HITACHI
Inspire the Next

WHITE PAPER

Hitachi Content Platform Architecture Fundamentals

Secure, Simple and Smart Web-Scale Object Storage Platform Delivers Superior Security, Efficiency and Interoperability

By Hitachi Vantara

December 2017

Contents

Executive Summary

Organizations are swiftly assessing and adopting new technologies and information management practices to defend against and ultimately transcend digital disruptors that are emerging across every industry. Pressure is mounting from internal and external influences alike. IT is in the position to define and lead the digital transformation strategy for the organization. Initiatives such as cloud, big data, mobile and social are no longer just buzz, but imminent.

With the IT landscape continuing to evolve, it gets even more difficult to ensure the right data is at the right place, at the right time. The scope of sources from which data is being created or accessed is no longer exclusive to traditional applications and workloads. New technologies, third-party applications and mobile devices mean data is everywhere and constantly changing. The challenge becomes how to retain security, control and visibility of that data, at all times.

Hitachi Content Platform (HCP) is a secure, simple and smart web-scale object storage platform that delivers superior security, efficiency and interoperability. It allows any organization to deliver unique, feature-rich private, hybrid or public cloud storage services at a cost comparable to public cloud. The rich feature set and extensive eco-system surrounding the platform allows organizations to improve efficiencies and optimize costs. They can choose to move data to lower-cost on-premises storage, off-site to a public cloud provider or to a combination of both.

HCP serves as the cloud storage platform for a tightly integrated portfolio of offerings built to service a wide range of information management use cases that span traditional and sustaining applications to emergent and disruptive technologies. The Hitachi Content Platform portfolio provides the ideal ecosystem to support existing content-centric applications and newer cloud use cases and workloads, simultaneously. It includes new functionality and tools that help businesses organize their data, extract intelligence and safely share it with a globally dispersed workforce, all through a single point of visibility and control.

Introduction

Organizations are swiftly assessing and adopting new technologies and information management practices to defend against and ultimately transcend digital disruptors that are emerging across every industry. Initiatives such as cloud, big data, mobile and social are no longer just buzz, but imminent.

With the IT landscape continuing to evolve, ensuring the right data is at the right place at the right time is a serious challenge. New technologies, third-party applications and mobile devices mean data is everywhere and constantly changing. The challenge becomes how to retain security, control and visibility of that data, at all times.

A secure, multipurpose and distributed object-based storage system, Hitachi Content Platform (HCP) is designed to support large-scale private and hybrid cloud repositories of unstructured data. The smart web-scale solution enables IT organizations and cloud service providers to store, protect, preserve, retrieve and distribute unstructured content with a single storage platform. HCP supports multiple levels of service and readily evolves with technology and scale changes. With a vast array of data management, data protection and content preservation technologies, the economical system can significantly reduce resource requirements, and even eliminate its own tape-based backups or backups of edge devices connected to the platform.

HCP obviates the need for a siloed approach to storing unstructured content. With massive scale, multiple storage tiers, Hitachi reliability, nondisruptive hardware and software updates, multitenancy and configurable attributes for each tenant, the platform supports a wide range of applications on a single physical HCP instance. By dividing the physical system into multiple, uniquely configured tenants, administrators create "virtual content platforms" that can be further subdivided into namespaces for further organization of content, policies and access. With support for leading APIs, thousands of tenants, tens of thousands of namespaces, petabytes of capacity in one system, and hybrid cloud configurations based on integration with leading public cloud services, HCP is truly cloud-ready.

This white paper describes how the Hitachi Content Platform portfolio provides the ideal ecosystem to support existing content-centric applications and newer cloud use cases and workloads, simultaneously. It also describes new HCP-based functionality and tools that help businesses organize their data: They can extract intelligence and safely share it with a globally dispersed workforce all through a single point of visibility and control.

Hitachi Content Platform Portfolio Overview

Distinctly unique from the competition, Hitachi Vantara offers an integrated portfolio of object storage portfolio of three products:

Hitachi Content Platform (HCP): An “Enterprise Cloud Storage Platform.” Hitachi Content Platform transforms existing investments into the cloud, including private clouds or hybrid cloud architectures. The rich feature set and extensive ecosystem surrounding the platform allows organizations to improve efficiencies and optimize costs by moving data to a choice of lower cost on-premises storage, off-site to public cloud providers, or a combination of both.

Hitachi Content Platform Anywhere (HCP Anywhere): A “Secure, Enterprise File-Sync-and-Share” solution. Hitachi Content Platform Anywhere enables a more productive workforce with a corporate-delivered file synchronization and sharing tool that allows for secure access and sharing across mobile devices, tablets and browsers.

Hitachi Data Ingestor (HDI): A "Cloud Storage Gateway.” Hitachi Data Ingestor enables remote and branch offices to be up and running in minutes with a low-cost, easy-to-implement, **file-serving** solution for both enterprises and service providers. HDI can also act as a caching solution for higher performing applications such as voice recording.

The Hitachi Content Platform portfolio provides the ideal ecosystem to support existing content-centric applications and newer cloud use cases and workloads, simultaneously. It also provides a central way for organizations to securely incorporate hybrid cloud storage on their terms to react faster to change and to optimize costs. The HCP ecosystem includes a number of integrated Hitachi Vantara products and solutions as well as an expansive set of independent software vendor (ISV) partners and broad protocol support.

HCP serves as the single foundation for storing data from multiple applications and data sources. Through its flexible architecture, it supports traditional applications, such as a variety of archiving applications (for example, file, email, recordings, database, Microsoft SharePoint and medical images). It also serves as the repository for newer Web 2.0, Amazon Simple Storage Service (S3)-enabled cloud, big data, mobile sync and share, remote and branch office, and file and open source application data – all from a single point of management.

With the combination of Hitachi Content Platform, Hitachi Content Platform Anywhere, and Hitachi Data Ingestor, the portfolio of solutions bridges the gaps between traditional IT, cloud and next-generation technologies that extend IT beyond the data center.

Common Use Cases

Hitachi Content Platform is a singular platform that was designed to handle compliance use cases as well as extreme density. Applications written to HCP work seamlessly, regardless of whether HCP is located in the same data center or in the cloud. With HCP, organizations can begin with a very small footprint and grow to have the largest density in the industry. Common HCP use cases are described below and depicted in Figure 1.

Figure 1. Hitachi Content Platform Use Cases

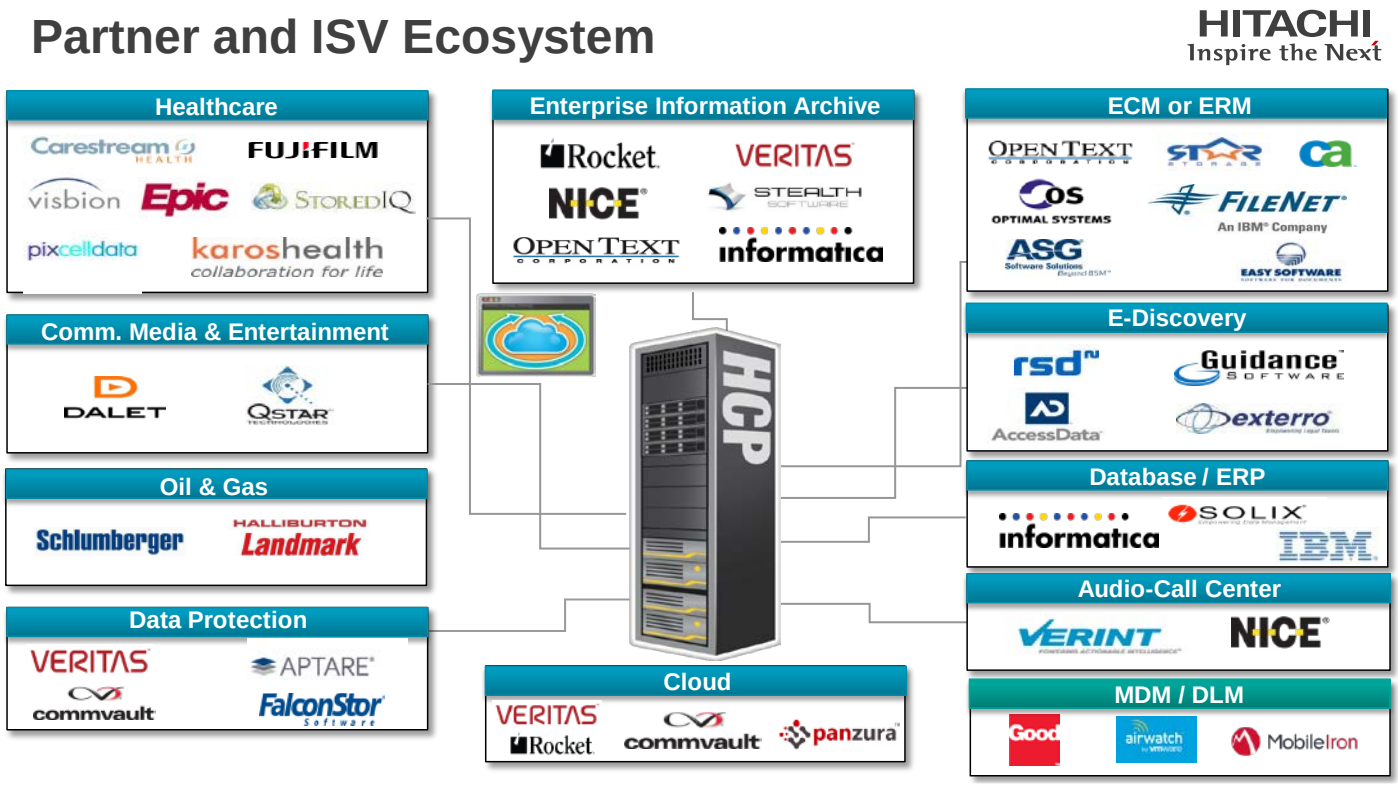

- **Archiving:** Archiving is a common and important use for object storage. Among the primary reasons for its use are to provide economical storage that can scale and advanced data protection, thus removing the need for backup. Over 140 different applications have directly integrated to HCP for archiving a variety of data types, such as email, call data records, document management data, healthcare records, medical images, media and entertainment files, and inactive files from file systems.
- **Regulatory Compliance and Discovery:** HCP adds value to archiving with an advanced set of features for retention management, legal hold and automatic data disposition that help organizations meet compliance regulations, such as SEC 17a4 and Dodd-Frank. Combining built-in custom metadata query and content search from applications, such as Hitachi Content Intelligence, HCP allows administrators to identify data subject to litigation and automatically execute legal holds on the data set to prevent deletion.
- **Backup Reduction and Optimization:** Typically, archiving is employed to remove data from the backup stream that will no longer be changing. Organizations are taking this concept further: Some actually store backup images to HCP where compression, efficient data protection and faster data recall rates provide value beyond backup storage to tape or expensive block deduplication appliances.
- **Storage for Cloud Applications:** Most new applications are being developed leveraging lightweight, web- and cloud-friendly REST APIs and targeting private, public or hybrid cloud storage. Object storage, such as HCP, is the most common type of storage presenting these APIs and optimizing for simple, horizontally scaling storage that can tie a rich set of metadata along with each file.
- **Unstructured Data Management:** HCP leads in capabilities that support policy-based data management capabilities. Policies can be defined to move data between different classes of storage within an HCP cluster or even external to the cluster, targeting popular public cloud services such as Amazon, Microsoft Azure and Google. In so doing, HCP can be a broker of these cloud services, moving data between the services while providing a level of API abstraction for the REST application. HCP provides methods to control geographical placement of data for disaster recovery protection and data distribution. Finally, having the most advanced architecture for custom metadata and built-in query capabilities, HCP is well positioned to act on this information and lead in a new age for data management.
- **Cloud Service Enablement:** Many enterprise IT organizations are revamping their service delivery models to align with growing public cloud service models. Similarly, Tier 2, regional or vertical industry-specific service providers are scrambling to do the same. For these groups, HCP object storage is an excellent storage choice, offering a variety of popular rest APIs and service differentiating features. Further, HCP is part of a tightly integrated portfolio of cloud service applications, including HCP Anywhere for file sync and share and Hitachi Data Ingestor for file cloud gateway capabilities. This portfolio provides a quick start to cloud service providing.
- **Big Data Storage:** By its very nature, big data involves massive quantities of mostly unstructured data. Organizations want to unleash strategic value from this data through analysis. HCP safely, securely and efficiently stores this data. Further, HCP's advanced metadata architecture can bring structure to the unstructured data, allowing analytics applications to query for specific subsets of data hastening analysis and improving results.
- **Remote Office and Branch Office (ROBO) File Services and Content Distribution:** Hitachi Data Ingestor (HDI) combines with HCP to deliver elastic and backup-free file services (NFS or CIFS) beyond the data center. When a file is written to HDI, it is automatically replicated to HCP. Once there, it's not only protected but also visible to other HDIs. This results in efficient content distribution that supports roaming home directories, where a user's permissions follow them to any HDI site. Recent files stay in the HDI file system until free space is needed, creating an "elastic cache." Periodically, HDI converts inactive files to pointers referencing the object on HCP. HDI drastically simplifies deployment, provisioning and management by eliminating the need to constantly manage capacity, utilization, protection, recovery and performance of the system.
- **File-Sync-and-Share Cloud Platform:** Hitachi Content Platform Anywhere provides organizations with a secure file synchronization, sharing and collaboration alternative to consumer-grade or less secure publicly consumed tools. With HCP Anywhere, you can enable a more productive workforce with a file synchronization and sharing tool, delivered from the corporate level, that allows for secure access across mobile devices, tablets and

browsers. End users simply save a file to HCP Anywhere and it synchronizes across their devices. These files and folders can then be shared via hyperlinks. Because HCP Anywhere stores data in HCP, it is protected, compressed, single-instanced, encrypted, replicated and access-controlled. HCP Anywhere also provides enterprise data mobility by enabling mobile access to data in NAS and Microsoft SharePoint storage.

Key HCP Values and Differentiators

Hitachi Content Platform architecture is composed of a comprehensive set of features and capabilities designed to allow organizations to ingest and track information across multiple sources and media types. It simplifies access, search and analysis. It also eases management, improves data protection and lowers costs. These values are enabled through:

- **Unprecedented capacity scale:** Start small (4TB) and scale to unlimited capacity. Deploy entirely as software-defined storage (SDS) based on virtual machines or as an appliance cluster. Scale capacity using Ethernet attached storage (HCP S series nodes) or Fibre Channel arrays.
- **Multiprotocol and heterogeneous access:** Accommodate legacy applications that use NFS, CIFS, SMTP or WebDAV, and those that use modern RESTful APIs including S3, Swift or REST for HCP. Users can write data using any supported protocol and then read data back with another.
- **Construct hybrid storage pools:** Using HCP's adaptive cloud tiering (ACT) functionality, manage a single storage pool using any combination of server disks, Ethernet attached HCP S nodes, SAN disks, NFS or a choice of one or more public cloud services, including Amazon S3, Google Cloud Storage, Microsoft Azure, Verizon Cloud, Hitachi Cloud Service for Content Archiving, or any other S3-enabled cloud service.
- **Multitenancy for application isolation:** With thin provisioning and capacity quotas, divide your storage resources into thousands of independent tenant and namespace areas, each with independent administration and assigned users.
- **Powerful service plans:** Define cradle-to-grave data management plans that govern an object's protection class, access speed and disposition policies.
- **Compliance storage modes:** Satisfy regulatory requirements that require immutable, undeletable "write once, read many" (WORM) storage, guaranteed authenticity or proof of chain of custody.
- **Extensible metadata and search**: Create and modify custom metadata at any time during an object's life cycle. Multiple authors can have separate sections of custom metadata. Use the API or search console to locate objects for application and analytical use, or to automatically apply legal holds.
- **Navigate technology transitions:** With Autonomic Technology Refresh (ATR), HCP boasts a 12-year track record of helping organizations perform online migrations from old to new hardware technology, preserving their application and API investments.
- **Performance:** A unique shared-storage architecture that is equally adept at serving small or large objects.
- **Global access topology:** Read or write data from any site and control where cached copies reside. Share data, but ensure it is hosted within country or continent boundary.
- **Portfolio breadth:** Go beyond the simple archive; choose tightly integrated sync-and-share, file gateway and backup solutions from Hitachi Vantara, or select applications from more than 100 ISV partners.
- **Data protection and security:** Capabilities include RAID-6, erasure coding, redundant copy control, AES256 Encryption, 2048-bit SSH service keys, SSL and certificates.
- **Monitoring:** The user interface and API provide visibility into hundreds of alerts and event logging, as well as chargeback reporting.
- **Data durability and efficiency:** Content verification services, sophisticated self-repair, multisite replication using geodistributed erasure coding, deduplication and compression.

- **Global support:** All HCP systems are eligible for free monitoring through the Hi-Track Remote Monitoring system. Hitachi's global service centers are staffed 24/7.

Architecture Overview

HCP cloud storage software is deployed on hypervisors or dedicated servers called access nodes (HCP G series nodes). As a cluster, these nodes virtualize and federate back-end capacity supplied by HCP storage nodes (HCP S series nodes), block, file or public cloud object sources. Each access node in the HCP system runs a complete software stack made up of the appliance operating system and the HCP core software (see Figure 2). All access nodes run an identical software image to ensure maximum reliability and fully symmetrical operation of the system. An HCP access node can serve as both an object repository and an access point for client applications and can take over the functions of other nodes in the event of node failure.

Figure 2. Example of Access Nodes in HCP

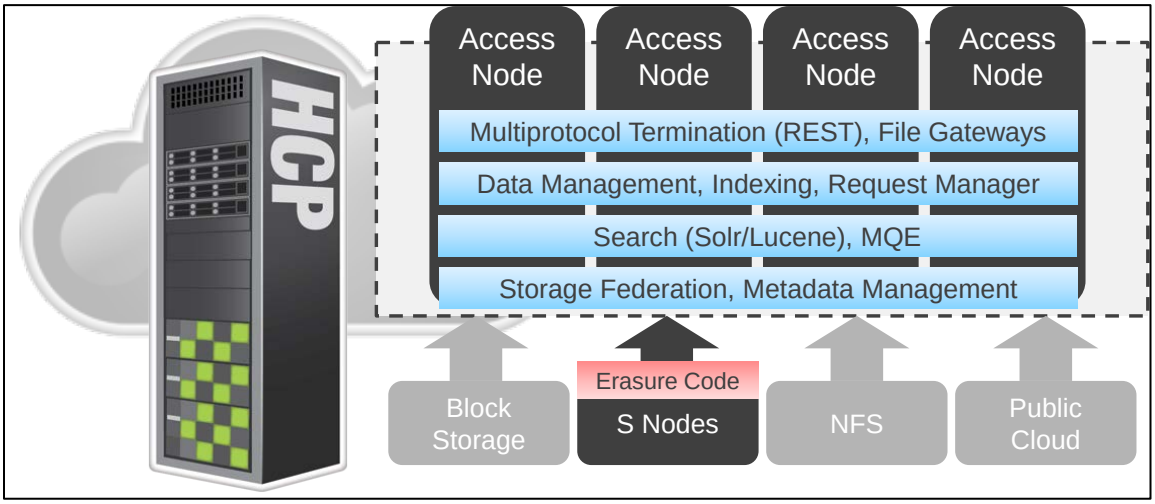

An HCP system is inherently a distributed system, spreading key functions, such as metadata management and storage placement, across all nodes. To process incoming client requests, software components on one node interact with components on other nodes through a private back-end network (see Networking section). All runtime operations are distributed among the access nodes. As such, no single node becomes a bottleneck since each node bears equal responsibilities for processing requests, storing data and sustaining the overall health of the system. They work cooperatively to ensure system reliability and performance.

The HCP distributed processing scheme allows it to scale linearly to accommodate capacity growth or more application clients. When a new node is added to the HCP system, the system automatically integrates that node into the overall workflow without manual intervention.

Access nodes (HCP G series nodes): HCP G series access nodes are in front of an HCP cluster. They perform the service and management functions and can be optionally enlisted to hold object data. All application data passes through access nodes. They always store a portion of object metadata and control object placement within the physical storage pool(s).

Storage nodes (HCP S series nodes): The optional HCP S series storage node appliances provide erasure code (EC) protected storage pools that are managed by access nodes. These Ethernet-attached nodes function as shared storage resources, which enable HCP to scale storage independently of compute resources. There are presently two S series node varieties called S10 and S30 (see Hardware Overview section).

Flexible: Start Small, Scale Big

HCP cloud software offers greater flexibility and choice by deploying as wholly virtual (via hypervisors), wholly appliance, or as a hybrid of both. In all cases, the object repository tolerates node, disk and other physical component failures. HCP architecture offers deployment flexibility that enables the platform to accommodate small workloads while also being able to scale efficiently to manage larger cloud configurations with ease.

Small Deployment: Shared Nothing Configuration. This deployment is composed entirely of access nodes that can each manage a private pool of internal storage in the 28TB range (see Figure 3). To create the HCP cluster, individual nodes are networked through a set of physical 10Gb Ethernet ports. In this deployment model, capacity scaling is achieved by adding nodes. Although ideal for small, sub-petabyte cloud designs, scaling only with access nodes may needlessly increase the compute capability of the cluster, since each increment of storage triggers a significant, potentially underutilized number of CPUs.

Figure 3. HCP Small Deployment

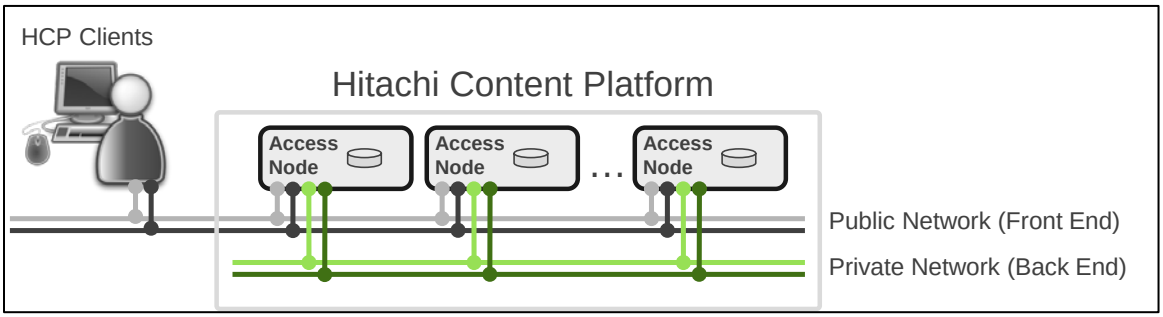

Large Deployment: Shared Storage Configurations. HCP architecture allows its compute and storage elements to scale independently through the use of shared storage (see Figure 4). Networked storage elements behind an HCP cluster can include any combination of on-site S nodes, Fibre Channel storage arrays [(such as Hitachi Virtual Storage Platform (VSP)], NFS, tape, optical or off-site public cloud storage.

Figure 4. HCP Large Deployment

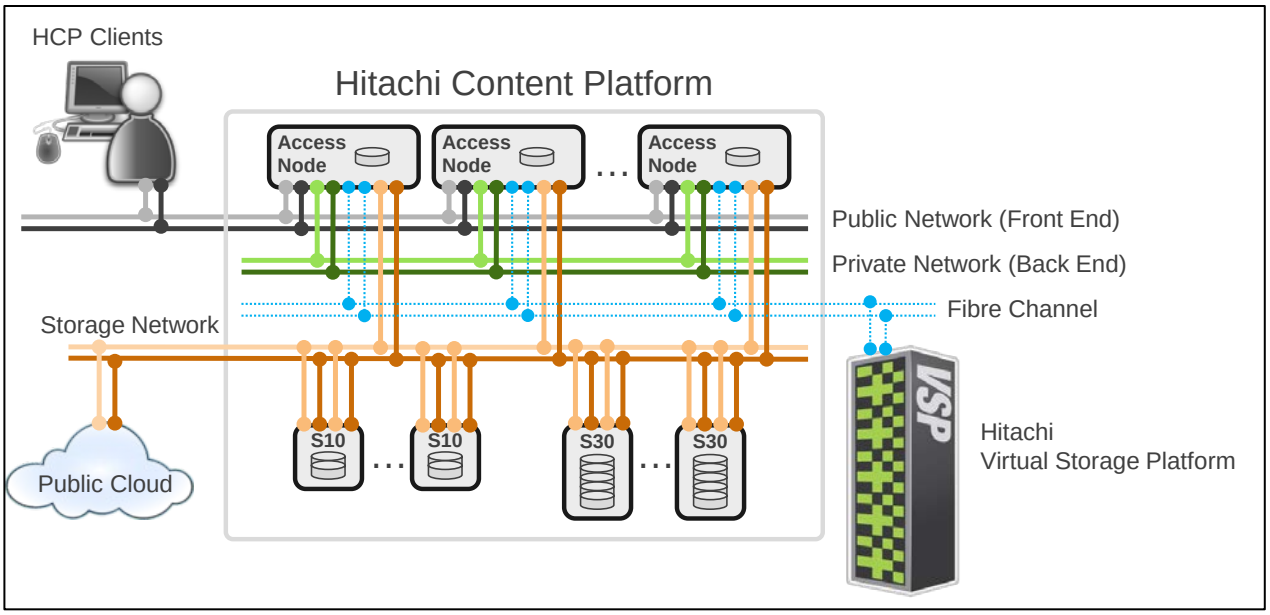

By incorporating support for off-site public cloud storage targets, HCP encourages adoption of hybrid cloud configurations, which can lower the costs of storing older less-active data. By trading a little performance and latency, organizations gain near instant capacity elasticity while retaining a single point of management for both new and old data.

The Power of Shared Storage

Shared storage lets organizations make hardware investments based on application need rather than an artifact of architecture design. For example, as the number of clients grows, there is generally a proportional increase on the HCP workload. HCP G series access nodes may be scaled to linearly improve small object performance and large object throughput, or increase CPU power available to HCP search and data services.

Alternatively, an organization may decide to tackle a new application that needs to store larger media or video files. In this case, HCP is not driving a lot of new I/O as much as it is directing many large files. In this case, additional HCP S nodes might be best to quickly add several petabytes to their virtualized storage pool.

In a pure Ethernet deployment, HCP G series nodes and HCP S series nodes are networked through a combination of physical 10Gb Ethernet ports and VLANS in a loosely coupled architecture. HCP S nodes are particularly well suited for storage scaling. These nodes can be scaled in convenient 125TB usable storage increments called half trays. The flexibility to deploy one or literally thousands of half trays gives organizations the ability to grow as needed, spreading capital investments over time.

Multisite Consistency

HCP installations may or may not span multiple geographical sites (see Figure 5). When considered from a single site perspective, HCP design favors consistency and availability as defined by Brewers CAP theorem[1]. The theory postulates that a distributed system of nodes can satisfy at most two of these three properties:

Figure 5. Node Properties at HCP Sites

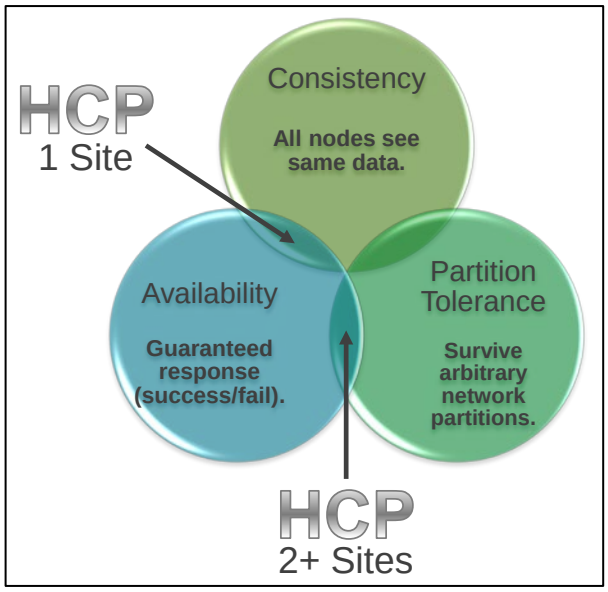

- *Consistency*. All nodes see the same data at the same time.
- *Availability*. Every request receives a response about whether it was successful or failed, guaranteed.
- *Partition tolerance.* The system continues to operate despite arbitrary message loss or failure of part of the system.

Within a **single site**, HCP will never return stale data, which is useful for applications that require strong data consistency. While HCP can handle many forms of partition failure, it does require that a majority of the HCP

[1] CAP Theorem, Eric Brewer, University of California, Berkley

access nodes (total nodes/2+1) be available and communicating with each other in order to take write requests. Reads can be processed with as few as one surviving node.

When an HCP deployment spans **two or more sites**, supporting an active-active global namespace, HCP favors data availability and partition tolerance over strict consistency; this is the favored model for public cloud deployments and is referred to as an eventually consistent model. In response to a whole site outage, HCP may deliver data from a surviving site that was not yet consistent with the failing site. This effect is a result of asynchronous replication, but minimized by HCP's **global access topology** (see Figure 6), which performs hyper-replication of metadata.

Figure 6. Global Access Topology Hyper-Replication With HCP

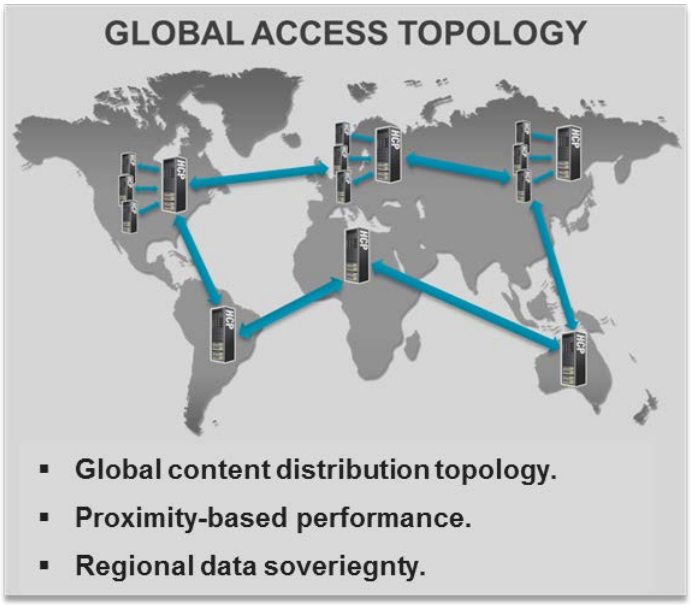

Hyper-replication is possible because each HCP system maintains a separate structure for object data versus object metadata. When an application writes an object, metadata is stored in a separate but parallel branch of HCP's internal file system. This physical separation enables many unique capabilities, including better data consistency between sites because HCP prioritizes metadata replication over replicating the actual object. Intersite consistency is thus less affected by network speed or object size. Participating sites are more quickly aware of new or modified objects. A physically separate structure for metadata is also key to HCP search, tiering and fencing capabilities, which are discussed later in this paper. When all sites are optimal, each HCP can respond to I/O requests with local resources and remain unaffected by the speed or latency of the WAN interconnecting sites.

Manage Exabytes of Capacity

HCP supports configurations exceeding 2.9 exabytes. Harnessing the full potential of HCP's scalable storage capabilities begins with great multitenancy management, delegation and provisioning features (see Figure 7). There is no need to prepurchase or reserve storage for specific applications. Rather, buy a modest amount upfront and grow capacity incrementally as demand increases. Manage in general terms, using quotas, and bill users by what they actually consume instead of what they might consume in the future. Offer them service options that appeal to their data usage patterns, such as versioning, compliancy and automated tiering plans to lower the costs of carrying older data.

System-level administration: These management roles cannot read or write data, but they do control how physical storage resources are virtualized and monitored. They design service plans to govern data placement, how it ages and how it is retired. These managers prioritize system services, create **tenants** and delegate control over capacity using a quota system.

Tenants provide management and control isolation at an organizational level, but are bounded by policies set forth by the system-level administrator. A tenant typically represents an actual organization such as a company or a department within a company that uses a portion of a repository. A tenant can also correspond to an individual person. An HCP can have many HCP tenants, each of which can own and manage many namespaces.

Figure 7. HCP Capabilities Manage Exabytes of Capacity

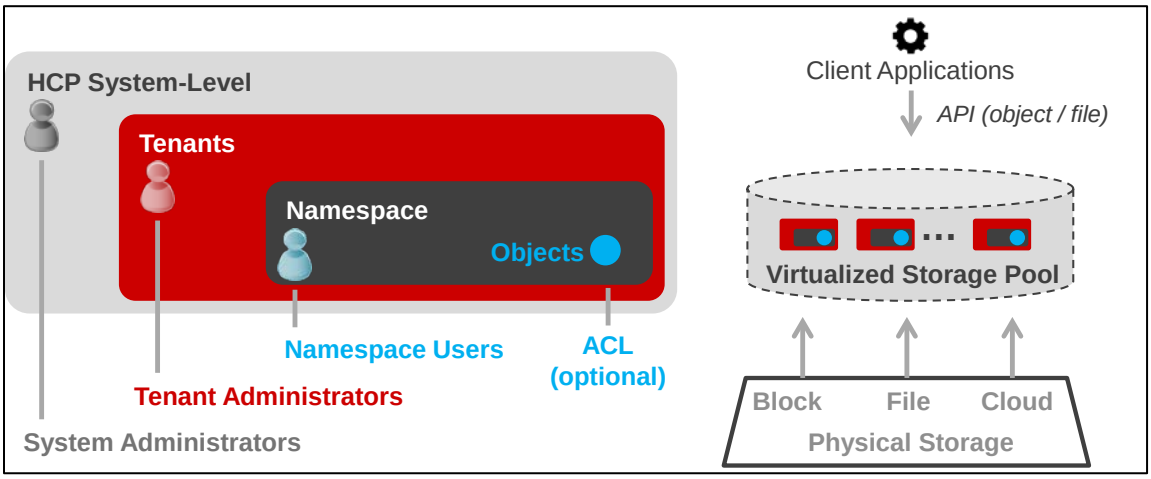

Tenant-level administration: There is a separate administrator for each tenant. They create and manage namespaces for application use at a micro level. They control namespace capacity through quotas, define user membership, access protocols and service policies. They further define which users can read, write, delete or search a namespace. The HCP system-level administrator controls the number of namespaces each HCP tenant can create.

A **namespace** is the smallest unit of HCP multitenancy capacity partitioning. Namespaces are thin provisioned and carved from the common virtualized storage pool. Namespaces provide the mechanism for separating the data stored by different applications, business units or customers. Access to one namespace does not grant a user access to any other namespace. Objects stored in one namespace are not visible in any other namespace. Namespaces provide segregation of data, while tenants provide segregation of management.

Applications access HCP namespaces through HCP REST, S3, Swift, WebDAV, CIFS (SMB 3.1.1), NFS v3 and SMTP protocols. These protocols can support authenticated and/or anonymous types of access. When applications write a file, HCP conceptually puts it in an **object** container along with associated *metadata* that describes the data. Although HCP is designed for WORM access of information, namespaces can be enabled with versioning to permit write and re-write I/O semantic (see software overview).

System-Level Dashboards and Notifications

With the web-based overview dashboard, the system-level administrator can quickly assess HCP cluster status (see Figure 8). The single pane summary displays color-coded health alerts, data services, major events and the total capacity consumed by all tenants and namespaces. Use one-click drill down into any of the 500+ alerts or events and electively choose to enable email notifications, SNMP or system logging (syslog).

Figure 8. HCP Web-Based Overview Dashboard

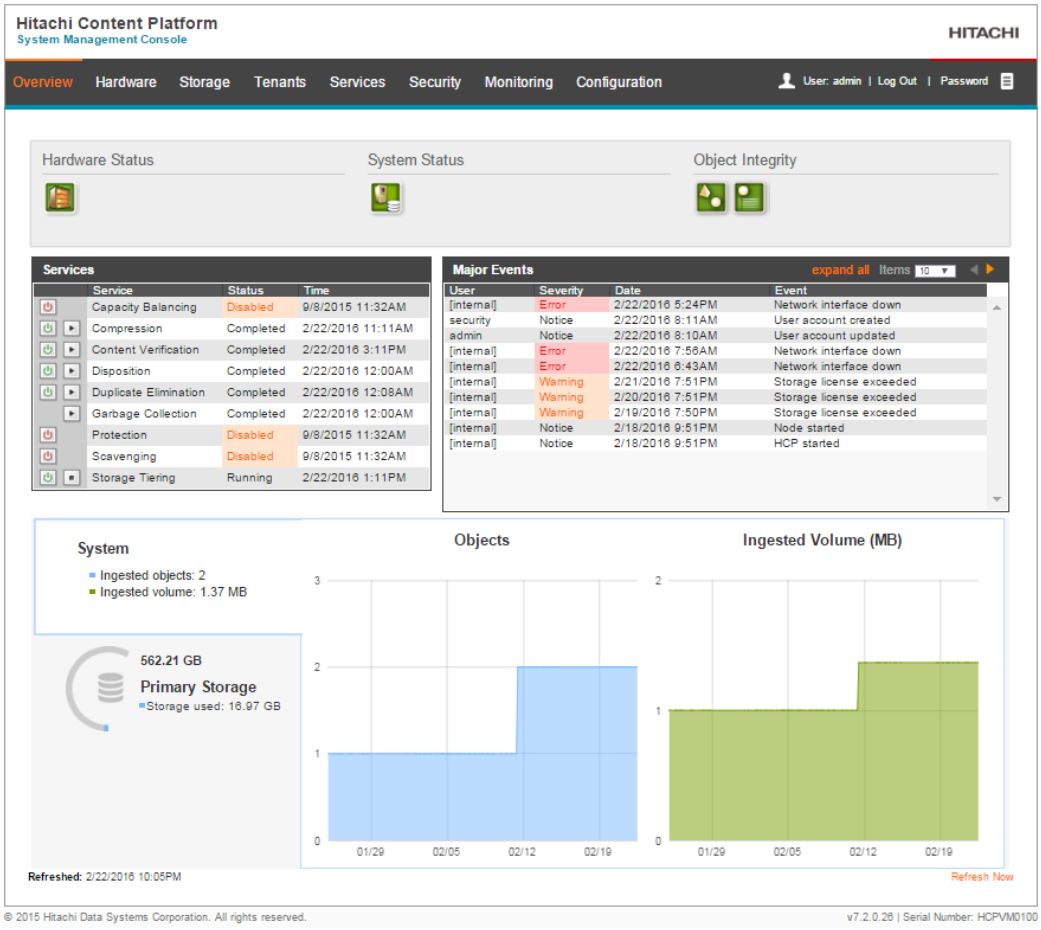

Wizards: HCP provides guided input templates to streamline virtually all major configuration tasks, notably for tenant creation and designing namespace service plans that control data placement and object life cycles.

Scheduling: Use a 7-day, 24-hour calendar to ensure post-process services such as tiering, deduplication, compression and verification run outside of your peak workload periods.

Tenant-Level Dashboards and Notifications

The overview dashboard for a tenant administrator displays a summary of events, and the sum total capacity consumed by all its defined namespaces (see Figure 9). The panel provides one-click drill down into any events, which are also forwarded to an email address.

Figure 9. HCP Overview Dashboard for Tenant Administrator

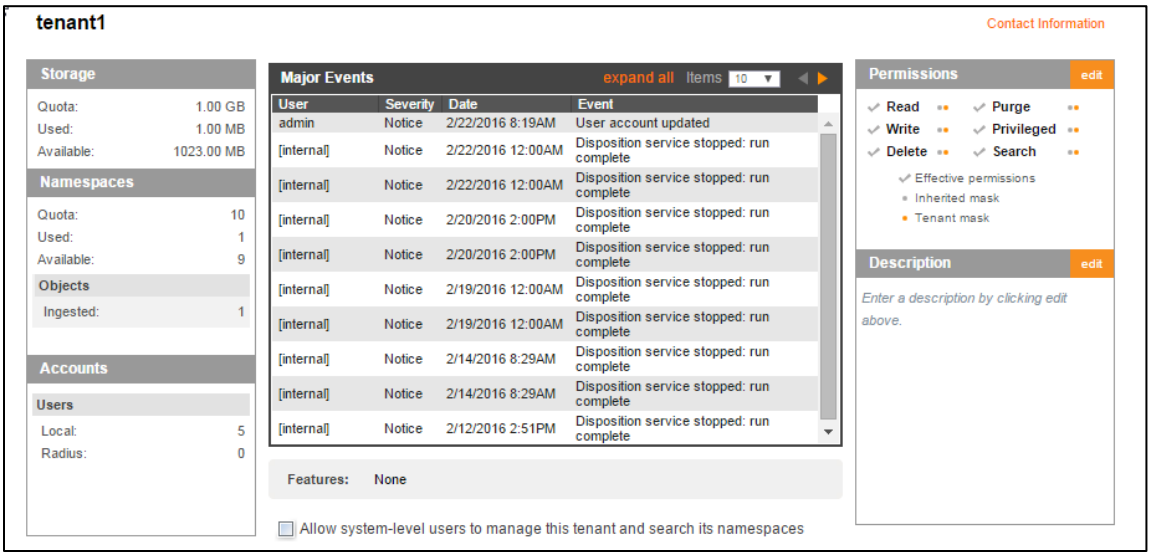

Namespace configuration templates: Each tenant administrator is delegated with authority over an allotted capacity. These templates help them create namespaces, configure permitted protocols, set capacity quotas and policies for retention, disposition, indexing and search. Optionally, configuration can be carried out through REST API or Microsoft PowerShell utilities.

Enterprise mode: The tenant administrator is always permitted to create namespaces with an enterprise retention policy. While normal users cannot delete objects under enterprise retention, a tenant-admin can be empowered to preform audit-logged privileged deletes.

Compliance mode: The tenant administrator can be permitted to create namespaces with a compliance retention policy. Objects under compliance retention cannot be deleted through any user or administrative action until their expiry date. Industry-specific regulations sometimes mandate immutable compliance modes to protect electronic business records. Utilize this mode with prudence since even experimenting can create permanent undeletable content.

Chargeback Reporting

Data consumption reports (see Figure 10) are available to both system and tenant administrators as an onscreen display or as a download. At the system level, the report will include a rollup of all tenants and their namespaces, while individual tenant administrators will receive a report limited to the namespace(s) they own.

Figure 10. HCP Data Consumption Report

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	systemName	tenantNa	namespac	startTime	endTime	objectCou	ingestedV	storageCa	bytesIn	bytesOut	reads	writes	deletes	deleted	valid
2	hcp.hcp-demo.com	splunk	folder	2/11/2016 21:25	2/11/2016 23:59	1	384686	385024	384686	0	0	1	0	FALSE	FALSE
3	hcp.hcp-demo.com	splunk	folder	2/12/2016 0:00	2/12/2016 23:59	1	384686	385024	0	0	0	0	0	FALSE	FALSE

Object Storage Software Architecture

Hitachi Content Platform is an object storage platform. Its architecture by nature means it is more efficient, easier to use, and capable of handling much more data than traditional file storage solutions. HCP automates day-to-day IT operations and can readily evolve to changes in scale, scope, applications, storage, server and cloud technologies over the life of data. In IT environments where data grows quickly or must live for years, decades or even indefinitely, these capabilities are invaluable.

Hitachi Content Platform eliminates the need for a siloed approach to storing unstructured content. HCP software receives unstructured data files via file or REST protocols and stores them as objects. All objects in the repository are distributed across all available storage pools according to policies assigned to the namespace. Externally, HCP presents each object either as a set of files in a standard directory structure or as a uniform resource locator (URL) accessible by users and applications via HTTP or HTTPS. In all cases, the software retains any file directory structure applied by protocols. Once ingested, the software offers a variety of services and tools to protect object integrity, manage object life cycles, search it, and ensure it is always available.

Object Container Structure

An HCP object is composed of fixed-content data (a user's file) and electronic "sticky notes" called metadata. Metadata describes the fixed-content data, including its properties. All the metadata for an object is viewable, but only some of it can be user-modified. The way metadata can be viewed and modified depends on the namespace configuration, the data access protocol and the type of metadata. HCP metadata types include *system metadata* and, optionally, *custom metadata and access control list (ACL)*. The structure of the object is shown in Figure 11.

Fixed-content data is an exact digital copy of a written file which is "fingerprinted" upon ingest using a hashing algorithm: MD5, SHA-1, **SHA-256 (default)**, SHA384, SHA-512 or RIPMD160. These files become immutable after being successfully stored in a virtual storage pool. If the object is under retention, it cannot be deleted before the expiration of its retention period (see compliance modes). If versioning is enabled, multiple versions of a file can be retained.

Figure 11. HCP Object

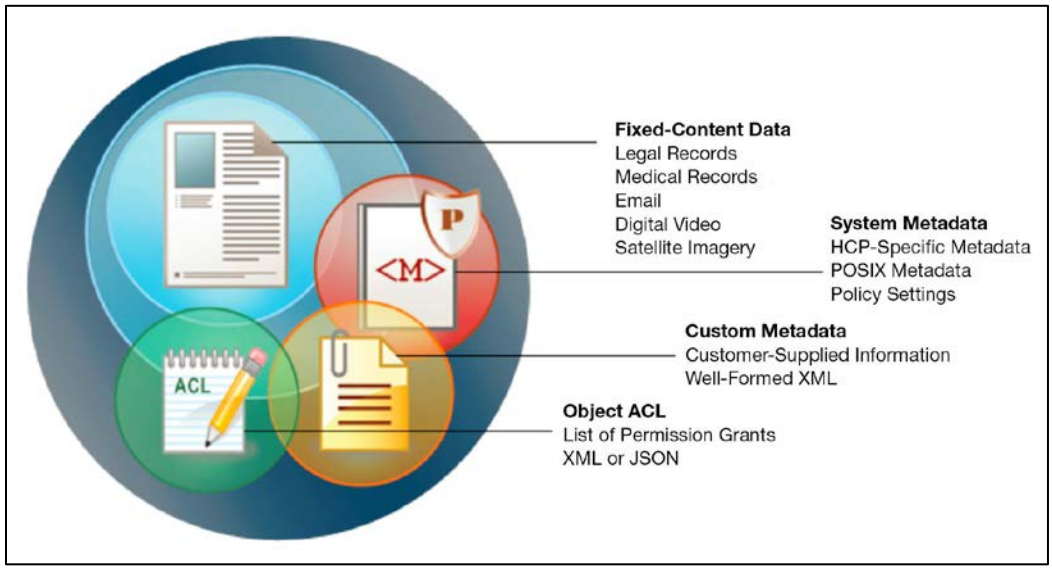

System metadata is composed of 28 properties that include the date and time the object was added to the namespace (ingest time), the date and time the object was last changed (change time), the cryptographic hash value of the object along with the namespace hash algorithm used to generate that value, and the protocol through which the object was ingested. It also includes the object's policy settings, such as number of redundant copies, retention, shredding, indexing and versioning. POSIX metadata includes a user ID and group ID, a POSIX permissions value and POSIX time attributes.

Custom metadata is optional, user-supplied descriptive information about a data object that is usually provided as well-formed XML. It is utilized to add more descriptive details about the object. This metadata can be utilized by future users and applications to understand and repurpose the object content. HCP supports multiple custom metadata fields for each object.

The Importance of Metadata

Custom metadata brings structure to unstructured content. It enables building of massive unstructured data stores by providing means for faster and more accurate access of content. Custom metadata gives storage managers the meaningful information they need to efficiently and intelligently process data and apply the right object policies to meet all business, compliance and protection requirements. Structured custom metadata (content properties) and multiple custom metadata annotations take this capability to the next level by helping yield better analytic results and facilitating content sharing among applications. Custom metadata can be divided into multiple partitions, so that multiple users and applications can work with the same object without impacting each other's metadata.

In many cases, the metadata is more valuable than the object itself. An individual X-ray is not that useful beyond the doctor and the patient. But when that image is stored alongside thousands of others, all with well-described metadata, trends can be discovered, connections made and insights revealed (see Search section).

Object ACL is optional, user-supplied metadata containing a set of permissions granted to users or user groups to perform operations on an object. ACLs control data access at an individual object level and are the most granular data access mechanism.

Store Objects

HCP access nodes share responsibility for knowing where content is stored. HCP stores fixed content file data separately from its metadata, placing them in separate parallel data structures. For scaling purposes, HCP nodes also maintain a hash index, which is a sharded database that is distributed among all HCP access nodes. The **hash index** provides the content addressable lookup function to find data. Each node is responsible for tracking a subset of the index called a region, which tells it where to find data and metadata.

Upon receiving a new file, any receiving node is able to write the fixed content file portion to storage it owns, as directed by the assigned service plan. It then computes a hash of the pathname, adds it to the object's system metadata along with the object's location, and forwards it to the node responsible for tracking the hash index region. HCP protects its index with metadata protection level of 2 **(MDPL2)**, which means it will store two copies, saved on different nodes. There is one authoritative copy, and at least one backup copy. A write is not considered complete until all MDPL copies are saved. The actual file is stored in a storage pool defined by the tenant administrator. Storage pools can be constructed with disks inside an HCP access node, HCP S nodes or SAN storage disks (see Hardware Overview).

Read Objects

Upon receiving a read request file, the HCP node computes a hash using the objects pathname. If it manages the particular hash index region, it can look up the object's location and fulfill the request. If it doesn't manage the hash region it can query the owner node for the files location (see Networking). In the case of a node failure, it can query the node with the backup hash index. In the case of a whole site failure, DNS can redirect the request to any surviving cluster participating in namespace replication.

Open Protocols and SDK

While HCP supports 100% of the **S3** and **Swift** API operations needed for CRUD programming (create, read, update and delete), many new cloud applications being developed still favor its native **REST** Protocol for HCP. This protocol is more full-featured than S3 and Swift, providing insight into HCP's physical infrastructure, its powerful retention, multitenancy, search and metadata capabilities.

To ease the transition to REST, developers can choose the **JAVA SDK** for HCP, complete with libraries, builder patterns and sample code. The SDK provides a fast track path to new applications that need to operate at a global scale, and with users who expect access from virtually any device that supports a network connection.

For those not quite ready to shed their old file access methods, HCP supports four legacy protocols that include **NFS v3, CIFS (SMB 3.1.1), SMTP** and **WebDAV**. Anything written with these APIs can also be accessed with any of the REST API, with directories and filenames intact.

Over 100 ISVs have applications compatible with HCP cloud storage, and enjoy access to Hitachi partner programs, where they can download HCP evaluation software and participate in forum discussions.

HCP Data Services

HCP software implements 13 background services, which are listed in the Table 1. These services work to improve the overall health of the HCP system, optimize efficiency and maintain the integrity and availability of stored object data. Services can run either continuously, periodically (on a specific schedule), or in response to certain events. The system-level administrator can enable, disable, start or stop any service and control the priority or schedule of each service. These controls include running a service for longer periods, running it alone or assigning it a higher priority. Control runtime system loading by limiting the number of threads that the service can spawn, using simple high, medium and low designations.

All scheduled services run concurrently but autonomously to each other, and thus each service may be simultaneously working on different regions of the metadata database. Each service iterates over stored content and eventually examines the metadata of every stored object. On a new HCP system, each service is scheduled to run on certain days during certain hours. If a particular service completes a full scan in the allotted period, the service stops. If it does not finish, the service resumes where it left off at its next scheduled time slot. After completing a scheduled scan interval, the service posts a summary message in the HCP system event log.

Table 1. HCP Background Services

Service	Description
Capacity Balancing	Attempts to keep the usable storage capacity balanced (roughly equivalent) across all storage nodes in the system. If storage utilization for the nodes differs by a wide margin, the service moves objects around to bring the nodes closer to a balanced state.
Compression	Compresses object data to make more efficient use of physical storage space.
Content Verification	Guarantees data integrity of repository objects by ensuring that a file matches its digital hash signature. HCP repairs the object if the hash does not match. Also detects and repairs metadata discrepancies.
Deduplication	Identifies and eliminates redundant objects in the repository, and merges duplicate data to free space.
Disposition	Automatic cleanup of expired objects. A namespace configuration policy authorizes HCP to automatically delete objects after their retention period expires.
Garbage Collection	Reclaims storage space by purging hidden data and metadata for objects marked for deletion, or left behind by incomplete transactions (unclosed NFS or CIFS files).
Scavenging	Ensures that all objects in the repository have valid metadata, and reconstructs metadata in case the metadata is lost or corrupted.
Migration	Migrates data off selected nodes or Hitachi storage arrays so they can be retired.
Protection	Enforces data protection level (DPL) policy compliance to ensure the proper number of copies of each object exists in the system.
Replication	Copies one or more tenants from one HCP system to another to ensure data availability and enable disaster recovery.
Shredding	Overwrites storage locations where copies of the deleted object were stored in such a way that none of its data or metadata can be reconstructed, for security reasons. Also called secure deletion. The default HCP shredding algorithm uses three passes to overwrite an object.
Storage Tiering	Determines which storage tiering strategy applies to an object; evaluates where the copies of the object should reside based on the rules in the applied service plan.
Geodistributed Erasure Coding (Geo-EC)	Geodistributed erasure coding can be applied when HCP spans three or more sites. This technology provides 100% data availability despite whole site level outages. Geo-EC deployments consume 25-40% less storage than systems deployed with simple mirror replication.

Autonomic Tech Refresh (ATR)

Autonomic tech refresh embodies the vision that software and data will outlive the hardware hosting it (see Figure 12). This built-in migration service enables organizations to move a “live” HCP onto new hardware; replace old servers, or siphon content from old storage as a rate-controlled background process, and write it to new storage. With ATR, there are no disruptions to customer business since applications continue to run normally. This forethought to maintenance is rather unique and testament to HCP’s long-term product support strategy.

Figure 12. Autonomic Tech Refresh: Choose, Review and Confirm

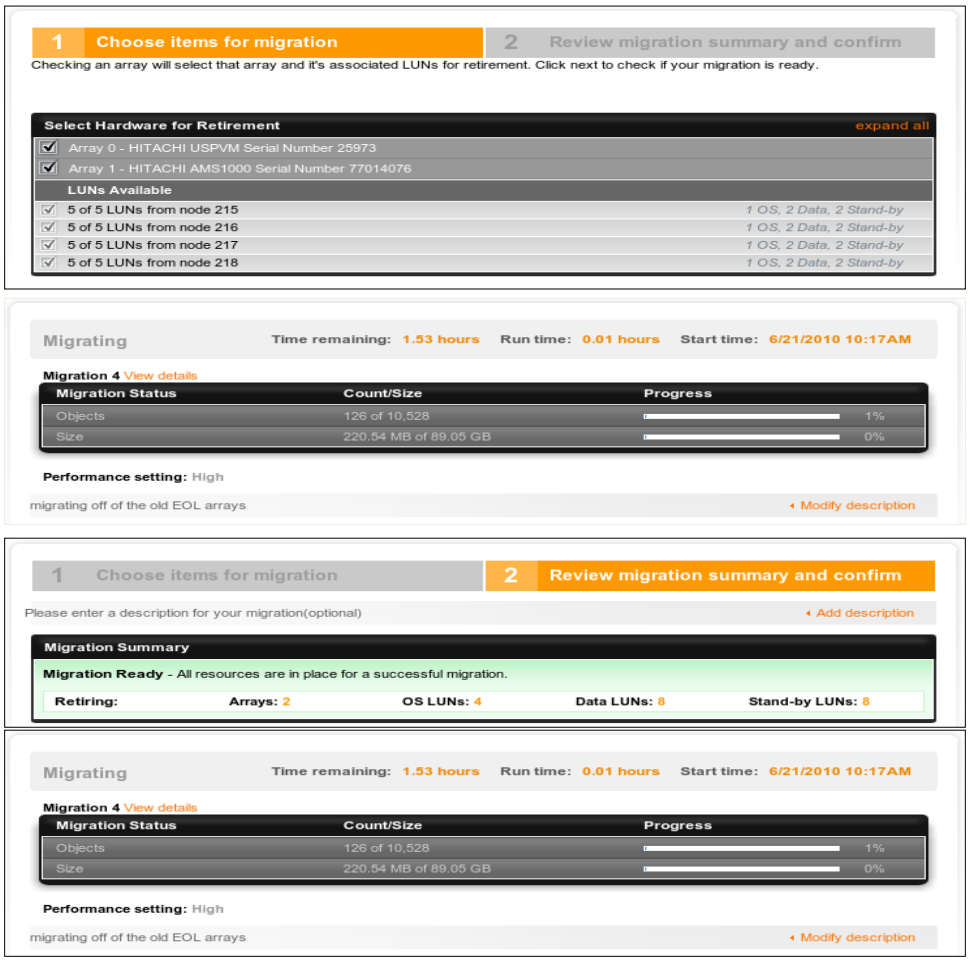

HCP Replication Topologies and Content Fencing

HCP offers multisite replication technology called global access topology. With these bidirectional, active-active replication links, globally distributed HCP systems are synchronized in a way that allows users and applications to access data from the closest HCP site for improved collaboration, performance and availability.

Metadata-Only Replication

Choose to replicate entire objects or just object metadata. A metadata-only strategy allows all clusters to know about all objects, but it controls placing object payload only where needed while saving on WAN costs.

Content Fencing

One practical use case for metadata-only replication is to create data fences, which allow organizations to share data, but ensure it stays hosted within a specific country or continent boundary. In this model, HCP replicates

metadata but withholds mass movement of data files. Applications at the remote end are able to see files and directory structures, search metadata fields and even write content. In all cases, the final permanent resting place for the object is at the source. Global access topology supports flexible replication topologies that include chain, star and mesh configurations (see Figure 13).

Figure 13. Chain, Star and Mesh Configurations

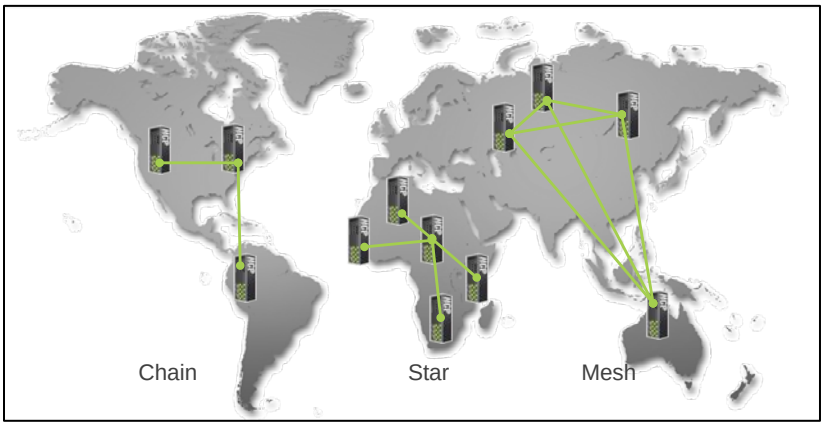

The replication process is object-based and asynchronous. The HCP system in which the objects are initially created is called the primary system. The second system is called the replica. Typically, the primary system and the replica are in separate geographic locations and connected by a high-speed wide area network. The replication service copies one or more tenants or namespaces from one HCP system to another, propagating object creations, object deletions and metadata changes. HCP also replicates tenant and namespace configuration, tenant-level user accounts, compliance and tenant log messages and retention classes.

Geodistributed Erasure Coding

With replication links established between three or more sites, HCP software offers a geodistributed erasure code service for greater storage efficiencies and cost savings. The service operates on objects to provide site-level disaster recovery using fewer bytes than a mirrored replica. The service is electively applied to a namespace (bucket) along with settings for activation. Administrators decide when objects should transition to a geodistributed state. New objects can remain as whole objects at all sites for a brief period, providing fast, low latency readback. When this period ends, the object transitions to an efficient geo-EC state that can reduce storage consumption by up to 40%, and keep objects eligible for compression and deduplication.

Number of Sites	2	3	4	5	6
Geo-EC Topology	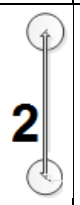	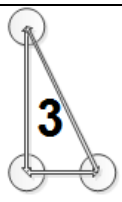	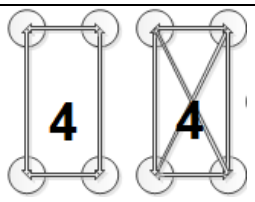	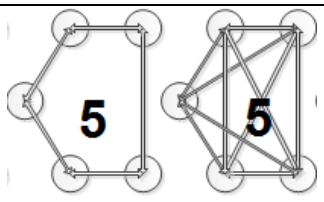	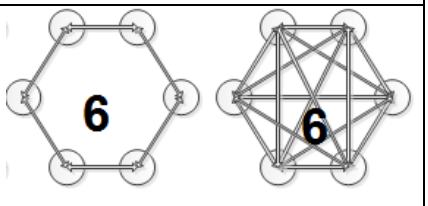
Capacity Savings	**0%**	**25%**	**33%**	**37%**	**40%**

Search

With Content Platform, you have access to metadata and content search tools that enable more elegant and automated queries for faster, more accurate results. Through these features you can gain a better understanding of the content of stored files, how content is used and how objects may be related to one another. This understanding can help you to enable more intelligent automation, along with big data analytics based on best-in-class metadata architecture.

HCP software includes comprehensive built-in search capabilities that enable users to search for objects in namespaces, analyze a namespace based on metadata, and manipulate groups of objects to support e-discovery for audits and litigation. The search engine (Apache Lucene) executes on HCP access nodes and can be enabled at both the tenant and namespace levels. HCP supports two search facilities:

1) A web-based user interface called the **search console** provides an interactive interface to create and execute search queries with “AND” and “OR” logic. Templates with dropdown input fields prompt users for various selection criteria such as objects stored before a certain date or larger than a specified size. Clickable query results are displayed on-screen. From the search console, search users can open objects, perform bulk operations on objects (hold, release, delete, purge, privileged delete and purge, change owner, set ACL), and export search results in standard file formats for use as input to other applications.

2) The **metadata query API** enables REST clients to search HCP programmatically. As with the search console, the response to a query is metadata for the objects that meet the query criteria, in XML or JSON format.

In either case, two types of queries are supported:

- An object-based query locates objects that currently exist in the repository based on their metadata, including system metadata, custom metadata and ACLs, as well as object location (namespace or directory). Multiple, robust metadata criteria can be specified in object-based queries. Objects must be indexed to support this type of query.
- An operation-based query provides time-based retrieval of objects transactions. It searches for objects based on operations performed on the objects during specified time periods. And it retrieves records of object creation, deletion and purge (user-initiated actions), and disposition and pruning (system-initiated actions). Operation-based queries return not only objects currently in the repository but also deleted, disposed, purged or pruned objects.

Multiple Metadata Annotations

Each HCP object supports up to 10 free-form XML metadata annotations up to 1GB total. This gives separate teams freedom to work and search independently. An analytics team may add annotations specific to their applications, which are different from the billings applications. XML annotation can provide significant advantage over simple key value pairs because the search engine can return more relevant results with XML. Consider Table 2:

Table 2. Metadata Annotation Example

This example XML record represents a single annotation	Key-value pair
<Record> <Dr>John Smith</Dr> <Patient>Jonn Smith</Patient> <Address>St John Smith Square</Address> </Record>	Dr=John smith Patient= Jonn Smith Address=St John Smith Square</Address>

Now imagine a search where you want the objects related to doctor named “John Smith.” The XML record allows you to pinpoint search results to this field, whereas the key value pair will produce a much larger set of search hits. As object count grows to millions and billions, key-value searches can quickly become slow and arduous.

Hardware Overview

HCP software is deployed on commodity x86 hardware or hypervisors as a distributed storage system. From a hardware perspective, each HCP system consists of the following categories of components:

- HCP G series access nodes (servers) or virtual machine instances.
- Networking components (switches and cabling).
- Infrastructure components (racks and power distribution units).
- Physical storage pools.
  - Access-based storage.
  - HCP S erasure code storage nodes (Ethernet-attached storage).
  - SAN-attached storage.
  - Extended or public cloud storage.

Access Nodes (HCP G Nodes)

If delivered as an appliance, the nodes are constructed from conventional x86 off-the-shelf 2U servers called an HCP G series node (see Figure 14). Each node is configured with multicore CPUs, DRAM, 10Gb networking, SSD (optional), and up to 12 internal SAS disks protected with RAID-6.

Figure 14. HCP G Nodes

- The most minimal HCP system will have four access nodes with local disk storage.
- Electively, HCP G nodes can be purchased with Fibre Channel adapters to utilize Hitachi SAN arrays for bulk HCP object storage.
- Access nodes can be deployed as virtual appliances (see virtual access nodes).
- Upgrade DRAM, SSD or internal disks anytime, as server space permits.

Flash-Optimized Option

HCP G series nodes may be populated with a pair of solid-state disks (SSD), configured in a RAID-1 mirror. The drives are not used as a cache, and do not contain user file data. Rather, they are used for the singular purpose of accelerating database performance related to operating HCP’s content-addressable hash index. The judicious and focused application of SSD was shown to improve read and write performance, especially when the number of objects managed per node grow very large (>100M objects).

Virtual Access Nodes

This deployment model lets administrators install virtual instances of HCP access nodes on hardware they supply. HCP access nodes (both virtual and physical) can run with as little as 1TB of licensed storage. At present, HCP supports running within KVM or VMware hypervisors, ESXi versions 5.5 or 6.x. Best practice installations would

ensure that virtual machines (VMs) reside on separate bare metal nodes to ensure the cluster operates with maximum availability. Capacity expansion can be accomplished with HCP storage appliances called S nodes or customer-owned Fibre Channel storage.

Storage Nodes (HCP S10, HCP S30)

Ethernet attached storage nodes (HCP S Nodes) provide erasure code (EC) protected storage pools that are managed by access nodes. These appliances allow an HCP to scale storage capacity independently from its compute capacity. All HCP S nodes feature two independent x86 controllers and SAS Enterprise disk drives. Control and data path occurs over bonded and redundant 10Gb Ethernet connections (four total ports). They come in two physical varieties:

- **HCP S10:** One 4U tray with 60 disks, two integrated x86 controllers.
  - 560TB (430TB usable).
- **HCP S30:** Up to sixteen 4U tray with 954 disks, and two 2U x86 servers.
  - 9.54PB of raw storage delivered in two racks.
  - A minimum HCP S30 would have three half trays (540TB raw), and grow in tray increments over time.

Protection: Reed-Solomon 20+6 Erasure Coding (see Figure 15). Immediately following insertion into the storage pool, new disks are divided into 64MB chunks called an extent. These extents are grouped into an EC-protected extent group by choosing and writing to 26 extents, each on a different disk. As objects are ingested, they are efficiently stored within the extent groups. The extent group forming constraint guarantees that any disk loss will only affect at most one extent from a particular extent group. With 20+6 erasure coding, objects within an extent group remain readable despite the simultaneous loss of six drives containing the data of that extent group.

Figure 15. HCP S30 Node With 90 drives: Sample EC Extent Group Allocation

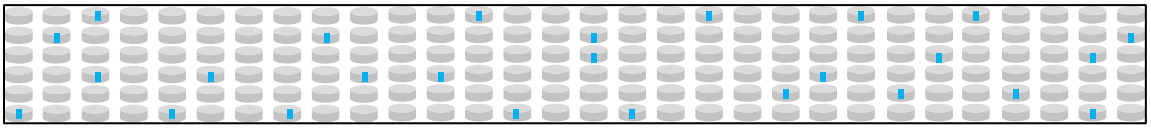

Common HCP S series node properties:

- On-demand initialization (no long formatting).
- Self-optimizing algorithms for extent group repair and auto rebalance.
- Rebuild duration is a function “bytes used” rather than “bytes available.” For example, if a 6TB disk fails, and it contained 1TB of written extents, the rebuild algorithm needs to rewrite only 1TB of extents.
- No “hot spares.” Extent groups have flexibility in placement, thus all drives are used all the time. When a drive fails, erasure coding just recreates missing data using the free capacity on remaining drives. It returns to optimal without adding new disks.
- Extend service intervals. Because failed disks no longer imply a degraded platform, it’s possible to defer replacement service until the number of disks with free capacity runs low. System automatically notifies what that point is approaching.
- Unlike RAID, rebuild I/O is not throttled by the write capabilities of one drive. When a drive fails, software repairs and redistributes erasure-coded extents to all active drives.
- Storage efficiency and rebuilds are file-size agnostic because HCP S series nodes utilize 64MB extents. These nodes can store literally hundreds of small files in a single extent. Thus, one write I/O to repair one extent member (64MB chunk), effectively repairs hundreds of files. In this way, HCP S series nodes are optimized for both large and small files.

Extended Storage (NFS and Public Cloud)

With **adaptive cloud tiering** technology, HCP access nodes can leverage storage from third-party sources to construct storage pools built with:

- On-site NFS or S3 compatible storage.
- Off-site public cloud storage sources Amazon S3, Microsoft Azure, Google Cloud Storage and Hitachi Cloud Services.

ACT makes it possible to construct hybrid HCP configurations that share resources between public and private clouds. With simple edits of an existing service plan, the system-level administrator can easily shift cluster composition:

- Easily scale HCP storage capacity up or down as needed. With elastic capacity, make room on-premises for a sudden capacity need by encrypting and tiering out older objects to third-party cloud.
- Gain cloud-broker capabilities. For example, easily switch public cloud allegiances towards the best long-term storage rates, or better reliability.
- Encrypt content you're not quite ready to part with and direct it to public cloud for the remainder of its retention period.

Service plan adjustments are transparent to client applications from a protocol perspective. However, policy changes may temporally increase total I/O workload if the changes trigger a background data migration. The impact can be mitigated by setting tiering-priority levels.

Networking

All HCP systems feature at least two physically separate Ethernet networks referenced as the private back-end and public front-end network. All networks are 10Gb capable, and all are constructed with a bonded port pair, which connect to independent switches.

Private back-end network: The two switches comprising the isolated back-end network carry traffic vital to internode communication; they are provided as part of an HCP appliance with no outside access. To ensure maximum network reliability and availability, the design requires two unstacked switches with options for optical or copper media. While a choice of 1G or 10G switches is offered, a pair of 1Gb switches provides ample communication bandwidth for many configurations.

Front-end network: This VLAN-capable network connects to a customer-supplied switch infrastructure. This network carries application read/write traffic, management traffic and HCP S node traffic if present using VLANs. The recommended front-end setup would include two independent switches that support 1/10Gb Copper (RJ45), or optical 10Gb SFP+ connections.

External communication with HCP is typically managed DNS, which round-robins client requests across all nodes to ensure maximum system throughput and availability. With DNS, clients reference a domain name rather than a specific node or IP address. Typically, a subdomain or delegate is defined within the corporate DNS and all nodes in the HCP cluster are listed as hosts. HCP uses load-balancing internally and manages all inbound read/write requests, passing them to the least busy node for execution. This network can be configured in native IPv4, native IPv6, or dual IPv4 and IPv6 modes where each virtual network will support either or both IP versions. IPv6 is mandated by many government agencies, and necessary to support very large scale networks.

Configurations Using Only HCP G Nodes

HCP systems with a server-only disk configurations use their internal disks to store files, metadata and the appliance software (see Figure 16). A minimal cluster consists of at least four HCP G series nodes populated with at least six disks to form one RAID-6 group (4D+2P). At any time, server capacity can be expanded with a second set of six disks to create another RAID-6 group.

In this configuration, the cluster is normally operating with a data protection level of 2 (DPL2), which means each file object (and custom-metadata) is stored twice utilizing separate protection sets on different nodes. Note: DPL is distinct and unrelated to MDPL, which controls the number of system-metadata copies. By default, the MDPL is also set at level 2. All of this ensures the cluster will remain fully functional in the event of a node loss.

A single site, four-node HCP G series node cluster with all disks populated, provides 56TB of usable capacity when storing objects with DPL2, and 112TB if storing with DPL1. DPL1 configuration should only be considered if deploying HCP across two sites using HCP replication. HCP licensing policy permits operation with as little as 4TB. Later the license can be increased to accommodate increased capacity needs. A license can also exceed the capacity of physically attached storage.

Figure 16. Configuration Using HCP G10 Nodes

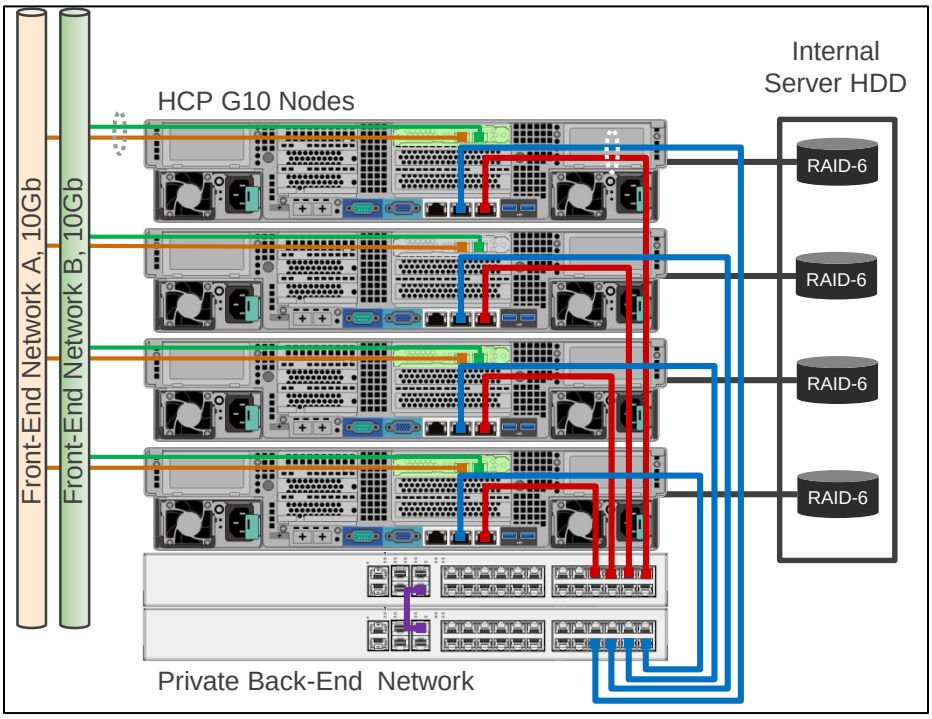

The capacity of HCP clusters configured with server-only storage can be expanded in two-node increments. Scaling capacity in this fashion may needlessly increase the compute capability of the cluster, since each increment of storage pulls in another server pair with 24 CPUs between them.

Situations that require significant capacity scaling should consider HCP S series nodes because they almost always deliver better economics in the form of lower cost per gigabyte.

HCP SAN-Attached Configurations

An HCP system with SAN-attached storage uses Fibre Channel arrays from Hitachi to store fixed-content file objects and custom metadata, but saves portions of system metadata on its internal disks (see Figure 17). A minimal cluster consists of at least four HCP G series nodes, each with a two-port Fibre Channel card and each with six disks to form one RAID-6 group (4D+2P). The internal disks hold only the appliance operating system and HCP's system metadata database.

Figure 17. HCP SAN-Attached Configurations

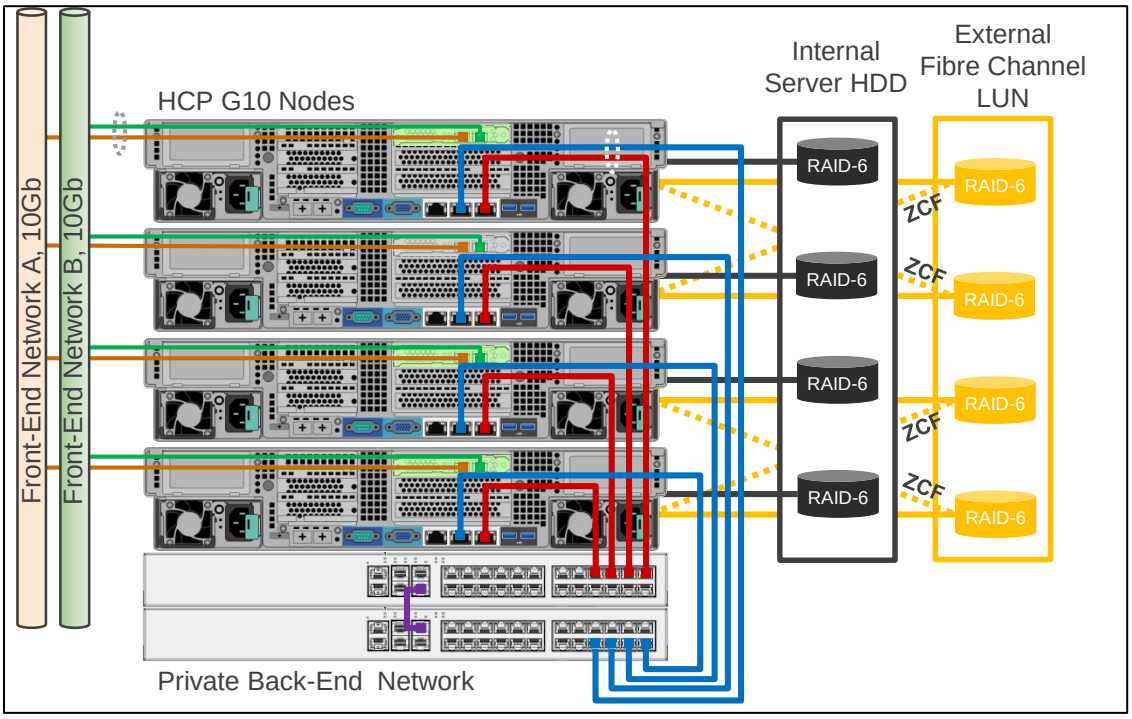

Fibre Channel arrays can supply each node with up to 1PB of storage capacity, providing much better storage density than the server-only storage configurations. Moreover, these configurations can be operated with a data protection level of 1 **(DPL1)** while maintaining high data availability because of the shared storage capabilities of a SAN. This is feasible because HCP software employs a feature called zero copy failover (ZCF). It works by pairing two nodes, and enabling them with multipathing, such that they can see and access one another's Fibre Channel storage LUNs, a process called cross-mapping. When the cluster is optimal, the LUNs managed by a node during normal operation are considered primary LUNs; the LUNs visible on the other node are considered standby LUNs. In the event of a node failure, cross-mapping of LUNs allows a surviving node to assume ownership of the standby LUNs. This shared storage configuration enables capacity scaling without increasing access node count. If scaling access nodes, they must be added in groups of two to facilitate ZCF pairing.

MDPL2 along with ZCF ensure the data stored in the cluster will remain fully functional in the event of a node loss.

Often, these configurations will not require fabric because many Hitachi arrays provide enough ports to directly connect up to 16 access nodes. Configurations exceeding these limits may be purchased with either Brocade or Cisco Switches (not shown in the diagram).

It merits mention that HCP S series nodes can also be utilized for capacity expansion, and offer very similar scaling and failover advantages of Fibre Channel.

Capacity Scaling With HCP S Node

Ethernet attached HCP S nodes provide capacity expansion options for any HCP configuration, including HCP for VMware vSphere-ESXi, HCP for KVM and previous generation nodes (HCP 500 or HCP 300) (see Figure 18). These appliances combine enterprise HCP software and commodity hardware to achieve more scale at a lower cost. HCP S series nodes easily rival and sometimes exceed the storage density found on Fibre Channel arrays. Moreover, as an Ethernet connected storage resource, they offer similar shared access properties of a SAN. As such, these configurations can be safely operated with a data protection level of 1 **(DPL1)**. To ensure cluster reliability, these clusters use MDPL2. MDPL2 along with HCP S series node storage pools ensure the cluster will remain available and fully functional despite loss of an access node.

In the example below, a single HCP S10 node is added to the front-end network using HTTPS and VLANS. HCP G series access nodes register the HCP S10 node as a storage component, and then add it to storage pool. In this configuration, HCP access nodes use their internal disks to store metadata, while file objects are stored on the HCP S series node(s). This additional capacity is then available to all nodes for new REST writes, or for tiering existing object stored on internal server disks.

Figure 18. Capacity Scaling With HCP S Node

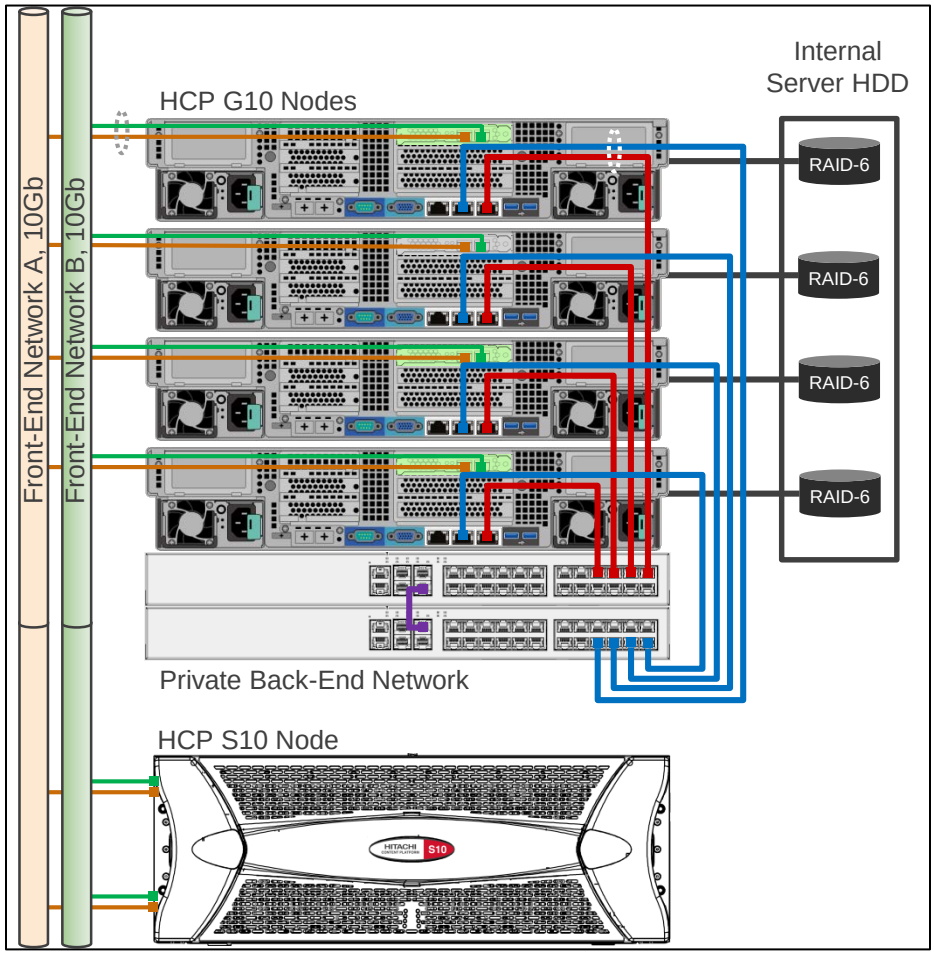

Four HCP G10 and HCP S10 Racking Options

There are numerous racking options that combine HCP G10 with HCP S10 (see Table 3).

The HCP “appliance” configuration (left) combines up to six HCP G10 and three HCP S10. For larger clusters, HCP G series nodes (middle) and HCP S10 nodes (right) are normally deployed in different racks.

Table 3. HCP G10 and HCP S10 Racking Options

		HCP Appliance HCP G10 and S10			HCP G10 Base (1/10G)			HCP S10 Only		
U				U			U			U
42	PDU Installation Area		PDU Installation Area	42		PDU Installation Area	42		PDU Installation Area	42
41		FCS (option)		41			41	Empty		41
40		HCP G10		40	Empty		40			40
39		or FCS (option)		39			39			39
38		HCP G10		38			38			38
37		or FCS (option)		37			37	HCP S10		37
36		HCP G10		36	HCP G10		36	4U 60 HDD		36
35				35	or FCS (option)		35			35
34		HCP G10		34	HCP G10		34			34
33				33	or FCS (option)		33	HCP S10		33
32		HCP G10		32	HCP G10		32	4U 60 HDD		32
31				31	or FCS (option)		31			31
30		HCP G10		30	HCP G10		30			30
29				29	or FCS (option)		29	HCP S10		29
28		1U Ethernet Switch		28	1U Ethernet Switch		28	4U 60 HDD		28
27		1U Ethernet Switch		27	1U Ethernet Switch		27			27
26		FC Switch (if necessary)		26	FC Switch (if necessary)		26			26
25		FC Switch (if necessary)		25	FC Switch (if necessary)		25	HCP S10		25
24				24	HCP G10		24	4U 60 HDD		24
23				23	or FCS (option)		23			23
22				22	HCP G10		22			22
21				21	or FCS (option)		21	HCP S10		21
20		Empty		20	HCP G10		20	4U 60 HDD		20
19				19	or FCS (option)		19			19
18				18	HCP G10		18			18
17				17	or FCS (option)		17	HCP S10		17
16				16	HCP G10		16	4U 60 HDD		16
15				15	or FCS (option)		15			15
14				14	HCP G10		14			14
13		HCP S10		13	or FCS (option)		13	HCP S10		13
12		4U 60 HDD		12	HCP G10		12	4U 60 HDD		12
11				11	or FCS (option)		11			11
10				10	HCP G10		10			10
9		HCP S10		9	or FCS (option)		9	HCP S10		9
8		4U 60 HDD		8	HCP G10		8	4U 60 HDD		8
7				7			7			7
6				6	HCP G10		6			6
5		HCP S10		5			5	HCP S10		5
4		4U 60 HDD		4	HCP G10		4	4U 60 HDD		4
3				3			3			3
2				2	HCP G10		2			2
1				1			1			1
U				U			U			U

HCP = Hitachi Content Platform, FCS = File and Content Solution, PDU = Power Distribution Unit

From a pricing standpoint, HCP S10 nodes are generally recommended when required capacity growth is < 1PB.

- Each HCP S10 tray provides up to 600 TB of raw capacity

HCP S30 Racking Options

An HCP S30 node is always deployed in a rack separate from the access nodes (see Table 4). This policy ensures an “entry” three-tray HCP S30 node (shown left) can be expanded over time. A full HCP S30 node spans two racks. HCP has presently qualified connecting up to 80 HCP S30 nodes to a single HCP cluster.

Table 4. HCP S30 Racking Options

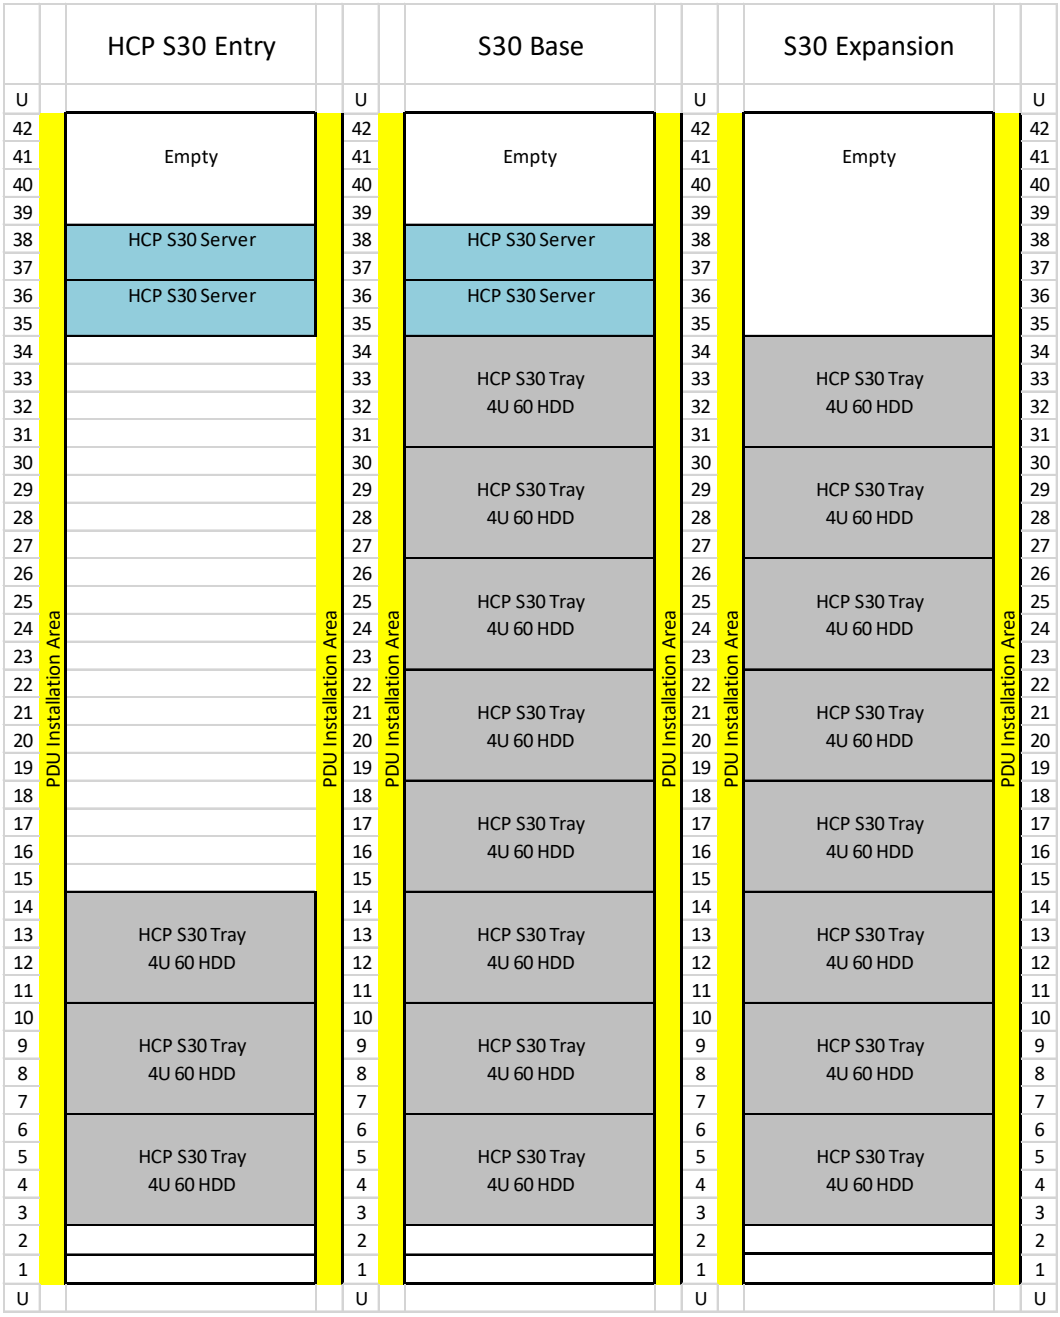

HCP = Hitachi Content Platform, PDU = Power Distribution Unit

From a pricing standpoint, HCP S30 nodes are recommended for capacity growth in 1PB or greater increments. Scale in single 4U tray increments, approximately 0.25PB each.

- Max node configuration: 16 full trays: 954 disks, 10TB.
- Raw 9.54PB (~7.34PB usable).

Security

With any cloud deployment, whether it's public, private or hybrid, security of data is paramount. Hitachi Content Platform is an extremely secure and efficient object storage with a comprehensive suite of data protection and

security must-haves. HCP is readily able to store and protect at scale, which is crucial to safeguarding all of the data all of the time for the business. Some of those features include:

Ports to Support Web Protocols and Administration: HCP cloud storage software will require ports to support web protocols and administration. For example, HCP needs TCP/IP ports 80 (HTTP) and 443 (HTTPS, also known as HTTP over TLS), 8000 (admin) to conduct both appliance management and data storage tasks.

Host-Based Firewalls: HCP follows security best practices and disables all external ports and processes that are not required by the software. Moreover, each HCP node runs a firewall that is designed to block all ports not associated with an active HCP service.

Secure Remote Service: All remote service is performed using SSH (with a 2048-bit key, which is limited to Hitachi Vantara's support organization). Organizations are encouraged to disable this SSH access unless service is needed, at which time, SSH access can be enabled with the system-level administrator.

SSL Server Certificates: HCP requires one SSL server certificate (self-generated or uploaded PKCS12) for each defined domain to prove authenticity to clients.

Encryption: If enabled, HCP utilizes a AES-256 block cipher with a key (or block) length of 256 bits. This is the cipher required for FIPS 140-1 compliance. Cypher keys are protected with Shamir Shared Secret mechanism (key built only into volatile memory of an HCP system, and can only be rebuilt with a quorum of nodes).

User Authentication: In addition to local user accounts, HCP supports enterprise identity services: Microsoft Active Directory, RADIUS and OpenStack keystone.

Per Object ACL (access control list): Utilize ACLs to control data access at an individual object level. ACLs provide a more granular data access controls that limit the permissions granted to users or user groups, as well as the operations they can use.

Dedicated Management Network: Administrative tasks can be isolated on VLANs or physically separate Ethernet ports available on HCP servers.

Conclusion

Hitachi Content Platform offers an intelligent solution to the unstructured data dilemma with an object storage system. It eliminates the limitations of traditional storage systems and can efficiently scale to virtually unlimited storage. Hitachi Content Platform allows IT to perform all essential data management tasks from a single system. And, it treats file data, file metadata and custom metadata as a single object whose life cycle is managed as a series of planned storage tier transitions. With secure multitenancy, HCP helps organizations manage huge pools of capacity by dividing it into smaller virtual object stores, and delegating partial control of these to owners, called tenants. Within bounds, the tenant owners are given authority to assign capacity quotas, set configurable attributes and choose service levels that support their application needs. This approach allows the object store to support a wide range of workloads, such as content preservation, data protection and content distribution from a single physical infrastructure.

Perhaps most exciting, we see markets rapidly adopting object storage as a technology that can be used for any and all storage workloads; it is no longer limited to large and infrequently accessed data sets. HCP cloud storage is the foundational part of a larger portfolio of solutions that includes Hitachi Data Ingestor for elastic, backup-free file services and Hitachi Content Platform Anywhere for synchronization and sharing of files and folders with the bring-your-own-device class of users. One infrastructure is far easier to manage than disparate silos of technology for each application or set of users. By integrating many key technologies in a single storage platform, Hitachi Vantara's 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, stay agile over the long term, and address future change and growth.

Additional Resources

- Gartner Critical Capabilities for Object Storage: https://www.gartner.com/doc/reprints?id=1-33E2S6I&ct=160413&st=sb
- IDC Hitachi Content Platform: End-to-End Portfolio for the 3rd Platform https://www. hitachivantara.com/en-us/pdf/analyst-content/hitachi-idc-technology-assessment-hcp.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-white-paper-hcp-solve-data-protectio-and-security-issues-for-big-data-and-cloud.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-whitepaper-achieve-secure-cost-optimized-data-mobility.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-data-mobility-to-secure-content-outside-the-data-center.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-data-mobility-to-secure-content-outside-the-data-center.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-white-paper-create-a-relational-distributed-object-store.pdf
- https://www.hitachivantara.com/en-us/pdf/white-paper/hitachi-white-paper-comparing-cost-and-performance-of-replication-and-erasure-coding.pdf
- https://community.hitachivantara.com/community/developer-network

Hitachi Vantara

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hitachivantara.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com
Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com

WP-553-C Hitachi Vantara December 2017