

Hitachi Data Ingestor

A bottomless, backup-free cloud on-ramp and filer, Hitachi Data Ingestor delivers file services to the edge and enables an advanced edge-to-core storage solution.

Reduce Cost and Complexity at the Edge and Simplify Cloud Deployments

Hitachi Data Ingestor (HDI) with Hitachi Content Platform (HCP) creates an integrated offering that provides distributed consumers of IT, such as remote offices, branch offices (ROBO) or cloud storage users, with a seamlessly scalable, backup-free storage solution. Deployed as a minimal-footprint or virtual appliance, HDI sends data from the edge to a core infrastructure, employing 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.

Bottomless, Backup-free Edge Device

Hitachi Data Ingestor provides a standard connection, or on-ramp, into the core data center without requiring application recoding and without changing the way users 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.

You no longer need to back up your edge devices with HDI. All files are automatically replicated to the core infrastructure in your data center. HDI appears as a bottomless storage device to users and applications, since it automatically moves content out of its internal cache and into the core to provide expanding storage capacity for new content. This reduces management time and cost by eliminating the need to constantly manage capacity, utilization and performance of the system. Once a file is moved into the core infrastructure, a pointer is left behind. By reducing files to pointers, HDI constantly frees up local storage capacity so that additional storage never needs to be added at the edge. The files can always be accessed, regardless of where they are physically located.

HDI supports corporate governance and compliance requirements through its built-in, file-level “write once, read many” (WORM) functionality, which is an option for

those who need it. Furthermore, independent software vendor applications can adopt HDI with minimal certification.

Deployment Options

Hitachi Data Ingestor is offered in 2 deployment options to support either local or remote Hitachi Content Platform access. The HDI node(s) may be deployed within the HCP rack as an integrated package, or separately. HDI offers a high availability (HA) dual-node cluster with external storage, a single node with internal storage, and a software-only configuration running on VMware vSphere Hypervisor that also supports HA architecture. In all configurations, HDI acts as a tiering solution, moving its resident files to HCP, and providing links to those files for on-demand recall.

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. Some of the features of HDI are described below:

- Provides local and remote access to a HCP for clients over CIFS and NFS
- Delivers seemingly bottomless storage capacity, back-ended by HCP
- Migrates content to a central HCP and maintains a local link to the migrated content
- Provides file restore
 - Users can retrieve previous versions of a file.
 - Users can retrieve deleted files.
 - File and directory access control is maintained within "history."
 - Feature is available on cluster, single node and virtual machine appliance configurations.
- Allows content sharing between HDI systems
 - Multiple HDIs can be read from a single HCP namespace.
 - One HDI has write capability, others have read capability.
- Enables NAS migration (CIFS) to HDI file system
 - CIFS data can be migrated from NAS systems to HDI.
 - Continuous migration into HDI is allowed.
- Provides a management API that enables integration with the HCP management user interface (UI) and 3rd-party or homegrown management user interfaces
- Supports Active Directory and LDAP authentication for HCP clients
- Supports HCP tenant and namespace features over CIFS and NFS
- Scales to 400 million files per HDI
- Employs intelligent local cache to accelerate access to HCP content over CIFS and NFS

HITACHI DATA INGESTOR (HDI) TECHNICAL SPECIFICATIONS

Architecture	2.4GHz Intel XEON ^620 processors per node; 12GB RAM per node; 4 gigabit Ethernet (GbE) ports per node; available as a dual-node or single-node appliance, or on VMware vSphere Hypervisor
Capacity (usable)	4TB; scales to 400 million files per HDI
Management	Initial setup wizard; management interface (GUI, SSL); CLI over SSH; simple network management protocol (SNMP); network time protocol (NTP); DFS
Name Services	Dynamic and static mapping between UNIX user and Microsoft Windows user; DNS; Active Directory; NIS
Security and Compliance	<ul style="list-style-type: none"> ■ LDAP v3 ■ POSIX ACLs, NTFS ACLs ■ HTTPS for communication with Hitachi Content Platform (HCP) over WAN ■ Aligns with HCP namespaces to enforce access and policy control at the file system level ■ WORM retention support available as an option, using industry-standard atime-based interface, or through user-settable auto-commit feature that requires no application programming ■ External Syslog server ■ RADIUS support
Backup	NDMP v4; backup and restore via HCP
Protocols Supported	CIFS SMB v1.0, v2.0; NFS v2, v3, v4; FTP
Network	Network link aggregation (IEEE802.3ad); Tagged VLAN (IEEE802.1q); link alteration

- Speeds cloud adoption by eliminating the need to rewrite applications or change user behavior

Summary

IT organizations serving distributed consumers benefit from the cloud on-ramp powered by Hitachi Data Ingestor. HDI coupled with Hitachi Content Platform creates an edge-to-core storage solution that is ideally suited for the challenges of supporting the needs of unstructured data in distributed IT models, such as ROBO and cloud service providers.

With HDI at the edge sending data to the core infrastructure in the data center,

distributed IT environments can reduce their dependence on local storage at the edge. In doing so, they can reduce the cost and complexity of management, backup, archiving, compliance and more.

At core IT data centers, Hitachi Content Platform ensures protection, security, reliable access, preservation, compliance, adherence to policies, and a host of other capabilities that simplify management and automate complex processes.

Combined, Hitachi Data Ingestor and Hitachi Content Platform propel the general concept of object storage forward, help rein in the cost of distributed IT and bring cloud within reach.



Hitachi Data Systems

Corporate Headquarters

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

Regional Contact Information

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

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

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

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

© Hitachi Data Systems Corporation 2012. All Rights Reserved. DS-168-E VA April 2012