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

Hitachi Data Ingestor: elastic, backup-free, cloud gateway that delivers file services beyond the data center.

Used with Hitachi Content Platform, Hitachi Data Ingestor provides file services to distributed consumers of IT, such as remote and branch offices (ROBOs) or cloud storage users.

HITACHI
Inspire the Next
Hitachi Data Systems

Reduce complexity and cost of file services beyond the data center, simplify cloud deployments

Hitachi Data Ingestor (HDI) is an elasticscale, backup-free cloud file server with advanced storage and data management capabilities.

With this solution, organizations can greatly reduce the cost and complexity of providing IT services to geographically dispersed locations or cloud consumers.

Elastic scale and backup-free

Hitachi Data Ingestor provides a standard connection, or gateway, into the core data center without requiring application recoding and without changing the way individuals interact with storage today. Because HDI acts as a caching device, it provides users and applications with seemingly endless storage and a host of newly available capabilities.

These abilities include roaming home directories, encryption of data at rest and in transit, and easy migration of data from existing storage to HDI. All files are automatically replicated to the core infrastructure in your data center.

HDI appears as a standard storage device to users and applications. However, for IT departments, it drastically simplifies the deployment, provisioning, management and protection of data at remote sites, branch offices and/or cloud service customer sites.

HDI automatically copies content out of its internal cache and into the core object storage cloud provided by Hitachi Content Platform (HCP). This action ensures robust data protection with easy recovery and provides everexpanding storage capacity for new content.

These capabilities reduce management time and cost by eliminating the need to constantly manage capacity, utilization, protection, recovery and performance of the system. Once a file is copied into the core infrastructure, it stays in the HDI file system until HDI detects that free space has reached a predetermined ceiling.

At that time, HDI reduces the least active files to pointers, effectively creating more local capacity and keeping the most important and frequently used files in local storage. Because HDI retains the pointers to relocated files, they can always be accessed via the cloud.

HDI supports corporate governance and compliance requirements through its built-in, file-level “write once, read many” (WORM) functionality option. And, independent software vendor applications can adopt HDI with minimal certification.

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

Features

Hitachi Data Ingestor presents a standards-based file system interface that is tightly integrated with Hitachi Content Platform to provide seamless access and a wide range of advanced storage features. HDI uses HTTP/HTTPS to securely move data over a local or wide area network and into HCP.

HDI features include:

- Provides local and remote access to HCP for clients over Server Message Block (SMB) and Network File System (NFS) protocols.
- Delivers elastic storage capacity with no need for local backup.
- Physically migrates content to a central HCP and maintains a local link to the migrated content.
- Provides file restore: Retrieves previous versions of a file or even deleted files; maintains file and directory access control within "history."
- Allows content sharing between HDI systems:
  - Single HCP tenant is used for multiple HDI systems, for simplicity.
  - Multiple HDIs can read from asingle HCP namespace.
  - One HDI has write capability; others have read capability.
- Allows NAS migration (SMB or NFS) to HDI file system using GUI or CLI: Automatically migrate file data from NAS systems to HDI.
- Supports roaming users by synchronizing content in SMB protocol home directories across a network of HDI systems.
- Provides optional AES 256-bit encryption of content for each HDI file system.
- Provides a management API that enables integration with the HCP management user interface and 3rd-party or homegrown management user interfaces.
- Supports Microsoft® Active Directory® and LDAP authentication for HDI clients.
- Supports HCP tenant and namespace features over SMB and NFS.
- Scales to 400 million files per HDI system.
- Employs intelligent local cache to speed access to HCP content over SMB and NFS.
- Speeds cloud adoption; no need to rewrite applications or change user behavior.

Deployment options

Hitachi Data Ingestor is offered in multiple deployment options: a high-availability, dualnode cluster with internal or external storage, a single node with internal storage, a software-only configuration running on VMware vSphere Hypervisor, and a small form factor remote server version ideal for small offices.

In addition, the single node, VMware and remote server configurations can be remotely configured, provisioned and managed using Hitachi Content Platform Anywhere (HCP Anywhere) and installed at the remote site by nontechnical personnel. Just plug it in, power it up, and it will import everything from HCP Anywhere.

In all configurations, HDI acts as a tiering solution, copying its resident files to HCP, and maintaining access to those files for on-demand recall.

Summary

IT organizations serving distributed consumers benefit from elastic scale, backup-free file services with Hitachi Data Ingestor. HDI coupled with Hitachi Content Platform and HCP Anywhere creates an integrated cloud file services solution to support distributed IT models, such as remote and branch offices and cloud service providers.

With HDI at the edge sending data to HCP in the data center, organizations can manage file services with HCP Anywhere. Distributed IT environments can reduce the cost and complexity of delivering file services beyond the data center.

Companies can empower their mobile workforces and take advantage of next-generation file services that enable content mobility, such as file sync and share, remote access to NAS data, and hybrid cloud storage.

Hitachi Data Ingestor Technical Specifications

Cluster [Hitachi Data Ingestor (HDI) integrated appliance]	▪ Server model: Hitachi Compute Rack 210H (CR 210H). ▪ CPU: 1x E5-2603 (1.80GHz, 4 core). ▪ Memory: 16GB. ▪ I/O: 1GbE and 10GbE (gigabit Ethernet). ▪ Storage: Hitachi Unified Storage 110 (8TB usable). ▪ Footprint: 5U.
Cluster (Diskless)	▪ Server model: CR 210H. ▪ CPU: 1x E5-2603 (1.80GHz, 4 core). ▪ Memory: 16GB. ▪ I/O: 1GbE and 10GbE. ▪ Storage: any Hitachi storage (any supported capacity). ▪ Footprint: 3U.
Single node (DAS)	▪ Server model: Hitachi Compute Rack 220S (CR 220S). ▪ CPU: 1x E5-2420 (1.9GHz, 6 core). ▪ Memory: 12GB. ▪ I/O: 1GbE and 10GbE. ▪ Storage: Internal HDDs (4TB, 8TB or 12TB usable). ▪ Footprint: 2U.

Virtual Machine Appliance	▪ VMware vSphere Hypervisor (ESXi): 4.1 or later.
	▪ Hardware requirements defined by VMware.
	▪ Hardware reference configuration:
	  ▪ CPU: 2x E5-620 (2.40GHz, 4 core).
	  ▪ Memory: 4GB.
	  ▪ Storage: up to 13 2TB LUNs.

HDI Remote Server	▪ ITX Mini Tower.
	  ▪ CPU: 1x G540 DUAL 2.5GHz.
	  ▪ Memory: 4GB.
	  ▪ Storage: Internal HDDs (1TB – 3TB usable).

People who rely on their IT infrastructure rely on Oncore IT