

# Hitachi Virtual Storage Platform G1x00 and F1500

80-06-4x or later

---

## Product Overview

This guide provides an overview of the Hitachi Virtual Storage Platform G1x00 (VSP G1x00) and Hitachi Virtual Storage Platform F1500 (VSP F1500) storage systems, including hardware components, general system specifications, software features, management interfaces, user documentation, and software management examples.

© 2018 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

**Notice:** Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

**Notice on Export Controls.** The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

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

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

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

---

# Contents

<b>Preface</b> .....	<b>5</b>
Product version.....	5
Accessing product documentation.....	5
Getting help.....	5
Comments.....	6
<b>Chapter 1: Introducing Hitachi Virtual Storage Platform G1000, G1500, and F1500</b> .....	<b>7</b>
<b>Chapter 2: Hardware features of VSP G1000, G1500, and VSP F1500</b> .....	<b>12</b>
Storage system architecture.....	12
Hardware components.....	13
Controller chassis.....	14
Drive chassis.....	15
Hardware features.....	16
<b>Chapter 3: Software components and features</b> .....	<b>18</b>
Hitachi Storage Virtualization Operating System (SVOS) RF.....	18
Management of SVOS RF systems.....	21
Overview of Storage Advisor.....	22
Hitachi Command Suite.....	23
Command Control Interface.....	26
Advanced global storage virtualization and software bundles.....	27
Hitachi Data Mobility software.....	27
High availability with global-active device.....	28
Nondisruptive migration.....	32
Hitachi Local Replication software.....	35
Hitachi Remote Replication software.....	36
Data-at-rest encryption.....	37
Management software.....	38
Hitachi Data Instance Director.....	38
Hitachi Automation Director.....	40
Hitachi Data Center Analytics.....	42
Hitachi Infrastructure Analytics Advisor.....	42

<b>Chapter 4: Software solution examples.....</b>	<b>44</b>
Enabling simple and efficient storage provisioning and unified management with Command Suite.....	44
Ensuring optimal storage performance and business application service levels with analytics .....	45
Maximizing business application performance and availability with data mobility .....	48
Delivering storage infrastructure as a service through automated workflows....	49
Data protection for business-critical Oracle databases .....	52
End-to-end performance troubleshooting using Infrastructure Analytics Advisor.....	54
Flexible reporting and analysis using Data Center Analytics.....	57
<b>Chapter 5: User documentation for VSP G1000, G1500, and VSP F1500.....</b>	<b>61</b>

---

## Preface

This guide provides an overview of the Hitachi Virtual Storage Platform G1x00 (VSP G1x00) and Hitachi Virtual Storage Platform F1500 (VSP F1500) storage systems, including hardware components, general system specifications, software features, management interfaces, user documentation, and software management examples.

### Product version

This document revision applies to:

- Hitachi Storage Virtualization Operating System: v7.2 or later
- VSP G1000, G1500, and VSP F1500: microcode version 80-06-4x or later
- Hitachi Storage Advisor: v2.2 or later
- Hitachi Command Suite: v8.5.2 or later
- Hitachi Infrastructure Analytics Advisor: v3.1 or later
- Hitachi Data Instance Director: v5.5 or later

### Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

### Getting help

[Hitachi Vantara Support Connect](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to [community.hitachivantara.com](https://community.hitachivantara.com), register, and complete your profile.

## Comments

Please send us your comments on this document to [doc.comments@hitachivantara.com](mailto:doc.comments@hitachivantara.com). Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara Corporation.

**Thank you!**

---

# Chapter 1: Introducing Hitachi Virtual Storage Platform G1000, G1500, and F1500

The Hitachi Virtual Storage Platform G1x00 (VSP G1x00) and Hitachi Virtual Storage Platform F1500 (VSP F1500) unified storage systems provide high performance, high availability, and reliability for always-on, enterprise-class data centers. VSP G1x00 and VSP F1500 deliver guaranteed data availability and feature the industry's most comprehensive suite of local and remote data protection capabilities, including true active-active metro-clustering. The storage systems scale to meet the demands of IT organizations' ever-increasing workloads. When combined with server virtualization, the mission-critical storage virtualization supports a new breed of applications at cloud scale while reducing complexity.

- **High-speed VSDs:** VSP G1500 and VSP F1500 are equipped with new virtual storage directors (VSDs) that use the latest generation of Intel Xenon 2.3-GHz 8-core microprocessor to efficiently manage the front-end and back-end directors, PCI-Express interface, shared memory, and service processor.
- **All-flash VSP F1500:** The VSP F1500 all-flash array is configured exclusively with the latest generation of flash module drives (FMDs) to provide performance optimized for intense I/O operations. Designed for flash-first, high-performance workloads and leveraging Hitachi's SVOS RF-based deduplication and compression, VSP F1500 offers up to five times greater ROI with unified support for SAN, NAS, and mainframe workloads.
  - Accelerated flash architecture delivers consistent, low-latency IOPS at scale.
  - Adaptive flash management distributes writes and rebalances load over time.
  - Hitachi FMDs deliver enterprise performance with superior functionality and greater cost value.

## Key features

Offering the ultimate in enterprise storage technology, VSP G1x00 and VSP F1500 provide the foundation for a software-defined infrastructure that delivers superior storage performance, improved resource utilization, automated, simplified management, and high availability.

- Global storage virtualization enables an always-on infrastructure with enterprise-wide scalability that provides a complete separation between host and storage. The scalability is independent of connectivity, location, storage system, and vendor. Remote data center replication support allows provisioning and management of virtual storage machines up to 100 meters apart.
- Integrated active mirroring enables volume extensibility between systems and across sites through the provisioning and management of active-active volumes up to 100 km apart. Combined with remote data center replication, this mirroring is an ideal solution for critical applications that require zero recovery point and recovery time objectives. Active mirroring is enabled by the global-active device feature.
- Capacity saving functions, including data deduplication and data compression, enable you to reduce the bitcost for stored data by deduplicating and compressing the data stored on internal flash drives. Data deduplication and compression are performed by the controllers of the storage system.
- Accelerated compression enables you to reduce your bitcost for the stored data by utilizing the compression function in the flash module compression (FMC) drives. Accelerated compression allows you to assign FMC capacity to a pool that is larger than the physical capacity of the FMC parity groups. The data access performance of the storage system is maintained when the accelerated compression function is used, as the compression engine is offloaded to the FMC drives.
- Unified storage with enterprise scalability allows you to centrally manage multi-vendor storage resources across all virtualized internal and external storage pools, whether deployed for SAN, NAS, or object storage.
- Centralized storage management software, including Hitachi Storage Advisor and Hitachi Command Suite, simplifies administrative operations and streamlines basic management tasks.
- Hitachi Accelerated Flash storage offers a patented data center-class design and rack-optimized form factor that delivers more than 600 TB per system. It supports a sustained performance of 100,000 8K I/O operations per second per device, with fast and consistent response time.
- Server virtualization integration with leading virtual server platforms provides end-to-end visibility, from an individual virtual machine to the storage logical unit, protecting large-scale multivendor environments.
- Customer-driven nondisruptive migration capability enables movement, copy, and migration of data between storage systems, including non-Hitachi storage systems, without interrupting application access and local and remote copy relationships.



## Storage Virtualization Operating System RF

Hitachi Storage Virtualization Operating System (SVOS) RF abstracts information from storage systems, virtualizes and pools available storage resources, and automates key data management functions such as configuration, mobility, optimization, and protection. This unified virtual environment enables you to maximize the utilization and capabilities of your storage resources while at the same time reducing operations overhead and risk. Standards-compatible for easy integration into IT environments, storage virtualization and management capabilities provide the utmost agility and control, helping you build infrastructures that are continuously available, automated, and agile.

SVOS RF is the latest version of SVOS. Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front-end to slower, less predictable arrays.

SVOS RF provides the foundation for superior storage performance, high availability, and IT efficiency. The enterprise-grade capabilities in SVOS RF include centralized management across storage systems and advanced storage features, such as active-active data centers and online migration between storage systems without user or workload disruption. Features of SVOS RF include:

- Advanced efficiency providing user-selectable data reduction
- External storage virtualization
- Thin provisioning and automated tiering
- Flash performance acceleration
- Deduplication and compression of data stored on internal flash drives
- Storage service-level controls
- Data-at-rest encryption
- Performance instrumentation across multiple storage platforms
- Centralized storage management
  - Simplified: Hitachi Storage Advisor
  - Advanced and powerful: Hitachi Command Suite, Command Control Interface

### Operating system support

The storage systems support IBM® mainframe operating systems as well as open-system operating systems, including VMware®, Windows®, Solaris®, AIX®, Linux®, HP-UX™, XenServer®, and OpenVMS™. For details about operating system support, see [https://support.hds.com/en\\_us/interoperability.html](https://support.hds.com/en_us/interoperability.html).

## Host connectivity

The storage systems support connectivity to mainframe hosts via FICON front-end directors and to open servers via Fibre Channel, iSCSI, and Fibre Channel over Ethernet (FCoE) front-end directors. The storage systems can be configured with a combination of all of these front-end directors to support both mainframe hosts and open servers simultaneously.

## Mainframe compatibility and functionality

Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500 provide full System-Managed Storage (SMS) compatibility and support the following functionalities in the mainframe environment:

- Sequential data striping
- Cache fast write (CFW) and DASD fast write (DFW)
- Enhanced dynamic cache management
- Extended count key data (ECKD) commands
- Multiple Allegiance
- Concurrent Copy (CC)
- Peer-to-Peer Remote Copy (PPRC)
- Extended Remote Copy (XRC)
- FlashCopy<sup>®</sup>
- Parallel Access Volume (PAV)
- Hyper Parallel Access Volume (HPAV)
- Priority I/O queuing
- Red Hat Linux for IBM<sup>®</sup> S/390<sup>®</sup> and zSeries<sup>®</sup>
- SUSE Linux for IBM<sup>®</sup> S/390<sup>®</sup> and zSeries<sup>®</sup>
- zHyperWrite for DB2<sup>®</sup>
- zHPF Extended Distance II (80-05-0x or later)
- FICON<sup>®</sup> Dynamic Routing
- FICON<sup>®</sup> Forward Error Correction (80-05-0x or later)

## Application solutions

Hitachi's portfolio of application solutions for VSP G1000, G1500, and VSP F1500 provides advanced automation, virtualization, scalability and data protection, enabling you to reach the highest application service levels with solutions made for resilience and speed.

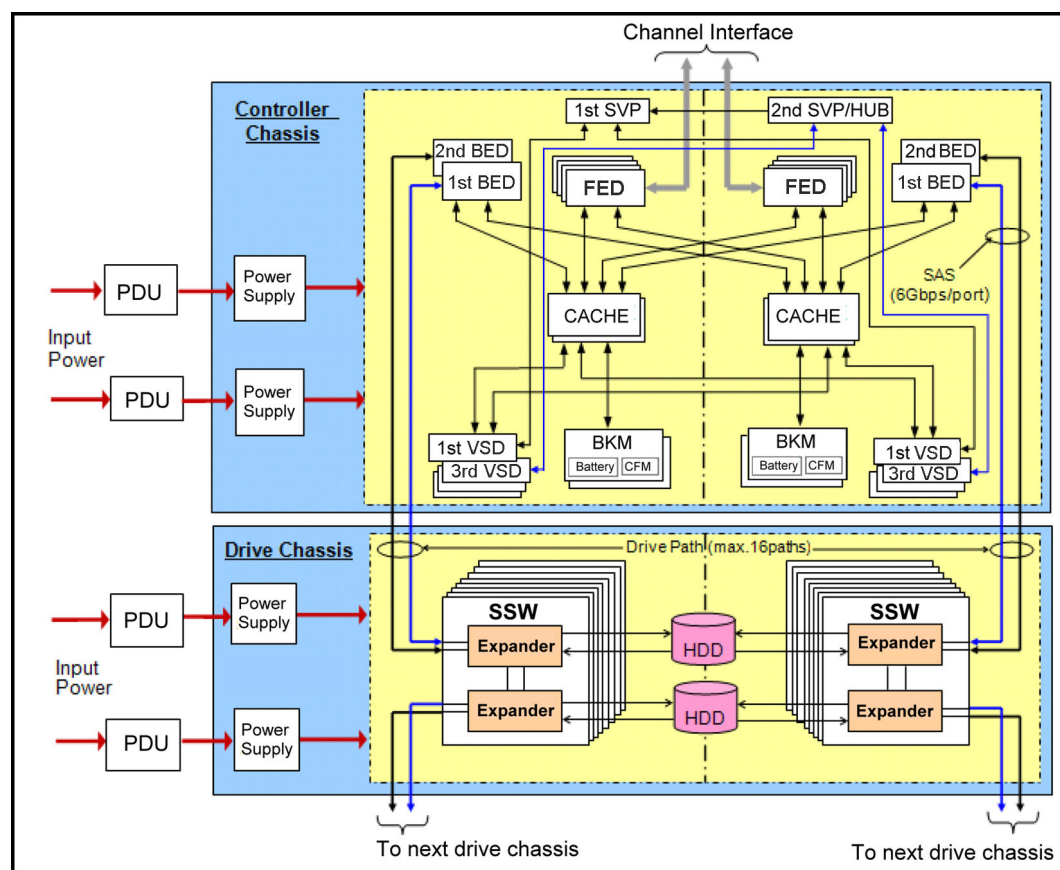
- **Microsoft:** Hitachi solutions for Microsoft® environments enable you to optimize Microsoft® Exchange and SharePoint growth, boost performance and utilization, decrease costs, and improve management and availability.
  - Microsoft Virtual ShadowCopy Service (VSS) for enhanced data protection
  - Microsoft Windows Offloaded Data Transfer (ODX)
  - Hitachi Infrastructure Adapter for Microsoft Systems Center Operations Manager
  - Hitachi Storage Adapter for Microsoft Storage Management Provider
  - Hitachi Storage Adapter for Microsoft System Center Orchestrator
- **Oracle:** Hitachi solutions for Oracle environments enhance consolidation, performance, and efficiency.
  - Database Infrastructure Evaluation Tool (DIET) analyzes your Oracle database environment and provides best practices and expert recommendations for improvement.
  - Hitachi Storage Adapters for Oracle Enterprise Manager, Oracle VM, Oracle Web Center, Oracle Automated Storage Reclamation Utility, and Oracle Database Cloning help you manage your database with less effort and better results.
  - Hitachi Storage Adapter for Oracle Recovery Manager integrates multiple protection services to maximize database availability.
  - Hitachi Dynamic Tiering offers finely tuned performance for Oracle, automatically keeping the most crucial data on the fastest storage.
  - Hitachi Dynamic Provisioning gives your Oracle applications the right amount and right type of storage to maximize performance and efficiency.
- **VMware support:** Supports the Virtual Volume (VVol) functions provided in VMware vSphere® 6.0 for enhanced compatibility in the VMware® environment.
  - Storage Manager for VMware vCenter™
  - vStorage API for Array Integration (VAAI)
  - Unified Storage Provider for VMware (VASA)
  - vStorage API for Multipathing (VAMP)
  - vStorage API for Data Protection (VADP)
  - Hitachi Storage Replication Adapter (SRA)

## Chapter 2: Hardware features of VSP G1000, G1500, and VSP F1500

Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500 are flash-optimized virtual storage platforms that deliver up to 4M IOPS of performance and offer the industry's only 100% data-availability guarantee.

### Storage system architecture

Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500 contain an improved, seventh-generation hierarchical star net architecture that is shown in the following illustration. In this architecture, the virtual storage directors (VSDs) are shared across the cache, front-end directors (FEDs), and back-end directors (BEDs), providing immediate processing power without wait time or interruption to maximize I/O throughput.



The storage systems provide a highly granular upgrade path, allowing the addition of drives to the drive chassis and components such as FEDs to the controller chassis as storage and processing needs increase. A storage system configuration can range from a single controller with no drives (diskless) to a dual controller managing up to 2,304 small form factor (SFF) drives or 1,152 large form factor (LFF) drives (or 576 FMDs) with homogeneous logic control, cache, and front-end and back-end interfaces, all mounted in standard Hitachi Vantara 19-inch racks. You can also use non-Hitachi standard racks that meet the specifications listed in the *Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500*.

## Hardware components


The VSP G1x00 and VSP F1500 storage systems can be configured with one or two controller chassis and up to twelve drive chassis that contain the drives. A controller chassis contains the control logic, processors, memory, and interfaces to the drive chassis and the host servers. A drive chassis contains drives, power supplies, and the interface circuitry that connects it to the controller.

The controller rack contains the controller chassis and up to two drive chassis. Additional racks can contain an intermix of 16U SFF or LFF drive chassis and/or 8U flash drive chassis as well as one or two Hitachi NAS (HNAS) file system servers. When the controllers of a two-controller system are housed in separate racks, the two controller racks can be placed up to 100 meters apart. In addition, the drive racks attached to a controller rack can be placed up to 100 meters from the controller rack.

VSP G1x00 support three types of drives: SFF hard disk drives (HDDs) and solid-state drives (SSDs), LFF HDDs and SSDs, and flash module drives (FMDs), including the second-generation FMD DC2 drives. The drives are mounted in a chassis that is specific to each drive type. VSP G1x00 also support a "diskless" configuration (no drive chassis, all external storage).

The VSP F1500 all-flash array is equipped with advanced high-density FMDs, providing up to 4M IOPS and 40 PB of capacity for multi-workload consolidation.

The following illustration shows a fully configured VSP G1000 with two controllers containing 8 virtual storage director (VSD) pairs (128 CPU cores), 2 TB cache, and 12 16U drive chassis containing 2,304 SFF drives with a physical storage capacity of approximately 2.7 PB. When configured with 6-TB LFF drives, the storage system provides a physical storage capacity of approximately 6.8 PB. When configured with 14-TB FMDs, the storage system provides a physical storage capacity of approximately 8.1 PB.




## Controller chassis

The controller chassis includes the logical components, memory, SAS drive interfaces, and host interfaces. The controller chassis is mounted at the bottom of the rack. If a two-controller system has two service processors (SVPs), both are mounted in controller chassis #0. A controller chassis includes the following maximum number of components:

- Two SVPs, or one SVP and one hub
- 1 TB cache memory
- Two cache backup modules, each comprising a set of two boxes that back up the cache in case of power failure
- Two cache path control adapters (CPA), each comprising a pair of redundant blades that host cache modules as well as providing the PCI-Express path to distribute data among virtual storage director (VSD) pairs, front-end directors and back-end directors. They mirror the data to prevent data loss if one blade fails.
- Four VSD pairs
- Four redundant power supplies with cooling fans
- Six front-end directors (FEDs), each comprising a pair of redundant blades
- Two back-end directors (BEDs), each comprising a pair of redundant blades
- Ten cooling fan assemblies

The following illustration shows the front and rear views of the controller chassis. All components are installed symmetrically on the left and right sides of the controller. The rear side of the controller includes 12 configurable I/O slots. Four of the slots support either FEDs or BEDs, and four of the slots support either FEDs or VSD pairs. Components can be added to these slots as needed to increase performance or storage capacity or to support more host systems. Note that it is not possible to install the maximum number of VSDs, FEDs, and BEDs all at the same time.



## Drive chassis

Hitachi Virtual Storage Platform G1x00 support three types of drives: SFF, LFF, and flash module drives (FMDs). The Hitachi Virtual Storage Platform F1500 all-flash array is configured with FMDs. Each drive type requires its own chassis.

Drive chassis	Description	Drive trays / Drives per tray	Maximum number of chassis / drives per system	
			Single controller (3 racks)	Dual controller (6 racks)
SFF	A 16U group of eight 2U drive trays. Each holds up to 24 vertically positioned 2.5-inch HDD and SSD drives, for a total of 192 SFF drives per chassis.	Eight 2U trays, up to 24 drives each	6 / 1,152 hard drives Up to 192 SSDs	12 / 2,304 drives Up to 384 SSDs

Drive chassis	Description	Drive trays / Drives per tray	Maximum number of chassis / drives per system	
			Single controller (3 racks)	Dual controller (6 racks)
LFF	A 16U group of eight 2U drive trays. Each holds up to 12 horizontally positioned 3.5-inch drives, for a total of 96 LFF drives per chassis.	Eight 2U trays, up to 12 drives each	6 / 576 hard drives Up to 192 flash drives	12 / 1,152 Up to 384 SSDs
FMD	An 8U group of four 2U drive trays. Each holds up to 12 horizontally positioned drives, for a total of 48 FMDs per chassis.	Four 2U trays, up to 12 drives each	6 / 288	12 / 576

## Hardware features

The Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500 offer state-of-the-art advances in hardware technology that improve reliability, serviceability, and access to drives and other components when maintenance is needed.

The key hardware features of the VSP G1000, VSP G1500, and VSP F1500 storage systems include the following:

- **High-speed 8-core CPUs** in the virtual storage directors (VSDs), expanded cache memory (up to 2 TB per 2-controller system), flexible installation, and increased drive types and capacities. The VSP G1500 and VSP F1500 storage systems are equipped with new virtual storage directors (VSDs) that use the latest generation of Intel Xenon 2.3-GHz 8-core microprocessor.
- **Hitachi Accelerated Flash FMD DC2 storage** offers a patented data-center-class design and rack-optimized form factor that delivers more than 2 PB per system, supporting a sustained performance of 100,000 8K I/O per second, per device, with low and consistent response time.
- **Hitachi FMD-HDE** (high-density with encryption) drives provide high-density storage capacities with accelerated compression and hardware-embedded encryption capabilities.
- **The latest 2.5-in. and 3.5-in. 6-Gbps SAS drives** support lower power consumption and higher density per rack with up to 2,304 drives in six 19-inch standard racks.
- **Hitachi NAS Platform hardware accelerated network protocols** support up to 2 Gb/sec throughput for sequential workloads and up to 1.2 million NFS operations per second.



- **Efficient caching** makes up to 2 TB global cache dynamically accessible by all connected hosts and Hitachi NAS Platform nodes.
- **Controller racks can be placed up to 100 meters apart**, providing maximum flexibility to optimize data center space usage and to provide ease of access for operation and maintenance. In addition, the drive racks attached to a controller rack can be placed up to 100 meters from the controller rack. See the detailed description of this feature and the cable diagrams in the *Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500*.
- **Accelerated compression** delivers a data compression capability that enables you to realize more virtual capacity in a parity group than the actual usable capacity.
- **High-temperature mode** is a licensed feature that allows the storage system to run at either standard temperature (60.8°F to 89.6°F / 16°C to 32°C) or at higher temperatures (60.8°F to 104°F / 16°C to 40°C) in a data center, saving energy and cooling costs.
- **External SAN storage virtualization** extends the useful life of existing assets by allowing legacy devices to transparently inherit new functionality from the storage system and SVOS RF. This capability eases migration and reduces costs and risk when moving from legacy Hitachi or non-Hitachi storage systems. Dynamic tiering of storage capacity (internal or external capacity) is supported as well as nondisruptive data migration between tiers and between vendors.

---

## Chapter 3: Software components and features

Hitachi VSP G1000, G1500, and VSP F1500 storage systems are powered by Hitachi Storage Virtualization Operating System RF (SVOS RF) and supported by Hitachi storage management software to enable you to effectively manage and centralize your software-defined infrastructure.

### Hitachi Storage Virtualization Operating System (SVOS) RF

Hitachi Storage Virtualization Operating System (SVOS) RF is the standard operating system for Hitachi VSP G series and VSP F series storage systems. An integrated software system, SVOS RF works with the virtualization capabilities of the storage systems and provides the foundation for global storage virtualization. SVOS RF delivers software-defined storage by abstracting and managing heterogeneous storage to provide a unified virtual storage layer, resource pooling, and automation. SVOS RF also offers self-optimization, automation, centralized management, and increased operational efficiency for improved performance and storage utilization. Optimized for flash storage, SVOS RF provides adaptive inline data reduction and keeps response times fast as data levels grow and automatically recovers services after failure to achieve zero downtime. Selectable services enable data-reduction technologies to be activated based on workload benefit.

SVOS RF provides the following base functionality for Hitachi VSP G series and VSP F series storage systems:

Function	Description
Thin provisioning	Dynamic Provisioning provides thin provisioning for simplified provisioning operations, automatic performance optimization, and storage space savings.
Data reduction	Data reduction functions include pattern detection and removal, accelerated compression provided by Hitachi Accelerated Flash, and selectable controller-based data deduplication and compression.
Global storage virtualization capability	Global storage virtualization enables active-active clustering environment spanning multiple matched Virtual Storage Platform family storage systems (supported externally attached storage).
External storage virtualization	Enables virtualization of external heterogeneous storage using Universal Volume Manager.

Function	Description
Resource partitioning	Resource Partition Manager supports secure administrative partitions for multitenancy requirements.
Cache partitioning	Virtual Partition Manager supports up to 32 cache partitions.
Multipathing and failover	Dynamic Link Manager Advanced provides advanced SAN multipathing with centralized management.
Performance monitoring	Performance Monitor provides an intuitive, graphical interface to assist with performance configuration planning, workload balancing, and analyzing and optimizing storage system performance.
Storage system-based utilities	Storage system-based utilities include LUN manager, customized volume size, Data Retention Utility, cache residency manager feature, quality of service controls, audit log, and volume shredder feature.
Standard management interface support	Management interface support includes SMI-S provider, SNMP agent, and REST.
Optimized storage for virtualized server infrastructure	A wide range of plugins and adapters are available to enhance virtual server infrastructure performance and administrator productivity. SVOS RF features integration of VMware applications (including VAAI, VASA, VAMP, VADP, SRA, VVOL) and Microsoft Windows applications (including VSS, ODX).
Mainframe compatibility	Mainframe compatibility features include IBM® PAV and HyperPAV, dynamic volume expansion (DVE), extended address volumes (EAV), Peer-to-Peer Remote Copy (PPRC), Extended Remote Copy (XRC), high-performance FICON® with multi-track, FICON® forward error correction, basic and IBM® GDPS HyperSwap®, FlashCopy® and space-efficient FlashCopy® (FCSE), zHyperWrite, and zHPF Extended Distance II.

### SVOS RF data reduction


Hitachi Storage Virtualization Operating System RF is designed to deliver superior adaptive data reduction and operational efficiency. To improve return on investment and allow greater VM consolidation, SVOS RF adaptive data reduction intelligence is optimized for highest system throughput and response time consistency. With multithreaded capabilities and quality of service (QoS) control, SVOS RF adaptive intelligence has the ability to slow down or pause the data reduction processing. It takes this action if the system reaches high processor utilization level and/or if elongated wait time is experienced for data on AFA cache that is ready to be written. SVOS RF also manages the data reduction services in use based on configuration. If flash modules are detected, FMD compression is used. If encryption is required, SVOS RF compression is used.

Data reduction capabilities include:

- Pattern detection and removal

Pattern detection identifies pre-defined repetitive binary patterns, including zeros, prior to compressing and identifying duplicates. This process reduces the volume of data to be processed by the compression and deduplication engine.

- Accelerated compression



The accelerated compression feature of Dynamic Provisioning delivers a data compression capability that enables you to realize more virtual capacity in a parity group than the actual usable capacity, providing improved storage optimization. You can enable accelerated compression at the parity-group level on HAF flash module drives (FMD DC2, FMD HD). When accelerated compression is enabled, the capacity of a parity group can be expanded up to several times. LDEVs created from an expanded-capacity parity group are used as Dynamic Provisioning pool volumes to create or expand a pool, and the data on these LDEVs is compressed before it is stored on the drives.

Implementation of accelerated compression requires careful planning, detailed calculations, and monitoring to verify the desired results. When accelerated compression is in use, both the used pool capacity and the used pool capacity reserved for writing must be monitored. Threshold values are set so that SIMs are reported when threshold values are exceeded, enabling you to expand the pool capacity or delete unwanted data before an error condition occurs (for example, pool full). For details about implementing accelerated compression, see the *Provisioning Guide for Open Systems*.

- Capacity saving

The capacity saving function includes data deduplication and compression. When the capacity saving function is in use, the controller of the storage system performs deduplication and compression to reduce the size of data to be stored, thereby reducing your bitcost for the stored data. Capacity saving can be enabled on DP-VOLs in Dynamic Provisioning pools. You can use the capacity saving function on internal flash drives only, including data stored on encrypted flash drives.

- Deduplication

The data deduplication function deletes duplicate copies of data written to different addresses in the same pool and maintains only a single copy of the data at one address. The deduplication function is enabled on a Dynamic Provisioning pool and then on the desired DP-VOLs in the pool. When deduplication is enabled, data that has multiple copies between DP-VOLs assigned to that pool is removed.

When you enable deduplication on a pool, the deduplication system data volume (DSD volume) for that pool is created. The DSD volume is used exclusively by the storage system to manage the deduplication function. A search table in the DSD volume is used to locate redundant data in the pool.

- Compression

The data compression function utilizes the LZ4 compression algorithm to compress the data. The compression function is also enabled per DP-VOL.

## Management of SVOS RF systems

SVOS RF includes the following software products to manage SVOS RF systems:

- Hitachi Storage Advisor (HSA)

To simplify operations, SVOS RF systems are managed by Hitachi Storage Advisor, a wizard-driven management application that lets IT staff configure resources in minutes and monitors the health of SVOS RF-managed storage at a glance.

- Hitachi Command Suite (HCS)

Hitachi Command Suite provides single-point management for all Hitachi physical and virtualized storage and is the interface for integration with other Command Suite software.

- Command Control Interface (CCI)

For more complex storage environments, CCI provides powerful command-line control and advanced functionality for Hitachi VSP G series and F series.

- Hitachi Storage Advisor Embedded (HSAE)

Hitachi Storage Advisor Embedded enables configuration and monitoring of a single SVOS RF system. Storage management is made easy with a simple and intuitive graphical user interface, high-performing REST API, and CLI embedded natively in the storage system.

## Overview of Storage Advisor

Hitachi Storage Advisor is a unified software management tool that reduces the complexity of managing storage systems by simplifying the setup, management, and maintenance of storage resources.

Storage Advisor reduces infrastructure management complexities and enables a new simplified approach to managing storage infrastructures. It provides intuitive graphical user interfaces and recommended configuration practices to streamline system configurations and storage management operations. You can leverage Storage Advisor to easily provision new storage capacity for business applications without requiring in-depth knowledge of the underlying infrastructure resource details. It provides centralized management while reducing the number of steps to configure, optimize, and deploy new infrastructure resources.



Some of the key Storage Advisor capabilities include:

- Simplified user experience for managing infrastructure resources. Visual aids enable easy viewing and interpretation of key management information, such as used and available capacity, and guide features to help quickly determine appropriate next steps for a given management task.
- Recommended system configurations to speed initial storage system setup and accelerate new infrastructure resource deployments.
- Integrated configuration workflows with Hitachi recommended practices to streamline storage provisioning and data protection tasks.
- Common, centralized management for supported storage systems.
- A REST-based API to provide full management programmability and control in addition to unified file-based management support.
- Storage Advisor enables automated SAN zoning during volume attach and detach. Optional auto-zoning eliminates the need for repetitive zoning tasks to be performed on the switch.

## Hitachi Command Suite

Hitachi Command Suite (HCS) is an application-centric storage management solution that simplifies administration of a common pool of multivendor storage. The software offers comprehensive management, control, and discovery for file, object, and block storage services, reducing complexity, costs, and risk in the storage infrastructure.

The base HCS product consists of Hitachi Device Manager, which provides centralized management of multiple Hitachi storage systems. By providing a single console for managing complex storage environments, Device Manager software unifies and simplifies storage management. Featuring an intuitive GUI, Device Manager supports multiple management views for primary and secondary storage, including physical, logical, host, and NAS and virtual server for provisioning and storage pooling.



**Note:** Key functions of Storage Navigator have been integrated into HCS to enable a unified interface for storage management.

HCS comprises the following optional components, each of which is licensed separately:

- **Hitachi Tiered Storage Manager:** Supports storage tiers of differing performance characteristics so that volume data storage costs and performance can be optimized.
- **Hitachi Replication Manager:** Adds remote replication capabilities and supports backup and disaster recovery.
- **Hitachi Tuning Manager:** Supports optimizing the performance of storage resources.
- **Hitachi Compute Systems Manager:** Supports centralized monitoring and management of hosts, including rebooting and power management.
- **Hitachi Dynamic Link Manager:** Supports the use of multiple paths between resources such as hosts and storage for path failover and load balancing.
- **Hitachi Global Link Manager:** Supports management of multipath management software between resources, such as hosts and storage.
- **Hitachi Automation Director:** Provides tools to automate and simplify end-to-end processes, such as storage provisioning, for storage and data center administrators.

At minimum, you must license Device Manager. Additional licensing can be added as needed for other storage management products. Related functionality becomes available in the HCS user interface in the form of activated menu choices and new or updated tabs and related screens and buttons.

### *Advanced SAN multipathing*

Hitachi Dynamic Link Manager offers robust multipath SAN connections between servers and storage systems. It provides fault-tolerant failover, failback, load balancing, and centralized path management, for improved information access, usability, and availability. Automatic workload balancing helps to maintain outstanding system performance across all available paths. If one path fails, Dynamic Link Manager automatically switches the I/O to an alternate path, ensuring that an active route to data is always available.

Dynamic Link Manager offers the following benefits:

**Business continuity**

- Improves system performance by spreading I/O request workload across available paths to ensure that no single path is overworked or underutilized
- Provides a high level of data availability through automatic path failover and failback, ensuring continuous access to application data, improved application performance, and reduced risk of financial loss due to failures of critical applications
- Improves availability and data access on storage systems in SAN environments, with path failover and I/O balancing over multiple HBAs
- With its health-check facility, monitors online path status at specified intervals, and places a failed path offline when an error is detected

**Productivity and process**

- Provides a centralized facility for managing path failover, automatic failback, and selection of I/O balancing techniques through integration with Hitachi Global Link Manager
- Eases installation and use through the auto-discovery function, which automatically detects all available paths for failover and load balancing
- Provides one path-management tool for all your operating systems
- Includes a command line interface (CLI) that allows administrators the most flexibility in managing paths across the network
- Provides manual and automatic failover and failback support

***Data replication***

Hitachi Replication Manager provides management capabilities to configure, manage, and monitor Hitachi replication products for local and remote sites. Replication Manager supports open systems and mainframe environments and provides support for multiple data centers and multiple storage systems at each data center. It simplifies and optimizes configuration, operation, task management, automation, and monitoring of the critical applications and storage components of your replication infrastructure. The following figure shows the Replication Manager interface.





Replication Manager offers the following benefits:

### Centralized management of a replication environment

Replication Manager can be used to manage storage systems and hosts at different sites. The status of copy pairs, the progress of copy operations, and performance information (such as data transfer delays between copy pairs and buffer usage when copying volumes) can be centrally managed from a single console.

### Integrated database backup management

Replication Manager supports creating backups of databases. Called application replicas, these backups are managed as a series of secondary volumes that are rotated on a scheduled basis. Replication Manager manages the relationships between backup objects and their associated logical units within storage devices, the relationships between primary and secondary volumes, and the backup history. Replicas can be mounted and dumped to tape using scripts executed through Replication Manager.

### Visual representation of replication structures

Replication Manager provides a centralized workspace where you can visually check the structure of copy pairs configured across multiple storage systems. Host and storage system relationships and copy pair definitions can be visualized using functional views. Copy pairs in complex configurations such as multitarget configurations and cascade configurations can be viewed as lists.

**Monitoring and immediate notification of error information**

Replication Manager provides capabilities to specify monitoring conditions for designated copy pairs and sidefiles. Alerts can be automatically generated when the conditions are satisfied. You can continue monitoring the system even when not logged in to Replication Manager because alerts can be reported in the form of email messages or SNMP traps. The status of application replicas is tracked and reflected in summary form so that you know to what extent the application databases are protected. These monitoring features allow you to work out advance strategies to handle potential problems such as the deterioration of transfer performance due to insufficient network capacity or blocked pairs caused by buffer overflows.

**Modification of replication structures**

Replication Manager provides capabilities to configure additional copy pairs as business operations expand and improve performance by expanding buffer capacity for copying volumes. You can also change pair states manually after error recovery. Using the wizards provided in the GUI, you can set up pairs while visually keeping track of complex replication structures.

**Monitoring and analyzing remote copy performance (write delay time)**

When using Universal Replicator, you can check copy performance visually and perform root cause analysis using the Replication tab of the Hitachi Command Suite GUI.

**Command Control Interface**

Command Control Interface (CCI) CLI software provides powerful command-line control for Hitachi Virtual Storage Platform family storage systems, enabling you to perform storage system configuration and data management operations by issuing commands to the storage systems.

CCI provides command-line control and advanced functionality for local and remote replication operations, including ShadowImage, Thin Image, TrueCopy, Universal Replicator, and global-active device. CCI commands can be used interactively or in scripts to automate and standardize storage administration functions, thereby simplifying the job of the storage administrator and reducing administration costs. For remote replication operations, CCI interfaces with the system software and high-availability (HA) software on the host as well as the software on the storage system. CCI provides failover operation commands that support mutual hot standby in conjunction with industry-standard failover products. Using CCI scripting, you can set up and execute a large number of commands in a short period of time while integrating host-based high-availability control over copy operations.

For VSP G series and VSP F series, CCI provides command-line access to the same provisioning operations that are available in Hitachi Device Manager - Storage Navigator. Because some provisioning operations can take time to process, CCI provides two ways to execute the configuration setting command: synchronously or asynchronously. Asynchronous command processing is used for operations that take time to process on the storage system. Once an asynchronous command has been issued, you can execute additional commands without having to wait for the asynchronous command to complete, and you can also monitor the completion status of asynchronous commands.

## Advanced global storage virtualization and software bundles

Optional SVOS RF features include best-in-class local and remote replication technologies as well as active-active metro clustering to provide rapid recovery from system and site-level events that could disrupt access to data. SVOS RF business continuity solutions are designed for maximum flexibility, enabling organizations to build a recovery strategy that spans multiple data centers and delivers to their specific SLAs.

Optional software products and packages for SVOS RF systems include:

- Hitachi Data Mobility package increases storage performance and lowers costs with automated data placement.
- Global-active device feature license enables active-active storage clusters that span data centers for business continuity and superior data sharing.
- Nondisruptive migration delivers large-scale migration capabilities that require less time and effort to execute and deliver continuous operations while ensuring application quality of service and maintaining data protection.
- Hitachi Local Replication package quickly creates space-efficient, point-in-time snapshots, eliminating the need for a traditional backup window and enabling fast recovery.
- Hitachi Remote Replication package includes synchronous and asynchronous replication providing zero RPO and near-zero RTO capabilities across three or even four geographically dispersed locations.
- Data-at-Rest Encryption software protects data at rest on internal storage media for enhanced data privacy and compliance.

### Hitachi Data Mobility software

By simplifying tiered storage management, Hitachi Data Mobility software delivers the highest storage performance for the most frequently accessed data while at the same time lowering costs by automatically optimizing data placement.

Hitachi Data Mobility software provides complete data movement capabilities. It combines two leading data mobility technologies with Hitachi Dynamic Tiering and Hitachi Tiered Storage Manager software. The combination enables intelligent placement of data within virtualized Hitachi storage environments while optimizing business application service levels.

- Hitachi Dynamic Tiering automates data placement and access in a tiered storage environment. It dynamically moves the most active data to the highest-performing storage tiers while moving less frequently accessed data to lower tiers. Hitachi Dynamic Tiering active-flash mode moves suddenly active data via synchronous promotion to higher-performing tiers in real time. In seconds to subseconds, active flash responds to workload demands based on current I/O activity. Active flash proactively preserves flash endurance by monitoring and demoting pages that exceed thresholds for heavy write I/O.
- Hitachi Tiered Storage Manager enables administrators to proactively match business application price, performance, and availability characteristics to storage resource attributes. Administrators can proactively create and pool different storage classes to maximize operational and cost efficiency and easily align them to specific business application needs. As storage service levels change over time, Tiered Storage Manager facilitates nondisruptive data migration between storage tiers and externally virtualized storage resources to match new application requirements. Through custom data management policies, Tiered Storage Manager helps you to properly monitor and control the automated and active behavior of Hitachi Dynamic Tiering.

## **High availability with global-active device**

Global-active device (GAD) uses volume replication to provide a high-availability environment for hosts across storage systems and sites. Global-active device provides data protection and minimizes data-access disruptions for host applications due to site or storage system failures, ensuring continuous, simplified operations in distributed environments. Efficient and scalable active-active design gives you continuous application availability for both traditional and cloud storage. Active-active stretched clusters over local and metro distances allow application access to replicated data from the shortest path, for the highest performance. Global-active device works seamlessly with other advanced capabilities of SVOS RF to simplify and improve disaster recovery operations and dramatically reduce return-to-operations time, enabling customers to meet strict service-level agreements for zero or near-zero recovery point objective (RPO) and recovery time objective (RTO).

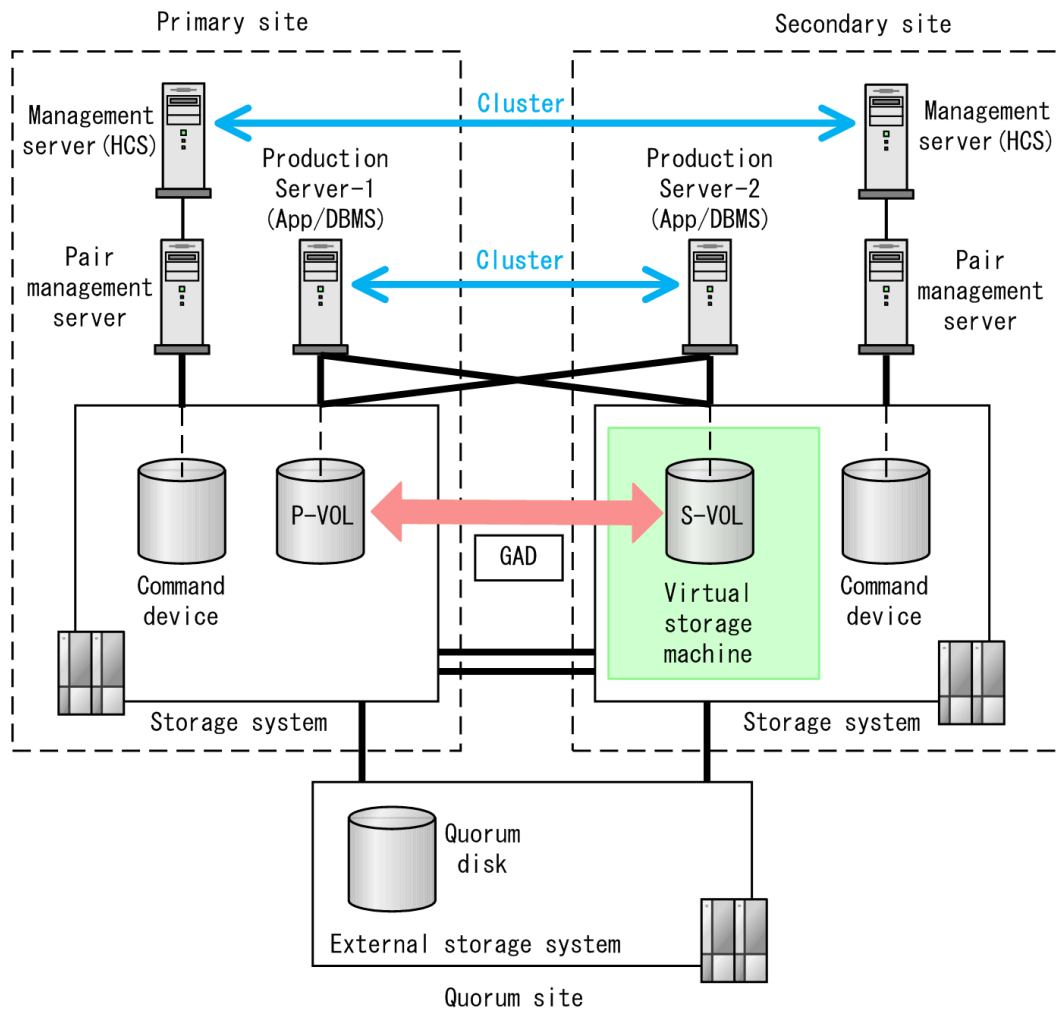
Establishing a global-active device pair has the following benefits:

- Continuous I/O  
If a primary volume becomes unavailable, the host continues to transparently access the secondary volume.
- Clustered failover  
You do not need to perform storage system tasks such as suspension or resynchronization of a global-active device pair due to a host failure.
- Virtual machine integration  
If a virtual machine is creating a high load at one site, you can move the load to the other site, eliminating the need for data migration.

### **How global-active device works**

A GAD pair consists of a primary data volume and a synchronous, remote copy on Hitachi VSP G series storage systems at the primary and secondary sites. A virtual storage machine is set up in the secondary VSP G series storage system using the physical information from the primary system. The GAD primary and secondary volumes are assigned the same virtual LDEV number in the virtual storage machine. As a result, the host treats the paired volumes as a single volume on a single storage system, with both volumes receiving the same data from the host.

The following figure shows an example GAD configuration.

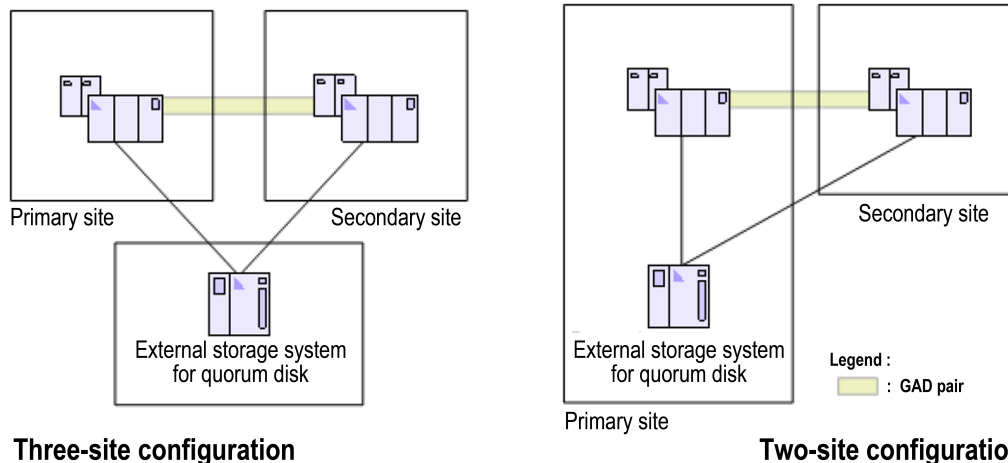


GAD pair volumes are monitored by a quorum disk (preferably located at a third site). The quorum disk acts as a heartbeat for the GAD pair, with the primary and secondary storage systems accessing the quorum disk periodically to check on the other storage system. In the event of a communication or hardware failure, the quorum disk determines which storage system is still accessible, allowing operations to continue without interruption.

The SAN multipathing software on the host runs in an active-active configuration. If the primary volume (P-VOL) or secondary volume (S-VOL) cannot be accessed, host I/O is automatically redirected to an alternative path. Native multipath software operates at campus distances using cross-site paths (as shown in the previous diagram). At metro distances, Hitachi Dynamic Link Manager (HDLM) offers increased performance using preferred paths (shortest possible route).

### Global-active device storage system configurations

Global-active device requires three storage systems: primary, secondary, and an external system used for the quorum disk. The configuration can be set up across one, two, or three sites.



- In a three-site configuration (recommended), each storage system is located at a separate site. This configuration provides maximum protection against system failures and site failures.
- In a two-site configuration, both the primary storage system and the quorum storage system are located at the primary site. This configuration provides a moderate level of protection against system and site failures.
- In a one-site configuration (not shown), all storage systems are located at the same site. This configuration protects against system failures but not site-wide failures.

For details about GAD configurations, requirements, and setup, see the following documentation:

- *Global-Active Device User Guide*
- *Hitachi Command Suite User Guide*
- Hitachi Command Suite Dynamic Link Manager documentation

### Combining global-active device and Universal Replicator

In a GAD system, the server accesses the primary site and the secondary site simultaneously and shares the same data between the two sites (at campus distance). If a failure occurs at one site, you can continue operations at the other site. However, if a failure occurs at both sites, for example due to a large-scale disaster, you cannot continue operations with the data redundancy provided by only global-active device.

To manage this situation, you can implement a 3-data-center (3DC) configuration by combining GAD and Universal Replicator (UR). This is called a GAD 3DC delta resync (GAD+UR) configuration. If a failure occurs at both the primary site and the GAD secondary site, the GAD+UR configuration enables you to continue operations using the UR secondary site (at metro distance).

For more information about GAD 3DC delta resync operations, see the following documents:

- *Global-Active Device User Guide*
- *Hitachi Universal Replicator User Guide*

- *Hitachi Command Suite User Guide*
- *Hitachi Command Suite Replication Manager User Guide*

### Setting preferred and nonpreferred paths using Asymmetric Logical Unit Access

When the paths connecting a server and a storage system in a GAD configuration contain a short-distance straight path and a long-distance cross path, I/O performance varies depending on the path. Using Asymmetric Logical Unit Access (ALUA), you can set the short-distance straight path as the preferred I/O path and the inefficient long-distance cross path as the nonpreferred path to improve overall system performance.

To use ALUA to set the preferred and nonpreferred paths for GAD pairs in a cross-path configuration, you first enable the ALUA mode on the storage system, which sets all paths as preferred paths, and then you set the asymmetric access status of the cross path as a nonpreferred path. For details and instructions, see the *Global-Active Device User Guide*.

## Nondisruptive migration

One of the biggest challenges during technology refresh cycles is to eliminate downtime and service disruption when the data used by the host is copied to a new volume on the new storage system and the host is reconfigured to access the new volume.

Nondisruptive migration makes it possible to relocate data from existing storage systems to newer storage systems without interrupting access by hosts. Data migration is accomplished using the global storage virtualization technology of the target storage systems. Resources on the source storage system are virtualized on the target storage system. From the perspective of the host, I/O requests continue to be serviced by the source storage system during the migration process.

The following storage system combinations are supported:

Source	Target
Hitachi Universal Storage Platform V/VM	VSP G1000, VSP G1500, and VSP F1500
Hitachi Universal Storage Platform V/VM	VSP G200, G400, G600, G800 VSP G350, G370, G700, G900
Hitachi Virtual Storage Platform	VSP G1000, VSP G1500, and VSP F1500
Hitachi Virtual Storage Platform	VSP G200, G400, G600, G800 VSP G350, G370, G700, G900
Hitachi Unified Storage VM	VSP G1000, VSP G1500, and VSP F1500
Hitachi Unified Storage VM	VSP G200, G400, G600, G800 VSP G350, G370, G700, G900



Nondisruptive migration offers these benefits:

- Data is migrated between storage systems without interrupting host applications.
- You can maintain data replication throughout the migration process by allowing the target storage system to inherit pair configurations before migrating the actual data.
- You can reduce the overall migration effort by importing configuration definition files instead of having to reconfigure pairs on the target storage system.
- The migration process is designed to be carried out in stages to reduce demands on network bandwidth.
- You can easily monitor migration project and migration job progress and status by reviewing both numerical and graphical data, which includes estimated information about how long the migration is likely to take.
- Up to seven source storage systems can be consolidated into a single target storage system.



**Note:** Because good planning is essential to smooth migration, we strongly recommend the nondisruptive migration planning service offered by Hitachi Vantara Global Solution Services (GSS).

### How nondisruptive migration works

The following workflow summarizes the stages of the migration process.

#### 1. Virtualization of source volumes

- A virtual storage machine is created in the target storage system, a representation of the source storage system that behaves exactly like its physical counterpart (with the same name and serial number).
- The source volume is mapped within the virtual storage machine as a virtual device (with the same LDEV ID as the source volume). This is known as the *target volume*.

#### 2. Switching of host I/O

The HCS nondisruptive migration workflow prompts you to perform the following operations manually:

- Initiate I/O between the target storage system and the host.
- Disable I/O between the source storage system and the host.

You must do this using path management software (such as Dynamic Link Manager), OS native multipath functions, or by changing the zoning configuration. When you confirm that the switch was successful, the I/O path is changed.

Initially, read and write requests continue to be processed by the source storage system. This is known as cache through mode, and is in effect while the volume on the source storage system remains connected to the host.

### 3. Unallocation of source volumes

To prevent the host from accessing the source volume through the source storage system, the HCS nondisruptive migration workflow reminds you to delete the LUN path between the source volumes and the host before continuing.

When you disable the connection between the host and the volume on the source storage system, the cache is switched to write sync mode. Thereafter, all read and write requests are handled by the target storage system, and data is written to both the source and target volumes.



**Note:** In write sync mode, when the host issues a read request while data is being written to the source storage system, the read processing does not start until the write is complete.

### 3. Re-creation of existing ShadowImage copy pairs

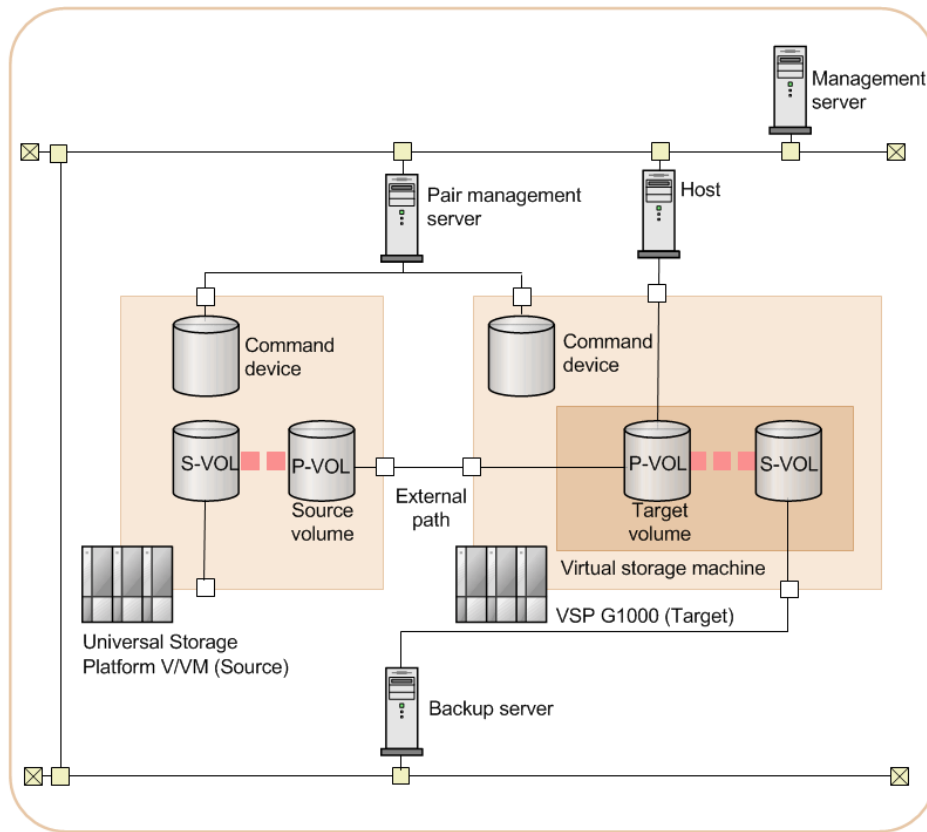
If you plan to migrate secondary volumes, the HCS nondisruptive migration workflow leads you through the process of re-creating the source secondary volumes on the target storage system.

### 4. Data migration

In this stage, the data is copied to its final destination on the target storage system.

#### Sample nondisruptive migration configuration

The following figure shows a nondisruptive migration configuration with secondary volumes and multiple servers. The term *backup server* is used because this server is responsible for running the scripts that copy the data from the primary to the secondary volumes.



Legend: ■ ■ ■ ■ ShadowImage

For a complete description of the nondisruptive migration feature, including requirements and setup, see the *Nondisruptive Migration User Guide* and the *Hitachi Command Suite User Guide*.

## Hitachi Local Replication software

Hitachi Local Replication software combines Hitachi ShadowImage<sup>®</sup>, Hitachi Thin Image, and Hitachi Replication Manager to deliver convenient and cost-effective full-volume data cloning for fast, point-in-time data copies. Hitachi Local replication ensures rapid

restart-and-recovery times by combining local mirroring of full volumes with fast, space-efficient snapshots.

- High-speed, nondisruptive local mirroring technology of Hitachi ShadowImage® replication software rapidly creates multiple copies of mission-critical information within all Hitachi storage systems in mainframe and open-systems environments. ShadowImage software keeps data RAID-protected and fully recoverable, without affecting service or performance levels. Replicated data volumes can then be split from the host applications and used for system backups, application testing and data mining applications, while business continues to run at full capacity.
- The high-speed, nondisruptive snapshot technology of Hitachi Thin Image snapshot software rapidly creates up to one million point-in-time copies of mission-critical information within any Hitachi storage system or virtualized storage pool, without impacting host service or performance levels. Because snapshots store only the changed data, the storage capacity required for each snapshot copy is substantially smaller than the source volume. As a result, Thin Image can provide significant savings over full volume cloning methods. Thin Image snapshot copies are fully read/write compatible with other hosts and can be used for system backups, application testing, and data mining applications while the business continues to run at full capacity.
- Part of Hitachi Command Suite, Replication Manager software configures, monitors, and manages Hitachi local and remote replication products for open systems and mainframe environments.

Application-consistent ShadowImage clones and Thin Image snapshots can be orchestrated using Hitachi Data Instance Director (HDID) software. HDID supports Microsoft® Exchange and SQL Server® as well as Oracle databases on Linux operating systems. These clones and snapshots can be easily created as part of a complete data protection workflow, using HDID's unique whiteboard-like interface. HDID can also trigger a ShadowImage clone or Thin Image snapshot on the remote side of a distance replication pair.

Hitachi Vantara Global Services Solutions provides Implementation Services for Hitachi ShadowImage® and Hitachi Thin Image software. These services help organizations improve testing and application deployment operations with high-speed, problem-free data duplication. Consultants tailor the configuration and integration of the local replication software to serve an organization's backup and recovery application requirements.

## Hitachi Remote Replication software

Hitachi Remote Replication software combines Hitachi TrueCopy®, Universal Replicator, and Hitachi Replication Manager solutions to enable remote data protection at up to four data centers. Providing continuous, nondisruptive, host-independent data replication, Hitachi Remote Replication software ensures the highest levels of data integrity for local or metropolitan areas. Copies generated by Hitachi Remote Replication software products can be used for the rapid recovery or restart of production systems on primary or secondary (disaster recovery) systems following an outage. They can also be

used for nondisruptive test and development, data warehousing, data mining, data backup, or data migration applications.

- Hitachi TrueCopy<sup>®</sup> enables synchronous remote replication of mission-critical data from a primary data center to a secondary data center at distances up to 300 km. TrueCopy delivers immediate zero-RPO and automated failover capabilities and is compatible with open-systems and mainframe environments.
- Universal Replicator features journal disk caching for achieving tight RPO time capabilities, even in the event of a network outage. Universal Replicator provides asynchronous remote copy, over any distance, for Hitachi VSP G series and VSP F series storage and is compatible with open-systems and mainframe environments. Deployed implementations can be configured with or without delta resync, which ensures replication consistency for the highest level of remote copy data integrity at any distance.
- Part of Hitachi Command Suite, Replication Manager software configures, monitors, and manages Hitachi local and remote replication products for open systems and mainframe environments.

TrueCopy and Universal Replicator can also be automated as part of an end-to-end, unified data protection, retention, and recovery management solution within Hitachi Data Instance Director (HDID) software. HDID can also automatically trigger Thin Image snapshots and ShadowImage clones from the remote copy of the data.

From remote copy planning to advanced implementation services, Hitachi Vantara Global Services Solutions can support the successful and timely deployment of the most resilient data protection infrastructures. Services to support TrueCopy and Universal Replicator software and other business continuity and disaster recovery solutions from Hitachi Vantara are available for both open-systems and mainframe environments.

## **Data-at-rest encryption**

Hitachi VSP G series and VSP F series storage systems provide a performance-friendly AES-256-XTS encryption capability on the back-end I/O module. This capability protects data at rest on internal storage media and volumes attached to those directors. When data is encrypted, information leakage can be prevented when replacing the storage system or the drives in the storage system. Similarly, the encryption capability provides an extra measure of protection and confidentiality for lost, stolen, or misplaced media that may contain sensitive information.

The data-at-rest encryption feature has two components: the encrypting back-end director hardware component and the Encryption software license. Encryption can be applied to some or all of the internal drives with no throughput or latency impacts for data I/O and little or no disruption to existing applications and infrastructure. Data-at-rest encryption includes integrated key management functionality that is both simple and safe to use, providing a unique encryption key for each individual piece of media internal to the array.

Data-at-rest encryption is configured and monitored through the Hitachi Command Suite and Device Manager - Storage Navigator management software, providing role-based access control (RBAC) for the separation of duties including enabling/disabling encryption as well as archiving encryption keys.

## Management software

The Hitachi approach to software-defined solutions enables you to effectively manage your IT infrastructure to align storage resources to rapidly changing business demands, achieve superior returns on infrastructure investments, and minimize operational costs. Hitachi's suite of management software delivers higher storage availability, mobility, and optimization for key business applications, automating storage management operations with integrated best practices to accelerate new resource deployments. Using Hitachi's storage management software, administrators are able to manage more storage capacity with less effort and ensure service levels for business-critical applications are met while increasing utilization and performance of virtualized storage assets.

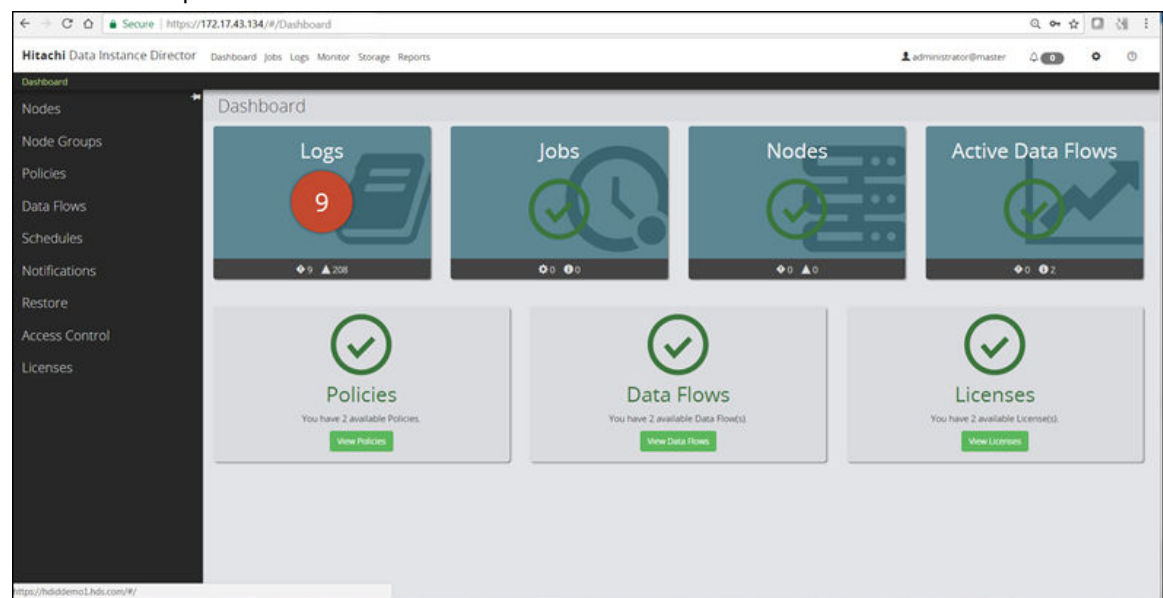
Management software for Hitachi VSP G series and VSP F series includes:

- Hitachi Data Instance Director (HDID)
- Hitachi Data Center Analytics (HDCA)
- Hitachi Automation Director (HAD)
- Hitachi Infrastructure Analytics Advisor (HIAA)

## Hitachi Data Instance Director

The enterprise copy data management platform enabled by HDID provides business-defined data protection, which simplifies the creation and management of complex, business-defined policies to meet service-level objectives for availability, recoverability, and retention.

HDID supports HUS VM, Hitachi VSP, VSP Fx00 models, and VSP Gx00 models, offering an orchestration layer for remote replication supporting global-active device, Hitachi TrueCopy® and Universal Replicator, local and remote snapshots and clones with Hitachi Thin Image and Hitachi ShadowImage®, continuous data protection, and incremental-forever backup.



HDID provides the following benefits:

## Operational recovery

HDID offers multiple approaches to meeting operational recovery requirements, allowing business service-level objectives for recovery to be met at optimal cost for differing criticality of data.

- *Storage replication-based operational recovery*

HDID configures, automates, and orchestrates local application-consistent snapshot and clone copies using the local replication capabilities of Hitachi Virtual Storage Platform (VSP) family, Hitachi Unified Storage VM (HUS VM), and NAS Platform (HNAS).

This integration provides the ability to create fast, frequent copies of production data, with no impact on the performance of the production system. Very aggressive recovery point objectives (RPO) can be easily achieved for Microsoft Windows platforms for Microsoft Exchange and Microsoft SQL Server, for Oracle database environments on Linux, IBM® AIX®, and Solaris, and for SAP HANA environments on Linux. HDID also provides storage-based protection of VMware vSphere environments natively for Hitachi Hitachi block storage systems and via Hitachi Virtual Infrastructure Integrator for Hitachi NAS Platform. Both types of vSphere integration allow vSphere administrators to apply protection policies without leaving the vSphere management interfaces. Other applications can also be integrated using the simple scripting interface.

Storage data snapshots and clones can be mounted and unmounted automatically as part of an HDID policy workflow. They can facilitate access to a current copy of production data for testing and development purposes, or back up to a target device such as a purpose-built backup appliance (PBBA) or tape.

- *Host-based operational recovery*

HDID includes several storage-agnostic technologies for protection of application and file system data. Continuous data protection (CDP) and live backup support Windows environments, with application-specific support for Microsoft Exchange and SQL Server. Batch mode backup is supported on Windows, Linux, IBM® AIX®, and Oracle Solaris systems.

## Disaster recovery

HDID provides choices for restoring operations at, or from, another location following a site level outage.

- *Storage-based disaster recovery*

HDID configures and automates global-active device active-active storage cluster, Hitachi TrueCopy<sup>®</sup> synchronous remote replication and Universal Replicator on block-based systems, and file replication on HNAS, to provide a copy of data in another location. HDID can also orchestrate application-aware snapshots of these remote replicas.

## Long-term retention

The governance copy services allow you to back up file data to Hitachi Content Platform (HCP) for Windows, Linux, IBM<sup>®</sup> AIX<sup>®</sup>, and Oracle Solaris systems. Unlike other data protection products, HDID places data in its original format; meaning that it can be read without HDID, which allows you to support corporate and regulatory data retention requirements. Because the data is readable, it is indexable with tools such as Hitachi Content Intelligence and can be used for analytics with tools such as Pentaho Data Integration.

## Unified management

One of the many benefits of HDID is its single-footprint platform. It enables you to layer, combine, and orchestrate backup, CDP, snapshots, and replication, along with access control and retention policies, to achieve the specific workflows and service levels each application requires.

The simple and easy-to-use graphical user interface (GUI) incorporates a powerful policy builder that resembles laying out business processes on a whiteboard. Using the GUI, you can easily create and change policies as needed, visualize data copy and movement processes, and align them with business management processes.

Additional features of HDID include:

- Block-level, incremental-forever data capture dramatically reduces the storage capacity needed for copy data, as compared to traditional full and incremental methods.
- Supports a range of storage repositories, including block, file, and object.
- Scales seamlessly to manage petabytes of data.

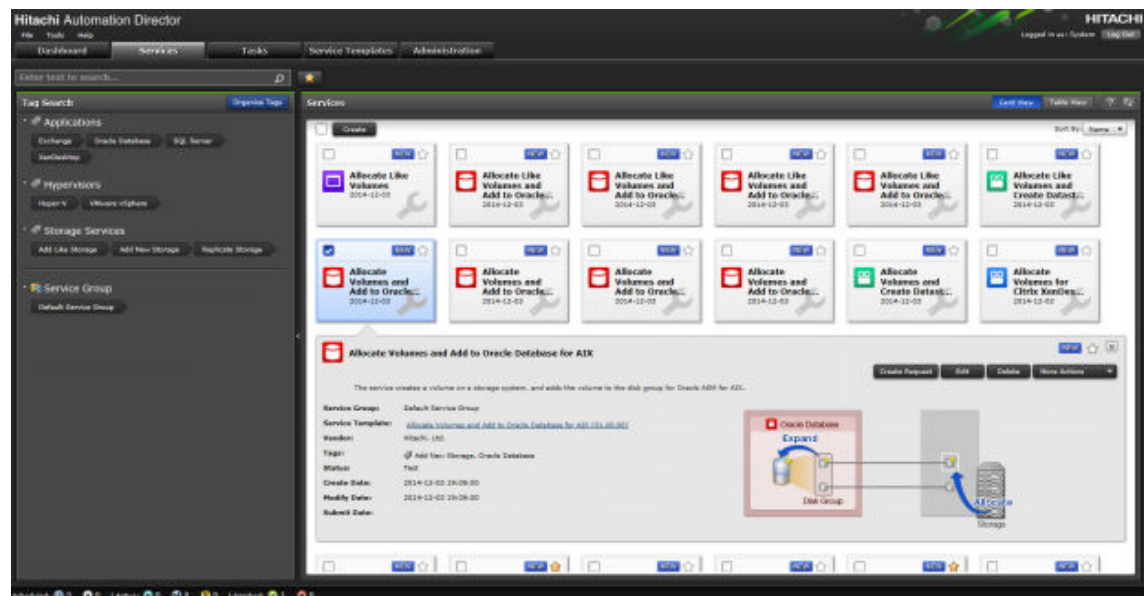
## Hitachi Automation Director



Hitachi Automation Director (HAD) is a software solution that provides tools to automate and simplify end-to-end processes, such as storage provisioning, for storage and data center administrators. The building blocks of the product are prepackaged automation templates known as service templates. These templates can be customized to your specific environment and processes creating services that automate complex tasks such as resource provisioning. When HAD is configured, it integrates with existing Hitachi Command Suite applications, including Hitachi Device Manager and Hitachi Tuning Manager, to automate common infrastructure management tasks by using your existing infrastructure services.

Some of the key features of HAD are:

- Automation services for intelligent provisioning of volumes from different storage classes.
- Preconfigured service templates that help you create customized automation services.
- Role-based access to defined services.
- Intelligent pool selection based on an algorithm that chooses the best pools in terms of performance and capacity.
- Common service management attributes that can be assigned and shared across all automation services.
- A REST API for application integration.
- The ability to create infrastructure groups based on customer needs and environment.



**Figure 1** Select a service on the Services tab to review details and create a request for provisioning.

HAD offers the following benefits:

- Provisioning is simplified through use of service templates that can automate workflow, resulting in additional OPEX savings.
- Service customization can be performed by skilled storage administrators, increasing the efficiency of resource usage and reducing human error.
- Simplified infrastructure management, including classification of storage systems and high-level grouping of resources, significantly improves storage management and provides efficient utilization of resources.
- The ability to customize predefined service templates by using the Service Builder tool, addresses an organization's changing needs.
- The REST API facilitates integration of HAD with relevant IT automation processes.

## Hitachi Data Center Analytics

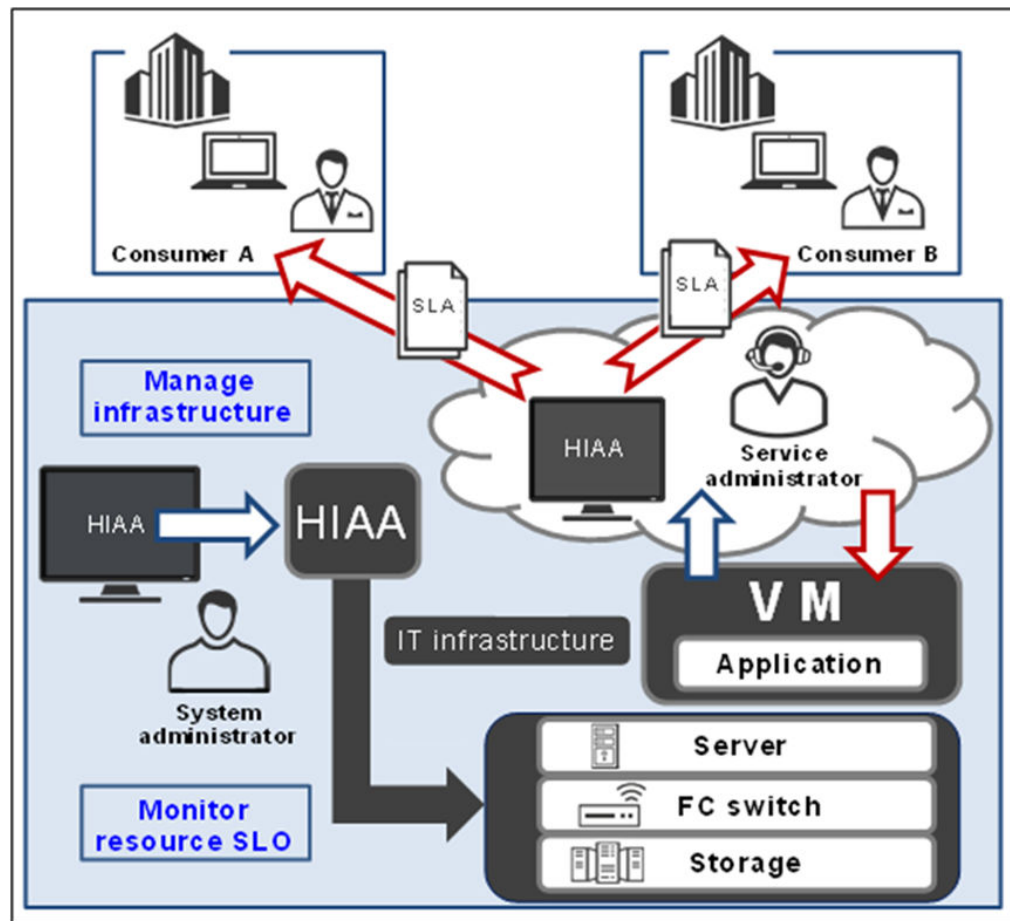
Hitachi Data Center Analytics (HDCA) is a storage performance analytics application that includes a highly scalable data repository and analytics engine for historical performance and capacity trending across the data center. HDCA provides deep and granular performance monitoring and reporting to aid users in identifying infrastructure bottlenecks and trends in order to optimize both application and storage system performance. This software enables a common, centralized storage analytics solution for Hitachi and multi-vendor storage environments, thus reducing the need for vendor-specific performance analytic tools. HDCA provides multi-vendor storage system support for Hitachi and third-party storage system environments.

## Hitachi Infrastructure Analytics Advisor

Hitachi Infrastructure Analytics Advisor (HIAA) is data center management software that monitors, reports, and correlates end-to-end performance from server to storage. HIAA supports monitoring of Hitachi VSP F series and VSP G series storage systems. With HIAA, you can define and monitor storage service-level objectives (SLOs) for resource performance. You can identify and analyze historical performance trends to optimize storage system performance and plan for capacity growth. When a performance hot spot is identified or a service-level threshold is exceeded, the integrated diagnostic engine aids in diagnosing, troubleshooting, and finding the root cause of performance bottlenecks.

Using HIAA, you register resources (storage systems, hosts, servers, and volumes) and set service-level thresholds. You are alerted to threshold violations and possible performance problems (bottlenecks). Using analytics tools, you find which resource has a problem and analyze its cause to help solve the problem.

The following figure shows how HIAA ensures the performance of your storage environment based on real-time SLOs.



The system administrator uses HIAA to manage and monitor the IT infrastructure based on SLOs, which match the service-implementation guidelines that are negotiated under a service-level agreement (SLA) with consumers.

HIAA monitors the health of the IT infrastructure using performance indicators and generates alerts when SLOs are at risk.

Having data center expertise, the service administrator uses HIAA to assign resources, such as VMs and storage capacity from registered storage systems, to consumer applications. This manages critical SLO violations and ensures that service performance meets the SLAs.

## Chapter 4: Software solution examples

The management software for the Hitachi VSP storage systems enables you to increase operational efficiency, optimize availability, and meet critical business requirements.

### Enabling simple and efficient storage provisioning and unified management with Command Suite

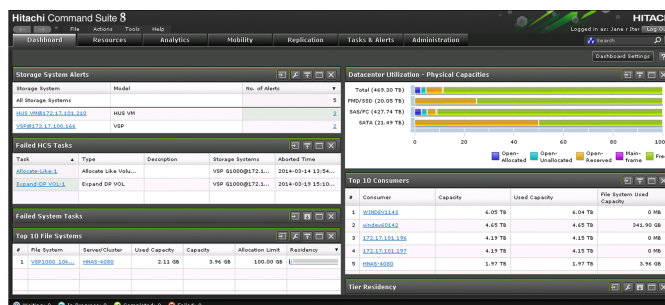

Today, financial institutions provide a wide array of services to their customers. These services must support both structured data (online and ATM transactions, such as withdrawing or depositing checks and cash) and unstructured data (such as email messages, SMS text messages, customer feedback, bank statements, and electronic forms). To meet the ever-increasing need for customer access to the services, the institutions must have a solution that meets the following needs:

- Ability to process customer transactions quickly and accurately. At the same time, provide access to online reports (such as account statements) and forms (such as for opening a new bank account or for applying for a mortgage).
- Flexibility to accommodate structured and unstructured data, and ability to access services no matter where the storage system resides.
- Centralized management of all storage repositories to reduce storage management costs and total cost of ownership.

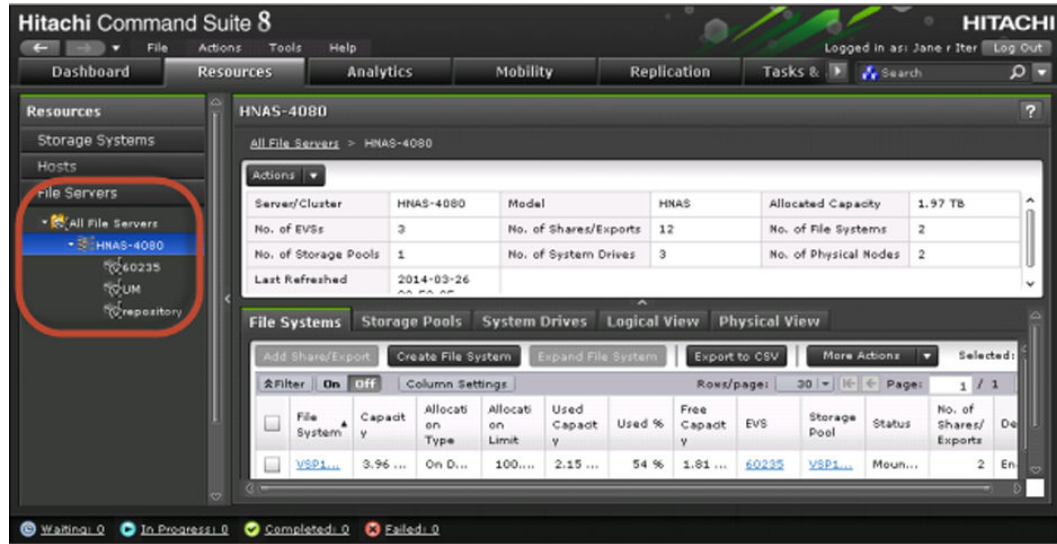
Overall, financial institutions require a platform with the breadth and flexibility to provide services wherever, whenever, and however customers need them.

#### Solution

Hitachi Command Suite (HCS) software consolidates block and file storage arrays to unify the management of all types of data, and provides a single, integrated view for all customers.



HCS natively discovers Hitachi storage systems, Hitachi NAS systems, and Hitachi Data Ingestor file appliance-based systems, displaying the correlation of File Module system drives with back-end physical volumes and File Module storage pools.



HCS discovers and displays related file systems, mount points, and share information for CIFS, and export information for NFS systems. It unifies block, file, and content data across all Hitachi storage and manages all virtualized heterogeneous storage assets.

HCS natively provisions storage to an HNAS cluster the same way as to a physical or hypervisor server, such as the VMware ESX server. It creates and manages file systems, CIFS shares, and NFS exports using the unified, common GUI. Reaching across file, block, content, and application environments, HCS improves business application availability and performance, and expedites access to critical data.

## **Ensuring optimal storage performance and business application service levels with analytics**

Banks offer several incentives to their customers. One such incentive is online banking, which customers have come to prefer. They see the need and growing importance of creating an excellent experience for their online customers. They must provide quick, 24/7 access to online banking services, and must do so across the many devices and platforms used by customers. Customers expect access to these services anytime and from anywhere. If the service is not fast, not available 24/7, and not consistent, customer loyalty can be negatively affected, resulting in bank account closures.

ATM machines provide another critical service to bank customers. ATM transactions have become an essential component of the banking industry. The problem is when ATM machines are not functioning.

Banks strive to keep their business-critical services available for customers, but often find the following problems still exist:

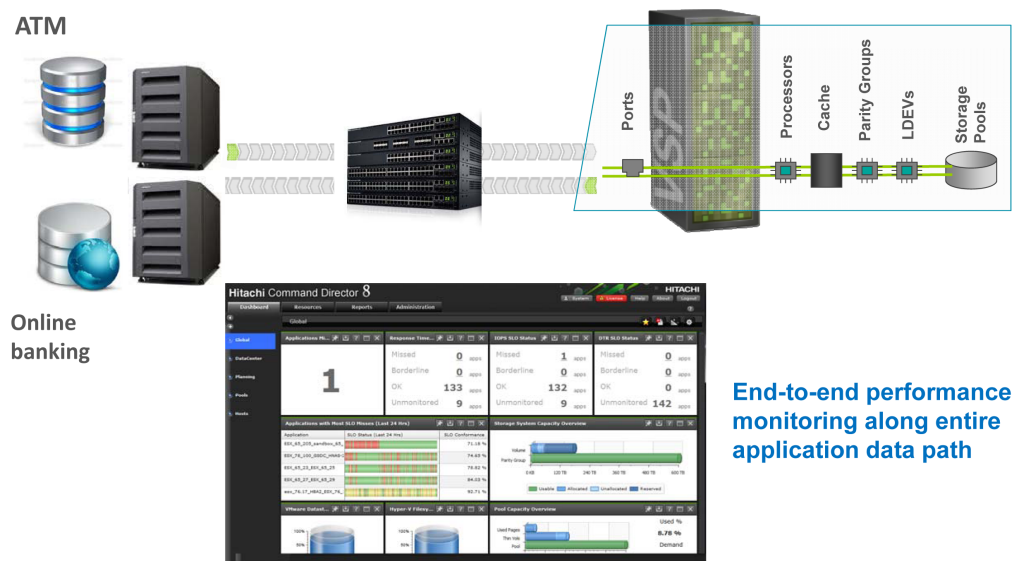
- Lack of performance baselines or benchmarks to analyze response time for online banking and ATM applications
- Insufficient root cause analysis (RCA) techniques that look deep into application performance problems, and ineffective existing techniques
- Absence of real-time monitoring capability and analysis of all elements in the customer environment
- No tools to help storage administrators analyze application performance or to determine if the storage is at fault
- Lack of custom reporting capabilities to obtain detailed storage capacity and performance metrics to gain insight into key storage system performance indicators
- Uncertainty whether critical business applications are meeting required storage service levels

## Solution

Use Hitachi Command Suite Analytics to monitor performance and meet storage service-level needs.

- To help banks determine how well their online banking service is performing, they must know the current level of performance and benchmark it against an industry best practice. Storage downtime affects system availability for online transactions. One of the best ways to avoid bottlenecks is through regular monitoring, system feedback, and on-demand customizable reporting based on parameters defined by users.

The parameters can be based on storage or files, such as EVS, FS, and VVOL utilization, and on capacity reporting, such as on tiers, users, and groups. Instead of reacting to bottlenecks after they occur, administrators can get alerts from HCS Analytics about potential bottlenecks before they occur. Administrators can identify problem performance trends at an earlier stage to avoid system downtime.



HCS Analytics performs end-to-end performance monitoring along the application's entire data path to quickly determine if storage is the source of application-performance degradation. With this monitoring information, storage administrators can take appropriate measures to remove upcoming bottlenecks and to improve storage (and ultimately application) performance.

- To ensure that critical business applications are meeting required storage service levels and comply with storage service-level requirements, storage administrators can use HCS Analytics to accurately monitor application storage levels and quickly resolve problems. Applications have varying service-level objectives (SLO) based on their business criticality. For important applications, such as online banking and ATM transactions, storage administrators can use HCS Analytics to provide the applications with appropriate storage resources in compliance with defined SLO requirements.

## **Management software**

To ensure business application performance and predictive growth, Hitachi Command Suite Analytics provides all the necessary capabilities to find storage resource trouble spots, identify the actual affected storage resources, and help determine the root cause of problems.

HCS Analytics features Tuning Manager: Hitachi Tuning Manager provides comprehensive storage performance monitoring required to maximize both business application and Hitachi storage system performance. It provides integrated performance analytics that can quickly identify, isolate, and find possible causes of performance bottlenecks. Within the HCS central management console, the integrated analytics capabilities provide the necessary first step to quickly address performance problems associated with Hitachi storage environments.

If additional performance details or diagnosis is required, Tuning Manager includes a web-based interface to provide deeper performance monitoring across a comprehensive range of performance and capacity metrics, with historical trending and custom reporting capabilities.

## **Maximizing business application performance and availability with data mobility**

Customer service is a top priority for major commercial and retail banks. They strive to maintain good relationships with, and retain current customers as well as attract new ones. They would also like to achieve faster response times for customer transactions involving personal banking or credit cards, and for potential customers inquiring about their services.

In addition to ensuring the timeliness of critical transactions, banks must provide customers with effective processing of mortgage applications from inception to closing.


Banks must optimize the cost of maintaining data gathered from numerous mortgage applications. While users can tolerate slightly slower response times that are required for transactional systems, they are quickly frustrated by consistently slow responses. In a fast-paced business, older and closed mortgage applications lose business relevance quickly, so it does not make sense to store them on fast storage. A lower tier of storage can be used to achieve effective, long-term archiving of inactive data (such as closed or inactive mortgage applications that companies maintain largely in response to legal requirements).

### **Solution**

A Hitachi Dynamic Tiering (HDT) pool is added to a storage system to support mortgage applications. Using Hitachi Command Suite Mobility, a custom policy is applied to the volumes in the HDT pool that supports the mortgage applications.

The policy is set to ensure that infrequently or never accessed mortgage applications are placed on the lowest cost storage, reducing the total cost of ownership. Conversely, the newest and still-active mortgage applications are promoted to the fastest tier and get the fastest response time.





### Management software

To optimize data access and application Quality of Service, Hitachi Command Suite Data Mobility software places data wherever and whenever it is needed. HCS Data Mobility features Dynamic Tiering, Tiered Storage Manager, and the file-tiering capabilities of the storage system.

- Hitachi Dynamic Tiering automates data lifecycle management at a low cost while delivering top-tier performance to the information most frequently accessed by the business. HDT manages the tiering dynamically. It monitors and manages space utilization at the page level rather than at the file or dataset level. This means that only frequently referenced parts of a file or dataset reside on the highest tier of storage, minimizing the amount of tier 0 storage required for the highly referenced data.

HDT identifies hot spots of frequent access and moves them to the highest tier of storage to improve storage performance. It also moves less frequently referenced pages to lower tiers of storage. All of this occurs with complete transparency to the application.

- Hitachi Tiered Storage Manager (HTSM) proactively matches application performance and availability needs to storage attributes for optimal placement.
- Intelligent file tiering improves performance in file-sharing environments by automatically separating metadata from user data, placing metadata on the fastest storage tier for improved response times, while keeping user data on less expensive storage tiers.

## Delivering storage infrastructure as a service through automated workflows

Financial institutions must provide services 24/7, with almost zero tolerance for outages and inaccessibility to data and information. Storage provisioning plays an integral part in data management. Organizations need to control the complexities associated with storage management and balance operational efficiency. A positive customer experience depends on how the data center is controlled and managed and on the ability to deliver applications in a consistent and timely manner. However, to achieve this objective, customers require a solution to alleviate these pain points:

- Manual storage provisioning processes, which can lead to human errors. Studies show that more than 40% of outages in a storage environment are caused by human error.
- Time-consuming operational inefficiencies
- Cost-inefficient storage provisioning, which can waste storage resources
- A requirement to know infrastructure and environmental details, which allows for no abstraction
- A requirement to manually analyze performance and capacity without any built-in intelligence or automation

### **Solution**

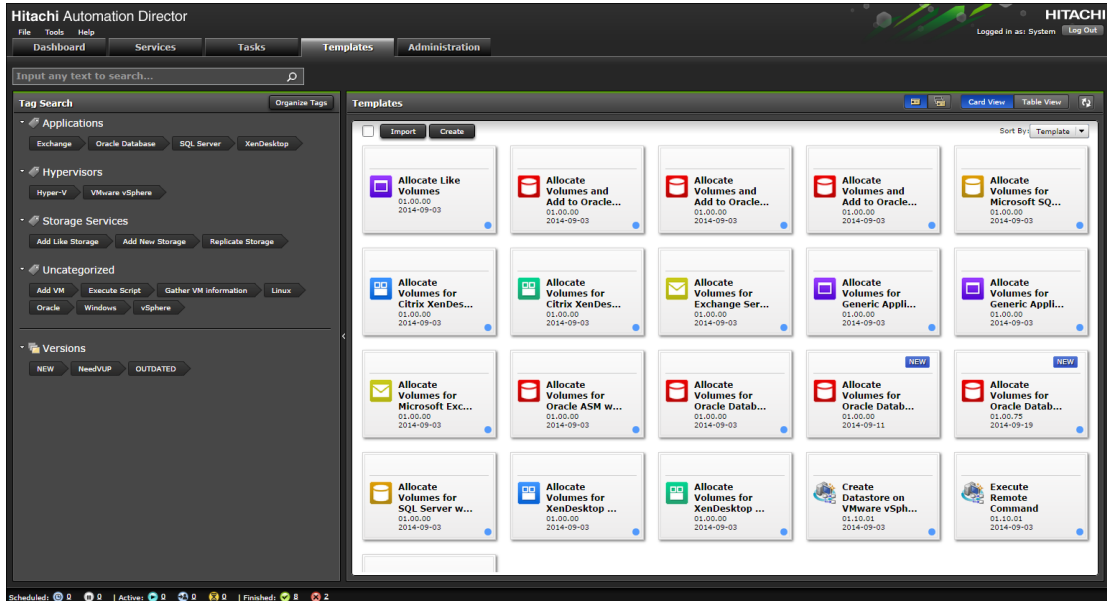
Hitachi Automation Director (HAD) automates manual storage provisioning processes and provides application-based provisioning services that require minimal user input and that intelligently leverage infrastructure resources. HAD provides the following solutions to alleviate the pain points that customers experience in the current environment:

- Implements intelligent automation workflows to streamline the storage provisioning process.
- Provides a catalog of predefined service templates and plugin components that incorporate Hitachi best practices in storage provisioning and that minimize human error.
- Provides customizable storage service templates requiring minimal input that administrative users can use to increase operational efficiency.
- Optimizes storage configurations for common business applications such as Oracle, Microsoft Exchange, Microsoft SQL Server and hypervisors such as Microsoft Hyper-V and VMware.
- Analyzes current storage pool capacity utilization and performance to automatically determine the optimized location for new storage capacity requests and to make storage provisioning more cost-efficient.

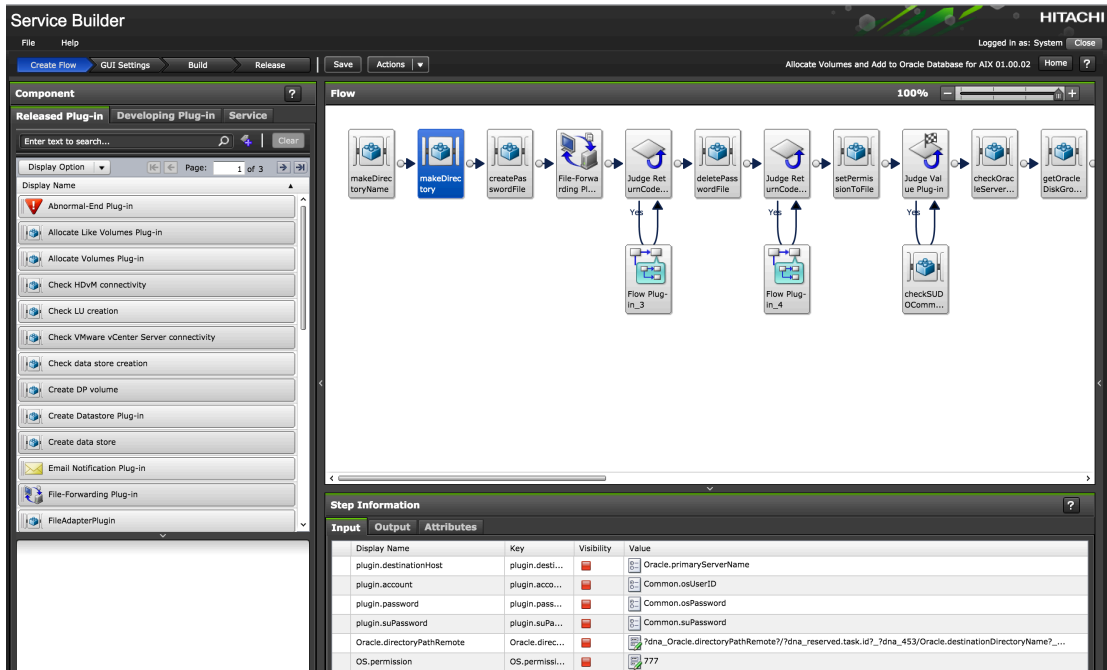
### **Management software**

HAD offers a web-based portal and includes a catalog of predefined workflows that are based on best practices for various applications. These workflows take into account infrastructure requirements for specific applications, including the appropriate storage tier. Capturing the provisioning process with predefined requirements in the workflow, a storage administrator can repeatedly provision infrastructure with simple requests.

After information for provisioning is submitted, the HAD intelligent engine matches the request with the appropriate infrastructure based on performance and capacity analysis. HAD expedites the provisioning process and enables smarter data center management. It provides a REST-based API to integrate provisioning workflows into existing IT management operation applications.



HAD includes a comprehensive tool, Service Builder, to create and modify existing workflows and plug-in components that automate the storage management tasks for a given operating environment.



HAD supports all native block storage systems and third-party storage systems through virtualization technology.

## Data protection for business-critical Oracle databases

Data protection and recovery operations are cited by most customers as one of their top three IT-related challenges. Meanwhile, traditional solutions cannot keep up with rampant data growth, increasing complexity, and distribution of infrastructure. Tighter data availability service-level requirements (backup window, recovery point objective, and recovery time objective) create an impossible situation for line of business owners.

The simple truth is that backup is broken in certain highly important areas, including critical 24x7 applications with large databases.

The business demands that critical data is protected with little or no data loss and with minimal or no performance or availability impact while the data protection occurs.

### Solution

Hitachi Thin Image (HTI) provides fast copies of the production data and Hitachi Universal Replicator (HUR) ensures that there is an asynchronous copy of the data on another storage system in a distant location. Hitachi Data Instance Director (HDID) orchestrates the HTI and HUR data protection activities through a business-objective-driven, whiteboard-like graphical interface, and ensures application consistency for both local and remote snapshots.


The HDID policy is defined in terms of recovery point objectives (RPO) and retention so that new application-aware snapshots are taken to meet each RPO and deleted after the retention period.

### Management software

Hitachi Data Instance Director (HDID) combines modern data protection with business-defined copy data management, simplifying the creation and management of complex data protection and retention workflows.

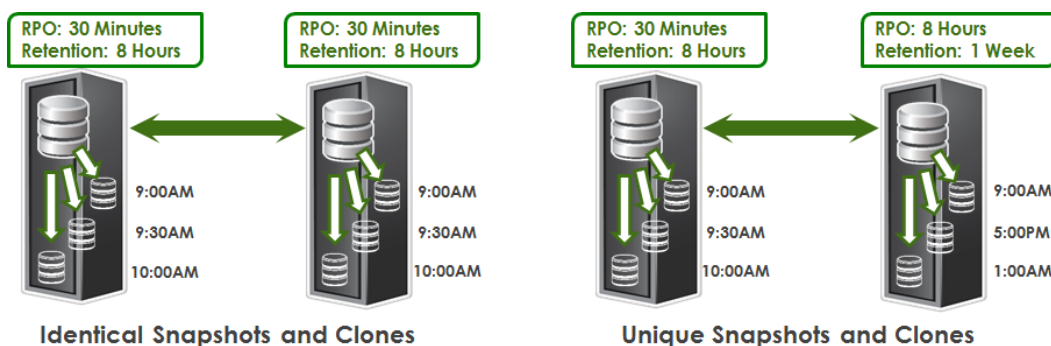
For simplified management, HDID provides a powerful, easy-to-use workflow-based policy engine, so that you can define a data protection workflow within 10 minutes:

- Service-level agreement (SLA)-driven policy enables administrators to define the data classification (such as SQL Server or Oracle), data protection operations, and required SLAs (RPO, data retention).
- Whiteboard-style data flow enables the administrator to define the copy destinations and assign policies to them using drag-and-drop operations. The topological view helps the administrator to visualize the data protection processes and align them with the management requirements.



You can use different methods to back up data across multiple sites, as described in the following table and figure.

Method	Description
Identical snapshots and clones	Provide identical RPO and data retention regardless of location. Keeping identical backups provides identical recovery options and procedures during a site failover, which simplifies the entire restore process.
Unique snapshots and clones	Provide flexible RPO and data retention based on differing business requirements between normal operation and a site failover. Keeping independent backups enables shorter RPOs and lower retention to be set on the local site for quick recovery, while protecting data longer on the remote site.



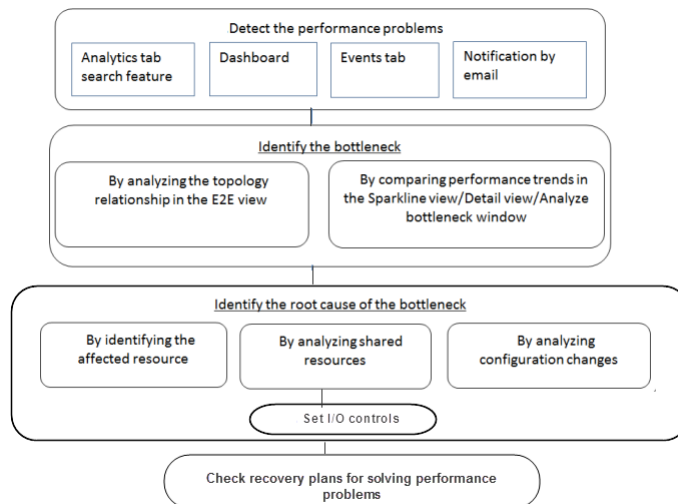
## End-to-end performance troubleshooting using Infrastructure Analytics Advisor

Hitachi Infrastructure Analytics Advisor (HIAA) provides analytical diagnostics to quickly identify, isolate, and determine the root cause of problems.

The traditional approach of troubleshooting performance problems in the unified infrastructure poses several challenges. For example, it can be difficult to identify performance problem in a storage infrastructure environment that includes various virtual machines, servers, network, and storage. The customers are challenged to accurately monitor storage performance and ensure service-level objectives are met, reduce efforts to troubleshoot performance hot spots, and efficiently report across a heterogeneous storage environment.

### Solution

HIAA offers an out-of-the-box analytics solution which lets you identify and troubleshoot performance problems at the component level. The topology view lets you view the graphical representation of the infrastructure components and their dependencies, which is crucial for troubleshooting the infrastructure performance problems. The troubleshooting aids helps in efficient root cause analysis. The analytics workflow is as follows:



## Workflow tasks

### ▪ Detect performance problems

You can view the threshold violations using the Dashboard tab and Events tab. You can configure the system to send email notifications when the threshold values are exceeded. You can also use the search feature in the Analytics tab to find the target resources for performance analysis.

### ▪ Identify performance bottleneck

The performance degradation in the user resources is caused by performance bottlenecks on the server, network, or storage components.

You can identify the resources causing the bottleneck in any of the following views:

- E2E view: The E2E topology view provides detailed configuration of the infrastructure resources and lets you view the relationship between the infrastructure components. You can manually analyze the dependencies between the components in your environment and identify the resource causing performance problems. By using the topology maps, you can easily monitor and manage your resources. You can use this view to monitor resources in your data center from applications, virtual machines, server, network to storage. In the topology view, if a resource has an alert associated with it, error indicators display on the resource icons. The color of the indicator corresponds with the severity of the alert.
- Sparkline view: In the Sparkline view, you can analyze the performance trend graphs of the target resource and the related resources. The Sparkline view displays performance trends for multiple resources in the same pane to enable a quick comparison between different resources. You can show trends of performance metrics of each resource and find the correlation with other resources.

- **Analyze the root cause of the bottleneck**

HIAA integrated troubleshooting aids provide guidance about how to find the root cause of the performance problems. The root cause can be due to the resource contention issues in the shared infrastructure, or due to configuration changes in the environment.


- **Identify affected resources:** You can identify the consumers, hosts, VMs, and volumes that use the bottleneck candidate. You can also verify the status of each resource. Based on the severity level displayed, you can troubleshoot the performance problems associated with the resources.
- **Analyze shared resources:** The performance problem arises in the shared infrastructure when an application or a resource uses the majority of the available resources and causes performance problems for other resources in the shared infrastructure. HIAA supports efficient optimization of the shared infrastructure by quickly identifying the resource contention problems.
- **Analyze related changes:** The configuration changes can sometimes be the source of the performance problem in your environment. HIAA supports the tracking of infrastructure configuration changes. Analyzing these changes and correlating them with the performance data lets you determine the effects of configuration changes on system performance and behavior.
- **Check recovery plans:** You can view the system-generated recovery plans for processor and cache performance bottlenecks.

For details, see the *Infrastructure Analytics Advisor User Guide*.

### **Management software**

HIAA provides comprehensive storage performance management to monitor, report, and correlate end-to-end performance from applications to shared storage resources that is required to optimize both business applications and storage systems. HIAA reporting capabilities enable you to monitor the infrastructure resources and assess their current performance and utilization. Reporting data provides you the information you need to make informed business decisions and plan for future growth. The advanced diagnostic engine aids in rapidly diagnosing, troubleshooting, and finding the root cause of performance bottlenecks.





## Flexible reporting and analysis using Data Center Analytics

In the fast-paced world of online transactions, many companies with global operations have invested in a sophisticated IT infrastructure that provides them a competitive edge. Monitoring and reporting features enable organizations to monitor applications closely and continuously to proactively identify any problems before they manifest into something more severe and requires immediate attention. Whether you are an IT manager for a bank, health care provider, or a government sector, proactive monitoring and reporting are useful in determining the performance trend of your system and addressing ways to improve customer service interactions in advance of customer feedback. To do this thoroughly requires a tool that can help track the health of you system at all hours and display the relevant metrics instantly in a report that you can share with your organization for assessment.

Hitachi Infrastructure Analytics Advisor integrates with Data Center Analytics to provide advanced reporting capability to continuously measure and analyze performance of your monitored resources. The up-to-date visual representation of your system's health enables you to share reports with others. You can create three types of reports:

- **Predefined reports:** provide high-level details at the application level and also a granular report that shows component-level performance data.
- **Ad-hoc reports:** enable you to combine related and unrelated metrics of any monitored resource in one report to review the overall performance impact.
- **Custom reports:** you create with a report builder.

All reports are included in the Reports dock, and are available when you select any storage system object in the storage systems hierarchy. Predefined reports differ based on your selection of the storage system object. An interactive chart and filtering resources enable you to view every detail in any report. You can also filter reports to display the most relevant data, and can print, create a PDF, and export a report to a CSV file.

## Overall and granular level reporting using pre-defined reports

Each node in the tree has predefined reports that cover important attributes of a metric to help your analysis of the resource. If you expand and click a node, for example, 609315f7 under Pools in the tree, the performance report displays. In this case, the Pool IOPS Vs. Response Time report displays and it only shows the metrics data for the 609315f7. No data for other Pools appear on the report.



## Compare node and metric with ad-hoc reports

On the reports, nodes are resources such as RAID Storage 302c7d0 and RAID Storage 302c6d6, and metrics such as cache usage and write pending rate. You can do a comparison between any nodes or between metrics of a single node or different nodes. In Add Report, type the report name in the field, then add specific metrics by dragging and dropping a node from the tree to either the axis section Y/Left or Y1/Right. The left and right axis boxes display the list of available resources, for example, virtual machines and hosts.

If, for example, you want to see a pattern for a storage node between two time periods, you can compare the reports on Storage IOPS to display in one view. Each graph line is color-coded and you can zoom in reports to get a better view.

You can also compare how one metric affects the other metrics. For example, you can create an ad-hoc report that compares IOPS with Response Time. This most commonly used report shows whether an increasing load on the system (IOPS) affects the performance (response time).



To create ad-hoc reports, you can combine the related and unrelated resource metrics and drag and drop the metrics into the report from the specific instances in the tree. For example, you can see the metrics for ports and volumes in one chart at any time. Attributes that are directly related, for example, IOPS and Response Time, usually have a built-in report from the Reports dock. Sometimes, the attributes can be unrelated (or indirect) such as the storage system cache usage from the file system transfer rate on a host can consume most of the storage from the array. You can add unrelated metrics and create a comparison chart.

### Custom reports

If the predefined charts and ad-hoc are not sufficient, you can create custom reports by building your own query. The Custom Reports feature is based on the Data Center Analytics query language. This regex-based expressive query language retrieves and filters the data in the Data Center Analytics database.

The Data Center Analytics query language allows complex analysis on the data in real time with constant run-time. The syntax makes it possible to traverse relations, identify the patterns in the data, and establish a comparison between metrics of a single component or multiple nodes.

The Data Center Analytics UI helps you build your custom query in the following three ways:

- Start with a predefined query and customize it as required.
- Build the query using the Build Query feature.
- Write the query directly using Data Center Analytics query language.

Build Query

vdisk[=capacityInKB rx .\*]&[@virtualDiskTotalReadLatency rx b [U1-U1000][1.]]

Select [VMware Virtual Disk](#) having [+ Configuration Filter](#) [+ Time Series Filter](#) X

[Capacity](#) like .\* X

and [virtual disk total Read Latency](#) values [between](#) 1 and 1000 for [at least](#) 1 consecutive samples X

[Add Resource](#) [Update](#) [Cancel](#)

## Chapter 5: User documentation for VSP G1000, G1500, and VSP F1500

The following tables list the user manuals for VSP G1x00 and VSP F1500. For unified configuration with Hitachi NAS Platform, the file-related documents are also listed.

Subject	Documents
General reference	<p><i>Hitachi Virtual Storage Platform G1000, G1500, F1500 Release Notes</i>, RN-8005xx-Mxxx</p> <p><i>Hitachi Storage Advisor Release Notes</i>, RN-94HSA002</p> <p><i>Hitachi Device Manager Release Notes</i>, RN-00HS266</p> <p><i>Hitachi Tiered Storage Manager Release Notes</i>, RN-00HS279</p> <p><i>Hitachi Data Instance Director Release Notes</i>, RN-93HDID000</p> <p><i>Hitachi Infrastructure Analytics Advisor Release Notes</i>, RN-96HIAA000</p>
Hardware reference	<p><i>Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500</i>, MK-92RD8007</p> <p><i>Hitachi Universal V2 Rack Reference Guide</i>, MK-94HM8035</p>
System installation	<p><i>Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500</i>, MK-92RD8007</p>
Management software installation and configuration	<p><i>System Administrator Guide</i>, MK-92RD8016</p> <p><i>Command Control Interface Installation and Configuration Guide</i>, MK-90RD7008</p> <p><i>Hitachi Storage Advisor Getting Started Guide</i>, MK-94HSA001</p> <p><i>Hitachi Command Suite Installation and Configuration Guide</i>, MK-90HC173</p> <p><i>Hitachi Command Suite System Requirements</i>, MK-92HC209</p> <p><i>Hitachi Command Suite Administrator Guide</i>, MK-90HC175</p> <p><i>Hitachi Command Suite Mainframe Agent Installation and Configuration Guide</i>, MK-96HC130</p> <p><i>Hitachi Infrastructure Analytics Advisor Installation and Configuration Guide</i>, MK-96HIAA002</p> <p><i>Hitachi Data Instance Director Quick Start Guide</i>, MK-93HDID002</p> <p><i>Hitachi SNMP Agent User Guide</i>, MK-92RD8015</p>

Subject	Documents
External storage	<i>Hitachi Universal Volume Manager User Guide</i> , MK-92RD8024
Host configuration	<i>Open-Systems Host Attachment Guide</i> , MK-90RD7037 <i>Mainframe Host Attachment and Operations Guide</i> , MK-96RD645 <i>Provisioning Guide for Open Systems</i> , MK-92RD8014 <i>Provisioning Guide for Mainframe Systems</i> , MK-92RD8013 <i>Hitachi Storage Provider for VMware vCenter (Block VASA Provider) Deployment Guide</i> , MK-90ADPTR010
Provisioning and mobility	<i>Hitachi Command Suite User Guide</i> , MK-90HC172 <i>Hitachi Storage Advisor User Guide</i> , MK-94HSA004 <i>Provisioning Guide for Open Systems</i> , MK-92RD8014 <ul style="list-style-type: none"> <li>▪ Accelerated compression</li> <li>▪ Data Retention Utility</li> <li>▪ Dynamic Provisioning</li> <li>▪ Dynamic Tiering</li> <li>▪ LUN Manager</li> <li>▪ Virtual LUN</li> </ul> <i>Provisioning Guide for Mainframe Systems</i> , MK-92RD8013 <ul style="list-style-type: none"> <li>▪ Dynamic Provisioning for Mainframe</li> <li>▪ Dynamic Tiering for Mainframe</li> <li>▪ Virtual LVI</li> <li>▪ Volume Retention Manager</li> </ul>
Data protection	<b>Open</b> <ul style="list-style-type: none"> <li>▪ <i>Hitachi Thin Image User Guide</i>, MK-92RD8011</li> <li>▪ <i>Hitachi ShadowImage<sup>®</sup> User Guide</i>, MK-92RD8021</li> <li>▪ <i>Hitachi TrueCopy<sup>®</sup> User Guide</i>, MK-92RD8019</li> <li>▪ <i>Hitachi Universal Replicator User Guide</i>, MK-92RD8023</li> <li>▪ <i>Global-Active Device User Guide</i>, MK-92RD8072</li> <li>▪ <i>Hitachi Data Instance Director User Guide</i>, MK-93HDID001</li> </ul>

Subject	Documents
	<p><b>Mainframe</b></p> <ul style="list-style-type: none"> <li>▪ <i>Hitachi ShadowImage<sup>®</sup> for Mainframe User Guide</i>, MK-92RD8020</li> <li>▪ <i>Hitachi TrueCopy<sup>®</sup> for Mainframe User Guide</i>, MK-92RD8018</li> <li>▪ <i>Hitachi Universal Replicator for Mainframe User Guide</i>, MK-92RD8022</li> <li>▪ <i>Hitachi Compatible XRC<sup>®</sup> User Guide</i>, MK-92RD8027</li> <li>▪ <i>Hitachi Compatible FlashCopy<sup>®</sup> / FlashCopy<sup>®</sup> SE User Guide</i>, MK-92RD8010</li> </ul>
Analytics and performance	<p><i>Performance Guide</i>, MK-92RD8012</p> <ul style="list-style-type: none"> <li>▪ Cache Residency Manager</li> <li>▪ Performance Monitor</li> <li>▪ Server Priority Manager</li> <li>▪ Virtual Partition Manager</li> </ul> <p><i>Hitachi Compatible PAV<sup>®</sup> User Guide</i>, MK-92RD8026</p> <p><i>Hitachi Infrastructure Analytics Advisor User Guide</i>, MK-96HIAA001</p> <p><i>Hitachi Infrastructure Analytics Advisor REST API Reference Guide</i>, MK-96HIAA003</p> <p><i>Hitachi Infrastructure Analytics Advisor Data Analytics and Performance Monitoring Overview</i>, MK-96HIAA004</p> <p><i>Hitachi Infrastructure Analytics Advisor Data Center Analytics User Guide</i>, MK-96HIAA005</p> <p><i>Hitachi Data Center Analytics Query Language User Guide</i>, MK-96HDCA005</p> <p><i>Hitachi Data Center Analytics REST API Reference Guide</i>, MK-96HDCA006</p> <p><i>Hitachi Storage Advisor RESTful API Reference Guide</i>, MK-94HSA003</p>
Security	<p><i>Hitachi Command Suite User Guide</i>, MK-90HC172</p> <p><i>Provisioning Guide for Open Systems (Data Retention Utility)</i>, MK-92RD8014</p> <p><i>Provisioning Guide for Mainframe Systems (Volume Retention Manager)</i>, MK-92RD8013</p> <p><i>Encryption License Key User Guide</i>, MK-92RD8009</p> <p><i>Hitachi Volume Shredder User Guide</i>, MK-92RD8025</p>
System maintenance	<p><i>Hitachi SNMP Agent User Guide</i>, MK-92RD8015</p> <p><i>Hitachi Command Suite Audit Log Reference Guide</i>, MK-92HC213</p> <p><i>Hitachi Audit Log User Guide</i>, MK-92RD8008</p>

<b>Subject</b>	<b>Documents</b>
Mainframe software	<p><i>Hitachi Cross-OS File Exchange User's Guide</i>, MK-96RD647</p> <p><i>Hitachi Cross-OS File Exchange Code Converter User's Guide</i>, MK-94RD243</p> <p><i>Hitachi Copy Manager for IBM® TPF Administrator's Guide</i>, MK-92RD129</p> <p><i>Hitachi Copy Manager for IBM® TPF Operations Guide</i>, MK-92RD131</p> <p><i>Hitachi Copy Manager for IBM® TPF Messages and Codes</i>, MK-92RD130</p>
Troubleshooting	<p><i>Hitachi Command Suite Messages</i>, MK-90HC178</p> <p><i>Hitachi Device Manager - Storage Navigator Messages</i>, MK-92RD8017</p> <p><i>Hitachi SNMP Agent User Guide</i>, MK-92RD8015</p>
Command line interface	<p><i>Command Control Interface Command Reference</i>, MK-90RD7009</p> <p><i>Command Control Interface User and Reference Guide</i>, MK-90RD7010</p>

**Table 1 Hitachi Command Suite optional component manuals**

<b>Subject of guide</b>	<b>Documents</b>
Software installation	<p><i>Hitachi Command Suite Tuning Manager Installation Guide</i>, MK-96HC141</p> <p><i>Hitachi Command Suite Command Director Installation and Configuration Guide</i>, MK-90HCMD002</p> <p><i>Hitachi Command Suite Compute Systems Manager Installation and Configuration Guide</i>, MK-91HC195</p> <p><i>Hitachi Command Suite Mainframe Agent Installation and Configuration Guide</i>, MK-96HC130</p> <p><i>Hitachi Command Suite Automation Director Installation and Configuration Guide</i>, MK-92HC204</p> <p><i>Hitachi Command Suite System Requirements</i>, MK-92HC209</p> <p><i>Hitachi Replication Manager Release Notes</i>, RN-00HS280</p> <p><i>Hitachi Tuning Manager Release Notes</i>, RN-00HS262</p> <p><i>Hitachi Compute Systems Manager Release Notes</i>, RN-91HC198</p> <p><i>Hitachi Command Director Release Notes</i>, RN-90HCMD003</p>
Management software setup	<p><i>Hitachi Command Suite Replication Manager Configuration Guide</i>, MK-90HC175</p> <p><i>Hitachi Command Suite Command Director Installation and Configuration Guide</i>, MK-90HCMD002</p> <p><i>Hitachi Command Suite Compute Systems Manager Installation and Configuration Guide</i>, MK-91HC195</p>



<b>Subject of guide</b>	<b>Documents</b>
	<i>Hitachi Command Suite Tuning Manager Agent Administration Guide, MK-92HC013</i>
SAN multipathing	<i>Hitachi Command Suite Compute Systems Manager User Guide, MK-91HC194</i> <i>Hitachi Command Suite Dynamic Link Manager (for AIX®) User Guide, MK-92DLM111</i> <i>Hitachi Command Suite Dynamic Link Manager (for Linux®) User Guide, MK-92DLM113</i> <i>Hitachi Command Suite Dynamic Link Manager (for Solaris) User Guide, MK-92DLM114</i> <i>Hitachi Command Suite Dynamic Link Manager (for Windows®) User Guide, MK-92DLM129</i> <i>Hitachi Command Suite Dynamic Link Manager (for VMware®) User Guide, MK-92DLM130</i> <i>Hitachi Command Suite Global Link Manager Installation and Configuration Guide, MK-95HC107</i> <i>Hitachi Command Suite Global Link Manager Messages, MK-95HC108</i>
Provisioning	<i>Hitachi Command Suite Tiered Storage Manager for Mainframe User Guide, MK-92HC207</i>
Data protection	<i>Hitachi Command Suite Replication Manager User Guide, MK-99HC166</i>
Analytics	<i>Hitachi Command Suite Tuning Manager Server Administration Guide, MK-92HC021</i> <i>Hitachi Command Suite Tuning Manager Agent Administration Guide, MK-92HC013</i> <i>Hitachi Command Suite Tuning Manager User Guide, MK-92HC022</i> <i>Hitachi Command Suite Tuning Manager Hardware Reports Reference, MK-95HC111</i> <i>Hitachi Command Suite Tuning Manager Operating System Reports Reference, MK-95HC112</i> <i>Hitachi Command Suite Tuning Manager Application Reports Reference, MK-95HC113</i> <i>Hitachi Command Suite Tuning Manager CLI Reference Guide, MK-95HC119</i> <i>Hitachi Command Suite Tuning Manager Getting Started Guide, MK-95HC120</i>

<b>Subject of guide</b>	<b>Documents</b>
	<i>Hitachi Command Suite Command Director User Guide, MK-90HCMD001</i>
Automation	<i>Hitachi Command Suite Guide Automation Director User Guide, MK-92HC205</i> <i>Hitachi Command Suite Automation Director Service Builder User Guide, MK-92HC222</i>
Troubleshooting	<i>Hitachi Command Suite Tuning Manager Messages, MK-95HC114</i> <i>Hitachi Command Suite Compute Systems Manager Messages, MK-91HC197</i> <i>Hitachi Command Suite Tiered Storage Manager for Mainframe Messages, MK-92HC210</i>
Command line interface	<i>Command Control Interface Installation and Configuration Guide, MK-90RD7008</i> <i>Hitachi Command Suite CLI Reference Guide, MK-90HC176</i> <i>Hitachi Command Suite Tiered Storage Manager CLI Reference Guide, MK-90HC177</i> <i>Hitachi Command Suite Tuning Manager CLI Reference Guide, MK-95HC119</i> <i>Hitachi Command Suite Compute Systems Manager CLI Reference Guide, MK-91HC196</i> <i>Hitachi Command Suite Command Director CLI Reference Guide, MK-90HCMD004</i>
API	<i>Hitachi Command Suite Tuning Manager API Reference Guide, MK-92HC218</i> <i>Hitachi Command Suite Command Director API Reference Guide, MK-90HCMD005</i> <i>Hitachi Command Suite Automation Director API Reference Guide, MK-92HC217</i>

**Table 2 Hitachi NAS Platform manuals**

<b>Subject of guide</b>	<b>Documents</b>
System installation and configuration	<i>Hitachi NAS Platform and Hitachi Unified Storage File Module System Installation Guide, MK-92HNAS015</i> <i>Hitachi NAS Platform Network Administration Guide, MK-92HNAS008</i>

<b>Subject of guide</b>	<b>Documents</b>
System management	<p><i>Hitachi NAS Platform Storage Systems User Administration</i>, MK-92HNAS013</p> <p><i>Hitachi NAS Platform System Access Guide</i>, MK-92HNAS014</p> <p><i>Hitachi NAS Platform Server and Cluster Administration Guide</i>, MK-92HNAS010</p> <p><i>Hitachi NAS Platform Storage Subsystem Administration Guide</i>, MK-92HNAS012</p>
Data management	<p><i>Hitachi NAS Platform Data Migrator Administration Guide</i>, MK-92HNAS005</p> <p><i>Hitachi NAS Platform File Services Administration Guide</i>, MK-92HNAS006</p>
Data protection	<p><i>Hitachi NAS Platform Replication and Disaster Recovery Administration Guide</i>, MK-92HNAS009</p> <p><i>Hitachi NAS Platform Backup Administration Guide</i>, MK-92HNAS007</p> <p><i>Hitachi NAS Platform Snapshot Administration Guide</i>, MK-92HNAS011</p> <p><i>Hitachi NAS Platform Antivirus Administration Guide</i>, MK-92HNAS004</p>
System maintenance and repair	<p><i>Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference</i>, MK-92HNAS017</p> <p><i>Hitachi NAS Platform and Hitachi Unified Storage Series 4000 Hardware Reference</i>, MK-92HNAS030</p>

**Hitachi Vantara**



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 95050-2639 USA  
[HitachiVantara.com](http://HitachiVantara.com) | [community.HitachiVantara.com](http://community.HitachiVantara.com)

Contact Information  
USA: 1-800-446-0744  
Global: 1-858-5474526  
[HitachiVantara.com/contact](http://HitachiVantara.com/contact)