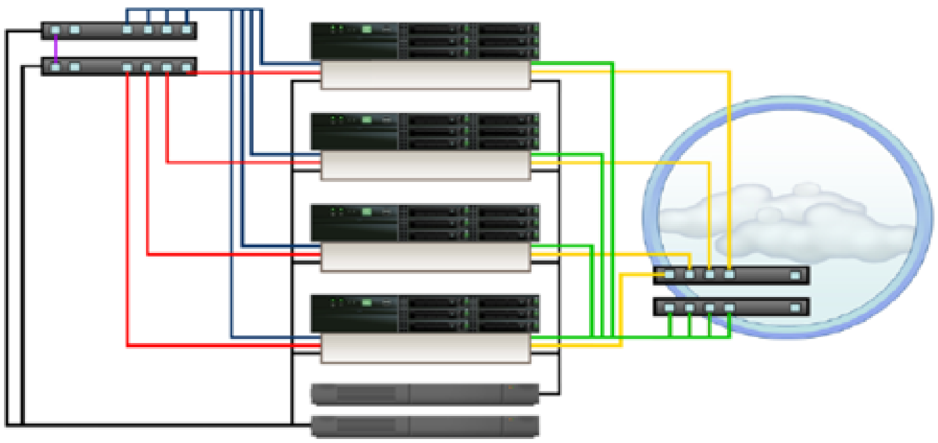

Cable	Connects from	Connects to
Red and blue Ethernet	Back-end network interface cards (NICs) in each node	Back-end switches
Green and yellow Ethernet	Front-end NICs in each node	Front-end switches
Purple Ethernet	Back-end switches	Each other
Black power	Each node	2 PDUs
	Each back-end switch	1 PDU

Repository Management

Repository management entails not only maintaining the integrity and security of stored data but also ensuring the continuous availability of that data, keeping the data in compliance with local regulations, and optimizing the use of storage and network bandwidth. HCP supports these requirements through the implementation of the system itself, configuration choices (both installation and runtime), automated processes, individual namespace and object settings, and object storage and retrieval options.

Data Integrity and Security

Hitachi Content Platform includes many features specifically designed to protect the integrity and

[1] Nodes in HCP SAIN systems are either individual servers or individual blades in Hitachi Compute Blade 320 (CB 320) servers. Some HCP SAIN systems include Fibre Channel switches between the nodes and the external storage.

ensure the security of stored data. In addition to WORM storage, these are:

- **Node login prevention.** HCP does not allow system-console logins on its nodes. This provides a basic level of protection not only for the stored data but also for the system software.
- **Secure Sockets Layer (SSL).** HCP can use SSL to ensure the privacy of HTTP and WebDAV access to namespaces. It always uses SSL to secure the management and search consoles. Additionally, use of the HCP management API requires SSL.
- **Content verification service.** Each data object has a cryptographic hash value that is calculated from the object data. The content verification service ensures the integrity of each object by periodically checking that its data still matches its hash value.
- **Scavenging service.** The scavenging service protects namespaces from the loss of system metadata. If the service encounters an object with invalid metadata, it restores the correct metadata by using a copy from another location.
- **Retention policy.** Each data object has a retention setting that specifies how long the object must remain in the repository before it can be deleted; this duration is called the retention period. HCP ensures that objects are kept until their retention periods expire. The only exception to this occurs in namespaces in enterprise mode. In these namespaces, users with explicit permission to do so can delete objects that are under retention. Such deletions are recorded in the tenant log.
- **Shredding policy.** Objects can be marked for shredding. When such an object is deleted, HCP overwrites its storage location in such a way as to completely remove any trace that the object was there.
- **Data access permission masks.** A data access permission mask determines which operations are allowed. Data access permission masks are set at the system, tenant and namespace levels. The effective permissions for a namespace are the operations that are allowed by the masks at all 3 levels.
- **Data access authentication.** To access the data in an HCP namespace, users and applications must present valid credentials. These credentials are defined by a data access account. The account specifies a username and password. It also specifies which namespaces the user or application can access and which operations the user or application can perform in each of those namespaces.

Data Availability

Hitachi Content Platform has features that help ensure the continuous availability of stored data:

- **Multipathing.** In an HCP system that uses external storage, a single node can connect to more than one port on a storage system, either directly or through multiple Fibre Channel switches. This creates multiple physical paths between the node and any given logical volume that maps to it. With this setup, if one component of a physical path connecting such a node to the storage system fails, the node still has access to the logical volume through another physical path. Multiple means of access to a logical volume from a single node is called multipathing.
- **Zero-copy failover.** In an HCP system that uses external storage, one node can automatically take over management of storage previously managed by another node that has failed. This process is called zero-copy failover. To support zero-copy failover, each logical volume that stores objects must map to 2 different storage nodes. The pair of nodes forms a set such that the volumes that map to one of the nodes also map to the other. This is called cross-mapping.

- **Protection service.** Each namespace has a data protection level (DPL) that specifies how many copies of each object HCP must maintain. When the number of copies of an object goes below the DPL (for example, because of a logical volume failure), the protection service automatically creates a new copy in another location to maintain the DPL.
- **Protection sets.** To protect data availability against concurrent node failures, HCP stores multiple copies of each object on different nodes in an automatically predetermined set of nodes, called a protection set. If a node (or one of its logical volumes) fails, objects stored on its associated volumes (or on the failed volume) are still available through other nodes in the set.
- **Replication service.** You can set up a 2nd HCP system to serve as a replica of the 1st, or primary, system. The replication service automatically copies selected tenants and namespaces on the primary system to the replica. Both configuration information and namespace content are copied. If the primary system fails, the replica can both maintain data availability and serve as a source for disaster recovery. HCP supports many-to-one replication configurations, in which multiple HCP systems send data to a single replica. It also supports chain configurations, in which 1 system replicates to a 2nd system, which, in turn, replicates to a 3rd system.
- **Read from replica.** If an object in a namespace that is being replicated cannot be read from the primary system (for example, because a node is unavailable), HCP can try to read the object from the replica. HCP tries this only if the namespace has the read-from-replica feature enabled and the object has already been replicated. If the namespace on the replica is being replicated to a 3rd HCP system in a chain configuration and HCP cannot read the object from the replica, HCP will try to read the object from that 3rd system.

Regulatory Compliance

Hitachi Content Platform includes features that enable you to comply with local regulations regarding data storage and maintenance:

- **Data privacy.** At HCP installation time, you can choose to encrypt all data and metadata stored in the repository, thereby ensuring data privacy in a compliance context. Encryption prevents unauthorized users and applications from directly viewing namespace content. Lost or stolen storage devices are useless to parties without the correct encryption key. HCP handles data encryption and decryption automatically, so no access or process changes are required.
- **Retention classes.** Some government regulations require that certain types of data be kept for a specific length of time. For example, local law may require that medical records be kept for a specific number of years. A retention class is a named duration that can be used as the retention setting for an object. When an object is assigned to a retention class, the object cannot be deleted until the specified length of time past its creation date. For example, a retention class named HlthReg-107 could have a duration of 21 years. Objects assigned to that class then could not be deleted for 21 years after they're created.
- **Retention mode.** A namespace can be created in either of 2 modes: compliance or enterprise. The difference between them has to do with which operations are allowed on objects that are under retention:
  - **In compliance mode**, objects that are under retention cannot be deleted through any mechanism. Additionally, retention classes (see above) cannot be deleted, and retention class durations cannot be shortened.

  - **In enterprise mode**, users and applications can delete objects under retention if they have explicit permission to do so. This is called privileged delete (see below). Also, in enterprise mode, authorized administrative users can delete retention classes and shorten retention class durations.
- **Privileged delete.** Some localities require that certain data be destroyed in response to changing circumstances. For example, companies may be required to destroy particular information about employees who leave. Privileged delete is an HCP feature that enables authorized users to delete objects even if they are under retention. This feature is available only in namespaces that are in enterprise mode. In compliance mode, objects can never be deleted while they are under retention. With each privileged delete operation, the user is required to specify a reason. HCP logs all these operations, including the specified reasons, thereby creating an audit trail.
- **Retention hold.** To support legal discovery, users and applications can place a hold on selected objects. While an object is on hold, it cannot be deleted through any mechanism, regardless of its retention setting.

Storage Capacity Optimization

Hitachi Content Platform uses these features to reclaim and balance storage capacity:

- **Compression service.** The compression service makes more efficient use of HCP storage by compressing object data, thereby freeing space for storing more objects.
- **Duplicate elimination service.** A repository can contain multiple objects that have identical data but different metadata. When the duplicate elimination service finds such objects, it merges their data to free storage space occupied by all but one of the objects.
- **Disposition service.** The disposition service automatically deletes objects with expired retention periods. To be eligible for disposition, an object must have a retention setting that is either a date in the past or a retention class with auto-deletion enabled and a calculated expiration date in the past.
- **Version pruning.** An HCP namespace can be configured to allow storage of multiple versions of objects. Version pruning is the automatic deletion of previous versions of an object that are older than a specified amount of time.
- **Garbage collection service.** The garbage collection service reclaims storage space both by completing logical delete operations and by deleting objects left behind by incomplete transactions.
- **Capacity balancing service.** The capacity balancing service ensures that the percentage of space used remains roughly equivalent across all the storage nodes in the system. This has the added benefit of balancing the processing load across the nodes.

Network Bandwidth Usage Optimization

Hitachi Content Platform offers these features to help reduce the use of network bandwidth by read and write operations:

- **Compressed data transmission.** Clients that use the HTTP protocol to communicate with

HCP can reduce network bandwidth usage by sending and receiving data in a compressed format. Before sending data to HCP, the client uses the publicly available gzip utility to compress the data. Upon receiving the data, HCP uncompresses it automatically before storing it. When requested to do so, HCP uses gzip to compress data before sending it to the client. The client then uses gzip to uncompress the data.

- **Combined data and custom metadata on reads and writes.** Clients that use the HTTP protocol for namespace access can store or retrieve both the data and custom metadata for an object with a single request. Issuing a single request instead of separate requests for the data and custom metadata reduces the network load. This feature can be used in conjunction with compressed data transmission to further reduce the network load.

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, please contact OnData today:

Call: 08451 489 248
Email: sales@oncoreit.com

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