



Product Overview

Hitachi Virtual Storage Platform G200, G400, G600, G800

Hitachi Virtual Storage Platform F400, F600, F800

© 2015, 2017 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 Data Systems 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 Data Systems Corporation at https://support.hds.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 Data Systems 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 to access relevant data; and
2. Verifying that 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

Preface.....	5
Product version	6
Accessing product documentation.....	6
Getting help.....	6
Comments.....	6
1 Introducing Hitachi Virtual Storage Platform Gx00 and Fx00.....	7
About the Hitachi Virtual Storage Platform Gx00 and Fx00 models.....	8
Application solutions.....	12
2 About the VSP Gx00 and VSP Fx00 hardware.....	15
Hitachi VSP Gx00 and Fx00 hardware overview.....	16
Hitachi VSP Gx00 specifications.....	18
Hitachi VSP Fx00 specifications.....	19
3 Software components and features of VSP Gx00 and VSP Fx00.....	21
Hitachi Storage Virtualization Operating System (SVOS).....	22
Capacity saving function: deduplication and compression.....	25
Accelerated compression.....	26
Dynamic tiering and active flash.....	27
High availability with global-active device.....	28
Nondisruptive migration.....	32
Centralized management.....	35
Overview of Storage Advisor.....	36
Hitachi Command Suite.....	37
Advanced SAN multipathing.....	39
Data replication.....	40
Data mobility.....	41
Hitachi Infrastructure Analytics Advisor.....	42
Product overview.....	42
Hitachi Data Center Analytics.....	43
Hitachi Data Instance Director.....	44

Hitachi Automation Director.....	45
4 Software management examples.....	49
Enabling simple and efficient storage provisioning and unified management with Command Suite.....	50
Ensuring optimal storage performance and business application service levels with analytics	51
Maximizing business application performance and availability with data mobility	54
Delivering storage infrastructure as a service through automated workflows.....	55
Data protection for business-critical Oracle databases	58
End-to-end performance troubleshooting using Infrastructure Analytics Advisor.....	60
Flexible reporting and analysis using Data Center Analytics.....	66
5 User documentation for VSP Gx00 and VSP Fx00.....	71



Preface

This guide provides an overview of the Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00) and Hitachi Virtual Storage Platform F400, F600, F800 (VSP Fx00) storage systems, including hardware components, general system specifications, software features, management interfaces, user documentation, and software management examples.

- [Product version](#)
- [Accessing product documentation](#)
- [Getting help](#)
- [Comments](#)

Product version

This document revision applies to the following product versions:

- VSP Gx00 and Fx00: firmware version 83-04-2x or later
- Hitachi Storage Virtualization Operating System: v7.1 or later
- NAS firmware: v13.1 or later
- Hitachi Storage Advisor: v2.1.1 or later
- Hitachi Command Suite: v8.5.1 or later
- Hitachi Infrastructure Analytics Advisor: v3.0 or later
- Hitachi Data Instance Director: v5.3 or later

Accessing product documentation

Product user documentation is available on Hitachi Data Systems Support Connect: <https://knowledge.hds.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 Data Systems Support Connect](#) is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html.

[Hitachi Data Systems Community](#) is a global online community for HDS 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.hds.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hds.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 Data Systems Corporation.

Thank you!

Introducing Hitachi Virtual Storage Platform Gx00 and Fx00

Based on Hitachi's industry-leading storage technology, the Hitachi Virtual Storage Platform G200, G400, G600, G800 (VSP Gx00 models) and the all-flash Hitachi Virtual Storage Platform F400, F600, F800 (VSP Fx00 models) include a range of versatile, high-performance storage systems that deliver flash-accelerated scalability, simplified management, and advanced data protection.

- [About the Hitachi Virtual Storage Platform Gx00 and Fx00 models](#)
- [Application solutions](#)

About the Hitachi Virtual Storage Platform Gx00 and Fx00 models

The Hitachi VSP Gx00 models and the all-flash VSP Fx00 models provide unparalleled performance, efficiency, and reliability. The Hitachi VSP Gx00 and Fx00 models feature complete system redundancy, hot-swappable parts, nondisruptive updates, and hardware-based compression, and they use the same common operating system, network file services, management, migration, virtualization, replication, and data protection services as the industry-leading Hitachi VSP G1000, G1500, and F1500 high-end arrays.

Key features

Key features of the VSP Gx00 and Fx00 models include:

- Up to 4M IOPS performance
- 100% data-availability guarantee
- Cloud-optimization
- Enhanced integration for VMware®, Windows®, and Oracle® environments
- Embedded NAS functionality
- Advanced active-active clustering, replication, and snapshots
- Active flash tiering and groundbreaking flash modules

Embedded NAS modules

Hitachi VSP G400, G600, G800 and VSP F400, F600, F800 can be configured with embedded NAS modules to provide native NAS functionality in a compact and powerful package. This space-optimized design retains Hitachi's unique field-programmable gate array (FPGA) NAS architecture, enabling you to easily deploy a scalable SAN, NAS, and cloud-optimized platform throughout your enterprise.

Cloud-optimization

With many enterprises implementing both private and public cloud services as part of their overall IT strategy, the ability to take advantage of this hybrid data migration solution is critical. The data migrator to cloud feature enables policy-driven, user-transparent, and automatic file tiering of less used (cold) files from unified models to private clouds, such as Hitachi Content Platform, and public clouds, such as Amazon S3 or Microsoft® Azure™. This approach frees up storage resources for more frequently accessed applications for Tier 1 storage, thus reducing overall storage expenditures.

Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System (SVOS) 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 provides the foundation for superior storage performance, high availability, and IT efficiency. The enterprise-grade capabilities in SVOS 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. Additional features of SVOS include:

- 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

Hitachi offers the following management software for centralized storage management:

- **Hitachi Storage Advisor:** Hitachi Storage Advisor is a common configuration management application for the Hitachi Virtual Storage Platform G series and F series, including VSP G1000, VSP G1500, VSP F1500, VSP Gx00, and VSP Fx00 storage systems. With simplified key administrative operations, Hitachi Storage Advisor makes storage management feasible for IT generalists and less time-consuming for experienced IT administrators. Based on open standard APIs, the interface provides easy configurations that abstract complex management operations into fewer steps. Recommended configuration practices can be used to streamline system configuration and unified storage provisioning operations.

Hitachi Storage Advisor enables a simplified, unified, and extensible approach to managing storage across data types, including both block and file storage requirements.

- **Unified onboarding:** Storage Advisor enables you to add block and file storage in a single step. The only requirements are the SVP IP address, user name, and password. When the file storage is added, the cluster is automatically registered in Storage Advisor.
- **Unified configuration:** Once a storage system is onboarded, all block and file resources can be configured and managed from a single storage system management page. File pool creation workflow incorporates best practices that simplify workflow and enhance usability. The file pools are

used to easily create virtual file servers, file systems, and shares and exports. File system creation automatically mounts and formats the new file system.

- **Unified reporting:** Capacity is reported for all aggregated storage systems in the dashboard. Capacity is also displayed for individual storage systems. Three views of capacity are available: file only, block only, or a unified view of block and file.
- **Hitachi Command Suite:** For more complex storage environments, Hitachi Command Suite (HCS) provides an application-centric storage management solution that simplifies administration of a common pool of multi-vendor 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.
- **Command control interface:** The command control interface (CCI) software enables you to perform storage system configuration and data management operations by issuing commands to enterprise storage systems. CCI provides command-line access to the same provisioning and storage management operations that are available in Hitachi Command Suite and Hitachi Device Manager - Storage Navigator as well as in-system replication, remote replication, and data protection operations. CCI commands can be used interactively or in scripts to automate and standardize storage administration functions, simplifying the job of the storage administrator and reducing administration costs.
- **Hitachi Infrastructure Analytics Advisor:** Hitachi Infrastructure Analytics Advisor (HIAA) is a data center management software that monitors, reports, and correlates end-to-end performance from server to storage. HIAA supports monitoring of Hitachi VSP G series and VSP F series storage systems. With Infrastructure Analytics Advisor, 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, integrated diagnostic engine aids in diagnosing, troubleshooting, and finding the root cause of performance bottlenecks.
- **Hitachi Data Center Analytics:** Hitachi Data Center Analytics (HDCA) delivers a highly scalable data repository and analytics engine to help you to easily identify data center performance trends and properly optimize application and storage system performance. With its heterogeneous storage system support, Data Center Analytics enables an efficient storage analytics solution that can span across a multivendor storage infrastructure. Data Center Analytics is available separately and also integrated with Hitachi Infrastructure Analytics Advisor in a software package called Hitachi Performance Analytics.
- **Hitachi Data Instance Director:** Hitachi Data Instance Director (HDID) unifies and simplifies copy data management. Through its unified, whiteboard-like, workflow-based policy engine, the data managed by HDID can be used for backup, archive, disaster recovery, test and development,

and many other purposes. HDID also automates and orchestrates Hitachi storage-based copy technologies, including snapshots, clones, and replication.

- **Hitachi Automation Director:** Hitachi Automation Director (HAD) enables storage infrastructure self-service with intelligent automated workflows that incorporate storage management best practices. Through infrastructure abstraction, common and repeatable storage management tasks can be simplified, improving reliability and helping to deliver new IT services quickly to the business.

Hitachi Accelerated Flash storage

Hitachi Accelerated Flash (HAF) storage delivers best-in-class performance and efficiency in the Hitachi VSP G series and VSP F series storage systems.

HAF features patented flash module drives (FMDs) that are rack-optimized with a highly dense design that delivers greater than 338 TB effective capacity per 2U tray based on a typical 2:1 compression ratio. IOPS performance yields up to five times better results than that of enterprise solid-state drives (SSDs), resulting in leading performance, lowest bit cost, highest capacity, and extended endurance. HAF integrated with SVOS enables leading, real-application performance, lower effective cost, and superior consistent response times. Running on VSP Gx00 models and VSP Fx00 models, HAF with SVOS enables transactions executed within sub-millisecond response even at petabyte scale.

- **Key features:** HAF delivers outstanding value compared to enterprise SSDs. When compared to small-form-factor 1.92-TB SSDs, the HAF drives deliver:
 - Up to three times the per-device random read 8 KB I/O performance
 - Up to five times the per-device random write 8 KB I/O performance
 - Up to 70% lower effective bit price
 - 60% lower response time at peak load
 - 60 times faster formatting speed
 - Superior data integrity
- **Second-generation flash modules:** Second-generation flash module (FMD DC2) drives are designed to support concurrent, large I/O enterprise workloads and enable hyperscale efficiencies. At their core is an advanced embedded multicore flash controller that increases the performance of multi-layer cell (MLC) flash to levels that exceed those achieved by more expensive single-level cell (SLC) flash SSDs. Their inline compression offload engine and enhanced flash translation layer empower the drives to deliver up to 80% data reduction (typically 2:1) at 10 times the speed of competing drives. With more raw capacity and inline, no-penalty compression, FMD DC2 enables up to five times more write IOPS than SSDs.

Nondisruptive migration

Hitachi offers nondisruptive migration capabilities in two options. You can choose the service that Hitachi Data Systems Global Services Solutions (GSS) offers through its migration center-of-excellence team, or self-service delivery built on Hitachi Command Suite management software. Both options accelerate data center infrastructure transformations by leveraging Hitachi's latest enterprise storage technology. These offerings enable 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 Command Suite nondisruptive migration option is designed for experienced administrators who prefer a self-service option for traditional one-to-one platform refresh.
- Nondisruptive migration from GSS migration center of excellence team is intended for more complex, large-scale, heterogeneous and replicated data center environments.

For more information about nondisruptive migration options, contact your Hitachi Data Systems team or visit <http://www.hds.com>.

Application solutions

Hitachi's portfolio of advanced storage solutions enables you to solve your application infrastructure challenges and achieve the highest application service levels with solutions made for resilience and speed. Make your business application investments count with proven solutions for converged, cloud, storage, server, database, and other applications.

VMware® support

Hitachi VSP Gx00 models and VSP Fx00 models are aligned with the VMware® software-defined storage vision, providing the following support:

- **Hitachi Storage Provider for VMware vCenter:** Hitachi Storage Provider works with VMware vSphere® API for Storage Awareness (VASA) to provide access to Hitachi Virtual Storage Platform Gx00 and Fx00, Hitachi NAS 4000 series, and VSP Gx00 and Fx00 models with embedded NAS modules. Storage Provider enables policy-based storage management using VMware Storage Policy-based Management (SPBM) and VMware Virtual Volumes (VVols). VVols are supported for both block and file on unified VSP Gx00 and Fx00 models. In the management operation, you can create a VVol datastore by selecting a Storage Container without any special knowledge of the storage system. You can create a virtual machine (VM) by setting policies (such as high IOPS and reliability) and can also verify whether the VM complies with these policies. Storage Provider provides a simplified method for VMware admins and storage admins to deliver effective storage that meets advanced VM requirements.

- **vSphere® Storage APIs - Array Integration (VAAI):** VAAI uses storage-native file cloning technology to drive the cloning of VMs from within VMware®'s standard GUI. This functionality is displayed in VMware® as "Hardware Acceleration". Cloning a VM using file clones is substantially faster and more space efficient than traditional host-based copy software.
- **Hitachi Storage Replication Adapter (SRA):** Hitachi Storage Replication Adapter (SRA) for VMware® Site Recovery Manager provides a disaster recovery (DR) solution that works with both your storage environment and your VMware® environment. Supporting both block and file, Hitachi SRA provides an automated replication solution between the production and recovery sites to replicate the storage hosting the VMs, enabling running VMs to seamlessly fail over to the recovery site. Arrays at both sites are "paired" during Site Recovery Manager configuration, and VMware® administrators use the Site Recovery Manager application to manage the configuration and definition of the DR plan.
- **Hitachi Virtual Infrastructure Integrator (V2I):** Hitachi Virtual Infrastructure Integrator (V2I) provides managed data protection/clone services for VMs and VMware® datastores and enhanced storage management and best practices validation for VM administrators. Hitachi V2I enables you to leverage VM-level storage hardware snapshots, resulting in improved RTO/RPO. Both VM-consistent and application-consistent backup are provided, and new VMs are automatically protected. Hitachi V2I also provides visibility to NFS storage services so you can see which VMs are sharing the same datastore and which NAS system is serving a specific datastore. Hitachi V2I API services are enabled for Hitachi Data Instance Director (HDID) integration.
- **vStorage API for Multipathing (VAMP):** Hitachi VSP Gx00 models and VSP Fx00 models support VAMP to provide enhanced control of I/O path selection and failover.
- **vStorage API for Data Protection (VADP):** Hitachi VSP Gx00 models and VSP Fx00 models support VADP to enable backup applications to perform file-level or VM-level backup of running VMs.
- **Hitachi Storage Manager for VMware® vCenter:** Hitachi Storage Manager for VMware® vCenter is a plug-in that provides integrated management of Hitachi storage systems within vCenter.

Microsoft Windows® support

Server virtualization integration with leading virtual server platforms gives you end-to-end visibility from an individual virtual machine to the storage logical unit and protects largescale multivendor environments. Support for Microsoft Windows® 2012 (including Microsoft Hyper-V) and Systems Center includes:

- Microsoft Virtual ShadowCopy Service (VSS)
- 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 Systems Center Orchestrator

Oracle® support

Hitachi Data Systems has developed and supported IT solutions for many of the world's largest companies with the most demanding Oracle® database environments, solutions that maximize business value, enhance your progress toward greater business outcomes, and ensure that you continue to see great performance from your Oracle® systems.

- Hitachi storage and server adapters for Oracle® databases provide integrated tools for converged infrastructure management and data protection. 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 drivers for Oracle® environments enhance consolidation, performance, and efficiency.
- The Database Infrastructure Evaluation Tool (DIET), available to Oracle® database administrators at no cost, analyzes your entire Oracle® database environment and provides best practices and expert recommendations on areas for improvement to ensure your storage, compute, and converged infrastructure operates at peak utilization.
- Hitachi Dynamic Provisioning gives your Oracle® applications the right amount and right type of storage to maximize performance and efficiency.
- Hitachi Dynamic Tiering offers finely tuned performance for Oracle®, automatically keeping the most crucial data on the fastest storage.

About the VSP Gx00 and VSP Fx00 hardware

The Hitachi VSP Gx00 models and the all-flash VSP Fx00 models are rack-mountable storage systems that incorporate state-of-the-art virtualization and advanced data-management and fault-tolerant technologies.

- [Hitachi VSP Gx00 and Fx00 hardware overview](#)
- [Hitachi VSP Gx00 specifications](#)
- [Hitachi VSP Fx00 specifications](#)

Hitachi VSP Gx00 and Fx00 hardware overview

The architecture of the Hitachi VSP Gx00 and Fx00 models accommodates scalability to meet a wide range of capacity and performance requirements. The storage systems can be configured with the desired number and types of front-end module features for attachment to a variety of host processors. All drive and cache upgrades can be performed without interrupting user access to data, allowing you to hot add components as you need them for pay-as-you-grow scalability.

The VSP Gx00 and Fx00 models have dual controllers that provide the interface to a data host. Each controller contains its own processor, dual in-line cache memory modules (DIMMs), cache flash memory (CFM), battery, and fans and is provided with an Ethernet connection for out-of-band management using Hitachi Device Manager - Storage Navigator. If the data path through one controller fails, all drives remain available to hosts using a redundant data path through the other controller. The VSP Gx00 and Fx00 models allow a defective controller to be replaced.

VSP Gx00 models support a variety of drives, including HDDs, SSDs, FMDs, and the second-generation Hitachi Accelerated Flash (HAF) flash module drives (FMD DC2) that are the foundation of the VSP Fx00 models. The FMD DC2 drives deliver improved response times, effective flash capacity, and return on investment. The VSP Gx00 models combine FMD DC2 flash with inline compression and Hitachi Dynamic Tiering active flash for improved responsiveness and efficiency, making it simple to move to all-flash gradually over time. The storage systems allow defective drives to be hot swapped without interrupting data availability. A hot spare drive can be configured to replace a failed drive automatically, securing the fault-tolerant integrity of the logical drive. Self-contained, hardware-based RAID logical drives provide maximum performance in compact external enclosures.

All models are compatible with most industry-standard, 4-post EIA 19-inch racks with square mounting holes. The storage system can be factory-configured and shipped in a Hitachi V2 Universal Rack (600 x 1200 mm) or shipped without a rack for installation into an existing compatible rack. The Hitachi rack comes with either two or four power distribution units (PDUs), depending on the PDU model selected.

Hitachi VSP Fx00 models

The Hitachi Virtual Storage Platform F400, F600, F800 all-flash arrays bring together all-flash storage and the simplicity of built-in automation software with the proven resiliency and performance of Hitachi VSP technology. The VSP Fx00 models offer 1.4 million IOPS to meet the most demanding application requirements. The VSP Fx00 models leverage the FMD DC2 second-generation flash-module drives, which deliver twice the capacity per tray over previous-generation flash module drives, resulting in a smaller

data-center footprint and lower maximum-performance cost. Fast inline data compression with an ASIC engine reduces space with no performance penalty.

Easy-to-use replication management is included with the VSP Fx00 models with optional synchronous and asynchronous replication available for complete data protection. The VSP Fx00 models range in storage capacity from 14 TB (raw) up to 1.4 PB effective flash capacity and provide an all-flash solution that works seamlessly with other Hitachi infrastructure products through common management software and rich automation tools.

Unified models with embedded NAS modules

Hitachi VSP G400, G600, G800 and VSP F400, F600, F800 can be configured with NAS modules to deliver native NAS functionality in a unified storage platform that combines the power of VSP storage and SVOS with the advanced capabilities of Hitachi NAS Platform. A single SVOS version provides all block and file features for the unified models, enabling you to leverage the advanced replication and disaster recovery features for protecting your virtual machines (VMs). The unified models offer a single platform that can satisfy block and file requirements for all deployments in solution-centric enterprise environments.

The unified VSP Gx00 and Fx00 models provide a two-node "cluster in a box" that is formed automatically on installation with no external cabling required, eliminating the complexity and additional costs of attached NAS hosts. A full SAN, NAS, and cloud-optimized storage system with embedded NAS modules is only 4U high. This space-optimized design retains the unique field-programmable gate array (FPGA) NAS architecture in a new, compact, and powerful package. Existing storage systems can be upgraded to add the NAS modules with no additional rack space required.

Service processor (SVP)

Hitachi VSP Gx00 and Fx00 models work with a service processor (SVP). The SVP provides out-of-band configuration and management of the storage system and collects performance data for key components to enable diagnostic testing and analysis.

The SVP is available as a physical device provided by Hitachi Data Systems or as a software application:

- The physical SVP is a 1U management server that runs Windows Embedded Standard 7.
- The SVP software application is installed on a customer-supplied server and runs on a customer-supplied version of Windows. For details about hardware and OS support, see the *Service Processor Technical Reference for VSP Gx00 and VSP Fx00 Models*.

Hitachi VSP Gx00 specifications

The following table provides the key specifications for the VSP Gx00 models.

Feature	VSP G200	VSP G400	VSP G600	VSP G800
Maximum internal capacity (raw)	2,472 TB	4,800 TB	7,200 TB	14,400 TB
Maximum external capacity (raw)	8 PB	16 PB		64 PB
Solid-state drive options (raw, SFF)	200 GB, 400 GB, 1.92 TB			
Flash module drive options	1.7 TB, 3.5 TB, 7 TB, 14 TB			
SFF hard disk drive options	15K RPM: 300 GB, 600 GB 10K RPM: 600 GB, 900 GB, 1.2 TB, 1.8 TB			
LFF hard disk drive options	7.2K RPM: 4 TB, 6 TB, 10 TB			
Maximum number of drives	264	480	720	1,440
Disk expansion trays	2U: 24 SFF (2.5"); 12 LFF (3.5"); 12 FMD 4U: 60 LFF (3.5") and SFF (2.5")			
Block module height	2U	4U		
Host interfaces (with no drives)	16 FC: 8, 16, or 32 Gbps 8 iSCSI: 10 Gbps 8 iSCSI: 10GBase-T	64 FC: 8, 16, or 32 Gbps 32 iSCSI: 10 Gbps 32 iSCSI: 10GBase-T		80 FC: 8, 16, or 32 Gbps 40 iSCSI: 10 Gbps 40 iSCSI: 10GBase-T
Maximum cache	64 GB	128 GB	256 GB	512 GB
Maximum LU size	256 TB	256 TB		256 TB
Maximum number of LUs	2,048	4,096		16,384
RAID levels	RAID-1+0, RAID-5, RAID-6			
Maximum RAID groups	84	240		480
Internal NAS module options	--	Each VSP G400, VSP G600, and VSP G800 can be ordered with two internal NAS modules. When fully configured with two internal NAS modules, each is 4U in height. Each NAS module has 6 x 10-GbE ports for a total of 12 x 10-GbE ports per VSP G400, VSP G600, or VSP G800.		

Feature	VSP G200	VSP G400	VSP G600	VSP G800
File module height	3U per node (Hitachi NAS Platform 4000 series)			
Nodes per cluster, cache per node	1-2 nodes, 48 GB per node	1-4 nodes, 48 GB per node		
File system size	1 PB pool, single namespace up to capacity			
Maximum file systems	128 file systems			
Maximum snapshots	1,024 per file system			
Protocols	NFS, SMB, FTP, iSCSI, and HTTP to the cloud			
Fibre Channel ports	4 x 8-Gbps ports per node			
Ethernet ports	4 x 10-GbE, 6 x 16-GbE ports per node			

Hitachi VSP Fx00 specifications

The following table provides the key specifications for the VSP Fx00 models.

Feature	VSP F400	VSP F600	VSP F800
Performance	Up to 600,000 IOPS 11 GB/sec bandwidth	Up to 800,000 IOPS 12 GB/sec bandwidth	Up to 1,400,000 IOPS 24 GB/sec bandwidth
Cache	128 GB	256 GB	512 GB
Connectivity	32 FC: 8, 16, or 32 Gbps 16 iSCSI: 10 GB/sec 16 iSCSI: 10GBase-T		48 FC: 8, 16, or 32 Gbps 24 iSCSI: 10 GB/sec 24 iSCSI: 10GBase-T
Configurations	Small with 8 flash modules and 1 spare Medium with 16 flash modules and 1 spare Large with 40 flash modules and 2 spares		
Capacity range	Available with 1.75-TB, 1.92-TB, 3.5-TB, 7-TB, or 14-TB modules From 14 TB to 280 TB raw, up to 1.4 PB effective flash capacity		
Expansion paths	Add a single shelf of 8 flash modules or a dual shelf of 25 flash modules		
Maximum flash modules	40		
Average data reduction ratio	5:1		
Internal NAS module options	Each VSP F400, VSP F600, and VSP F800 can be ordered with two internal NAS modules. When fully configured with two internal NAS modules, each is 4U in height.		

Feature	VSP F400	VSP F600	VSP F800
	Each NAS module has 6 x 10-GbE ports for a total of 12 x 10-GbE ports per VSP F400, VSP F600, or VSP F800.		
High availability	<p data-bbox="878 266 1211 291">Active-active N+1 architecture</p> <p data-bbox="875 315 1214 340">100% data availability warranty</p> <p data-bbox="773 363 1317 388">Active mirroring option for 100% data accessibility</p>		

Software components and features of VSP Gx00 and VSP Fx00

Hitachi VSP Gx00 models and VSP Fx00 models are powered by Hitachi Storage Virtualization Operating System (SVOS) and supported by Hitachi storage management software to enable you to effectively manage and centralize your software-defined infrastructure.

- [Hitachi Storage Virtualization Operating System \(SVOS\)](#)
- [Centralized management](#)

Hitachi Storage Virtualization Operating System (SVOS)

Hitachi Storage Virtualization Operating System (SVOS) is the standard operating system for Hitachi VSP G series and VSP F series. An integrated software system, SVOS works with the virtualization capabilities of the storage systems and provides the foundation for global storage virtualization. SVOS delivers software-defined storage by abstracting and managing heterogeneous storage to provide a unified virtual storage layer, resource pooling, and automation. SVOS also offers self-optimization, automation, and increased operational efficiency for improved performance and storage utilization.

SVOS features include:

- **Global storage virtualization capability** of global-active device provides active-active clustering environment spanning multiple storage systems.
- **Hitachi Universal Volume Manager** enables virtualization of external heterogeneous storage.
- **Hitachi Dynamic Provisioning** provides thin provisioning for simplified provisioning operations, automatic performance optimization, and storage space savings.
- **Hitachi Dynamic Tiering** delivers automated tiering to optimize performance in real time.
- **Capacity saving functions** reduce your bitcost for the stored data by deduplicating and compressing the data.
- **Hitachi local replication**, including Hitachi Thin Image and Hitachi ShadowImage, delivers convenient and cost-effective local point-in-time (PIT) data copies using storage-based, space-efficient snapshots and full-volume clones.
- **Cascaded snapshot configuration and full-copy cloning** enables creation of snapshot trees and full-copy clones from snapshots using Hitachi Thin Image.
- **Hitachi remote replication**, including Hitachi TrueCopy and Hitachi Universal Replicator, enables disaster recovery and business continuity spanning up to three data centers using a combination of synchronous and asynchronous replication plus advanced journaling and delta synchronization to minimize data loss.
- **Hitachi Resource Partition Manager** supports secure administrative partitions for multitenancy requirements.
- **Cache partition manager** feature supports up to 32 cache partitions.
- **Performance monitor feature** provides a graphical interface to assist with performance configuration planning, workload balancing, and analyzing and optimizing storage system performance.
- **Data-at-rest encryption** provides encryption of data at rest for internal storage to protect sensitive information.
- **Storage-system-based utilities** include LUN Manager, Virtual LVI (customized volume size), Hitachi Data Retention Utility, volume port

security, Hitachi Server Priority Manager, audit log, command control interface, and volume shredder feature.

- **Standard management interface support** includes SMI-S provider, SNMP agent, Volume Shadow Copy Service (VSS), and REST.
- **Hitachi Dynamic Link Manager Advanced** provides advanced SAN multipathing with centralized management.
- **Hitachi Device Manager - Storage Navigator** provides VSP element storage management functions.
- **Hitachi Storage Advisor** minimizes technology exposure, streamlines storage service-level administration, and simplifies storage management operations.
- **Hitachi Command Suite** provides a single point of management for all Hitachi physical and virtualized storage, and it acts as the interface for integration with other management systems.

SVOS for NAS

SVOS for NAS is specifically designed for the VSP Gx00 and Fx00 models with embedded NAS modules, for NFS, SMB, and iSCSI protocols. SVOS for NAS includes native file deduplication, snapshots, two enterprise virtual NAS server licenses, NDMP, virtual server security, anti-virus, read caching, and a tiered file system for efficient unified storage management.

The NAS server can support data on an external server using Hitachi Universal Volume Manager (UVM). UVM presents storage on external storage arrays to the server as if the storage is local. To subsequently migrate data from the external storage onto the local storage, the server also supports Hitachi Tiered Storage Manager (HTSM). Using UVM instead of Universal Migrator enables the NAS server to preserve snapshots, quotas, and ACLs. UVM also has the ability to replicate a whole span (storage pool) in a single operation.

Key capabilities of the SVOS for NAS software include:

- **Continuous availability:**
 - Zero RTO and RPO for sites in case of a node, storage, or site failure.
 - Flexibility for environments and sites up to 100 km.
- **Support for VMware VVOL:**
 - Increases storage efficiency through VM-centric storage allocation.
 - Automated provisioning of VMs delivers quicker adjustment to business changes through Hitachi policy-driven management.
 - Support for mapping individual VMs to virtual machine disks (VMDKs) delivers increased granularity and resource utilization rates.
- **Multitenancy:**
 - Enables independent enterprise virtual servers (EVSs).
 - Supports hosting multiple assignments on one Hitachi NAS Platform on the same IP address; delivers true separation.
- **Superior capacity efficiency:**

- Support for 1-PB file system.
- Primary storage deduplication to eliminate copies of redundant data.
- Support for FMD-based compression for the VSP F series.
- **Intelligent file tiering:**
 - Policy-based hierarchical storage management feature spans Hitachi NAS Platform and Hitachi Content Platform.
- **Enhanced high availability:**
 - Active-active clustering with cluster read caching for scalable, read-intensive NFS workload, incremental block replication (IBR), Hitachi NAS Replication high-speed replication, and synchronous disaster recovery service.
 - Optimized file system pre-mount checks and improves NVRAM replay time for faster cluster failover.
 - Nondisruptive cluster upgrades to remove updates and reduce downtime.
- **Virtualization services:**
 - VMware vStorage APIs for Array Integration (VAAI) adapter divests storage operations from VMware vSphere to Hitachi NAS.
 - Virtual volumes, virtual servers, and cluster namespace unify the directory structure while simplifying storage capacity management tasks.
 - Optional Hitachi Virtual Infrastructure Integrator simplifies backup, restore, and cloning operation from VMware vSphere to Hitachi NAS.
- **Data management services:**
 - Centralized GUI management, pointer-based snapshots, Hitachi NAS replication, writable snapshots, quick file restore, hard and soft quotas (volume, group, or user), NAS data migrator feature, scalable file systems, storage pools, policy-based management, and transparent data migration and relocation.
- **Protocols supported:** Hitachi NAS can support various protocols, including:
 - Internet Content Adaptation Protocol support for virus scanning.
 - IPv6 support: Connect using an IPv6 address or a host name resolving to an IPv6 address through the external system management unit (SMU) software or SMU command line interface (CLI).
- **Complete network protocol support:**
 - Server Message Block (SMB) 1.0, 2.0, 3.0; SMB 3.0 encryption support; Network File System (NFS) with UDP v2 and v3 or TCP v2, v3 and v4; NDMP v2, v3 and v4; File Transfer Protocol (FTP); Secure File Transfer Protocol (sFTP); File Transfer Protocol, Secure (FTPs); iSCSI. SMB2.1 signing, and SMB secure negotiation are supported.
- **Management and other protocols:**
 - HTTP, SSL, SSH, SNMP v3, NIS, DNS, WINS, NTP, and email alerts.

Capacity saving function: deduplication and compression

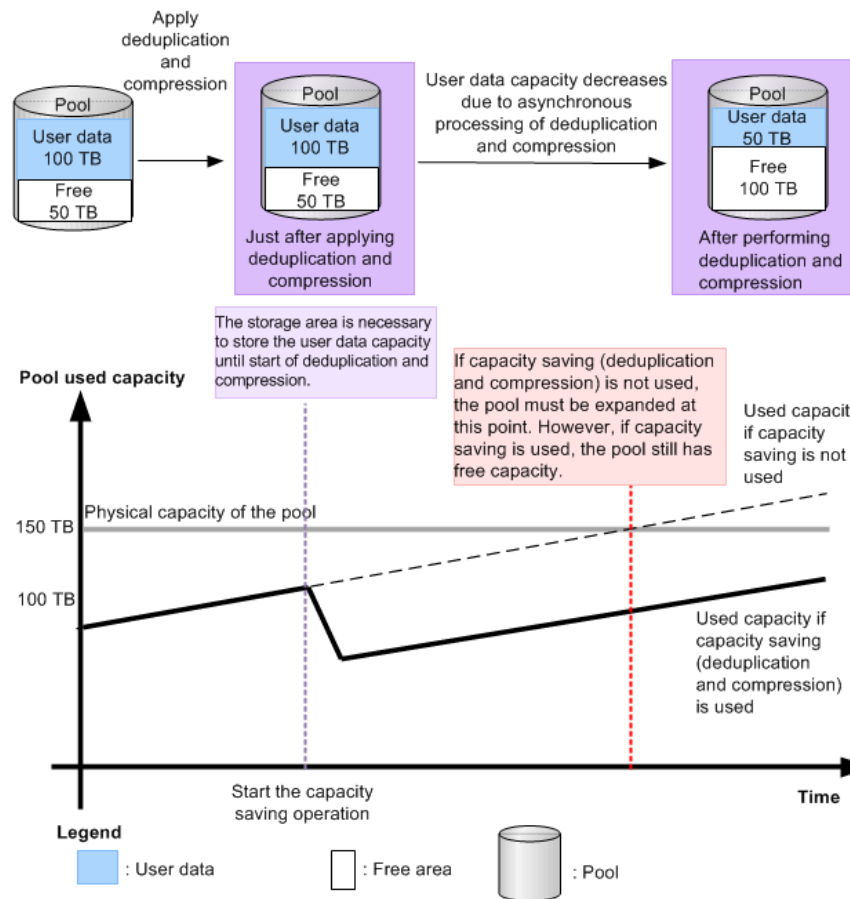
The capacity saving function includes data deduplication and data compression. When the capacity saving function is in use, the controller of the storage system performs deduplication and compression to reduce the size of data to be stored, thereby reducing your bitcost for the stored data. Capacity saving is available for use on internal flash drives only.

- **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 you enable deduplication on a pool, the deduplication system data volume (DSD volume) for that pool is created. The deduplication system data volume is used exclusively by the storage system to manage the deduplication function. A search table in the deduplication system data 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 can be enabled per DP-VOL.


The compression and deduplication processing is performed asynchronously for pages that store data, and the free area of the pool can be increased, thereby reducing the cost of purchasing drives over time.



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.

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.



Accelerated compression is enabled at the parity-group level on FMD DC2 flash module drives. You can enable or disable accelerated compression from launch points in the management software to Hitachi Device Manager - Storage Navigator or by using command control interface (CCI). Accelerated compression is supported on Hitachi Virtual Storage Platform G series and F series with SVOS 6.4 or later.

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*.

Dynamic tiering and active flash

Hitachi Dynamic Tiering (HDT) is based on Hitachi Dynamic Provisioning. It further simplifies tiered storage management by automating fine-grained, page-based movement of data to the most appropriate storage media according to workload and usage patterns. It automates management, maximizes service levels, and minimizes storage costs.

Dynamic Tiering places the host volume's data across multiple tiers of storage. There can be up to three tiers (high-, medium-, and low-speed layers). Dynamic Tiering determines tier usage based on data access levels. It automatically allocates pages with high I/O load to the upper tier, which

contains higher speed drives, and pages with low I/O load to the lower tier, which contains lower speed drives.

HDT tracks page usage within each tier over a configured time period. At the end of each time period, pages that have been accessed very frequently might be moved from a slower media to high-performance flash media.

Active flash is a feature enhancement to HDT that monitors a page's access frequency level and promotes pages that suddenly became busy in real time rather than waiting for a specific time period to elapse.

Dynamic Tiering and active flash offer the following benefits:

Reduced storage costs

- Reduces media costs and drive counts through self-optimized use of storage tiers
- Achieves space efficiency through thin provisioning
- Eliminates manual data classification
- Eliminates manual data movement between tiers
- Reduces operational overhead
- Reduces space, power, and cooling requirements

Improved performance

- Optimizes data placement automatically for performance using an I/O rate-based heat index
- Gives SSD-class performance to information stored largely on less expensive tiers by automatically moving the most accessed data to the highest (SSD) tier
- Supports the highest efficiency and throughput through granular page-based data movement
- Uses wide striping across the entire pool
- For applications that have high I/O to particular pages, active flash improves HDT performance by automatically moving the most frequently accessed data in real time.

Efficient administration

- Simplifies management of up to three storage tiers as a single volume
- Automatically moves the most active data to the highest performing tier
- Automatically adjusts to dynamic workloads and capacity requirements
- Moves pages up and down for optimal placement
- Significantly reduces administration time

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 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).

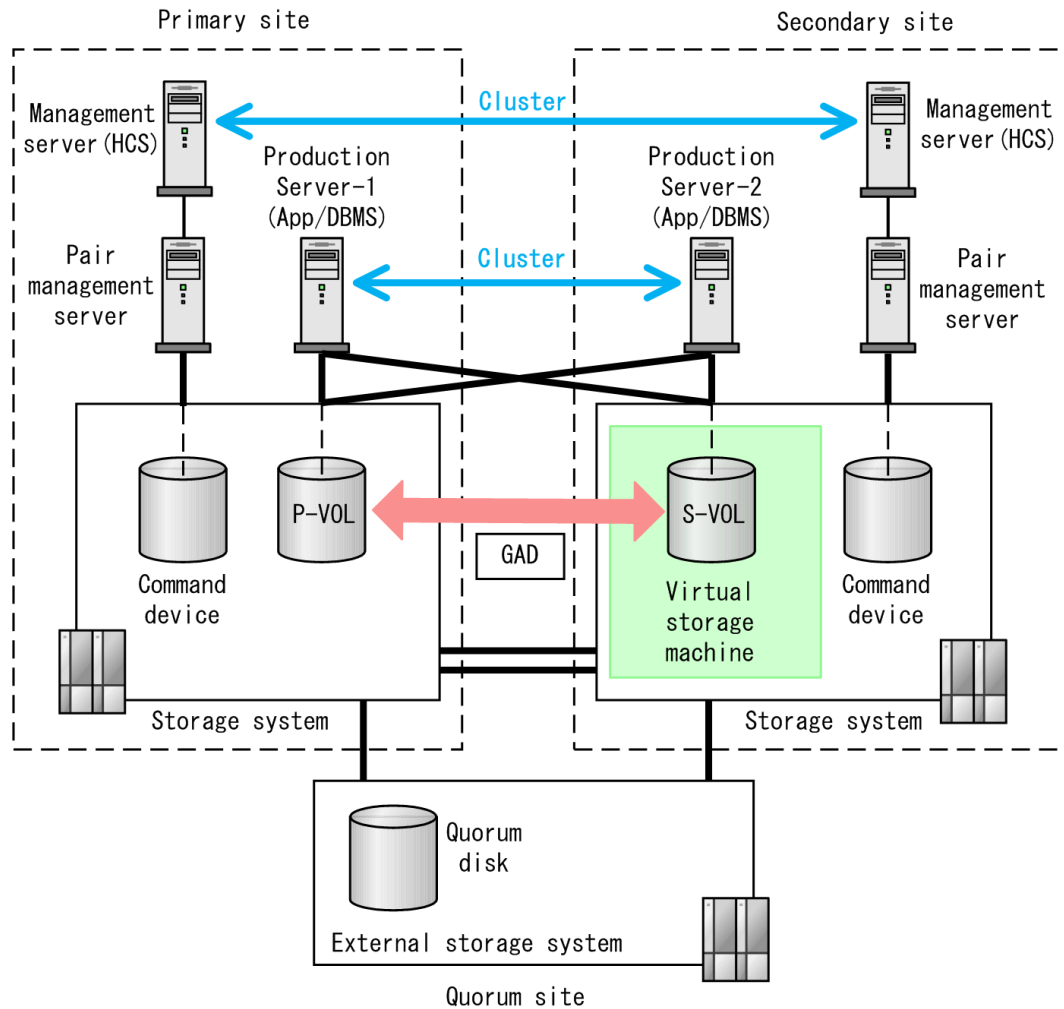
Establishing a global-active device pair has the following benefits:

- 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*

GAD enhanced for NAS

GAD Enhanced for NAS takes advantage of the GAD feature to cluster two VSP Gx00 or Fx00 systems with NAS modules across two sites. This synchronous disaster recovery configuration, also referred to as a stretched cluster, creates a four-node cluster stretched across two sites within 100 km of each other.

For more information about this special configuration, contact your Hitachi Data Systems representative.

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 Gx00 models
Hitachi Virtual Storage Platform	VSP G1000, VSP G1500, and VSP F1500
Hitachi Virtual Storage Platform	VSP Gx00 models
Hitachi Unified Storage VM	VSP G1000, VSP G1500, and VSP F1500
Hitachi Unified Storage VM	VSP Gx00 models

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 Data Systems 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*.

Centralized management

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 Gx00 and Fx00 models includes:

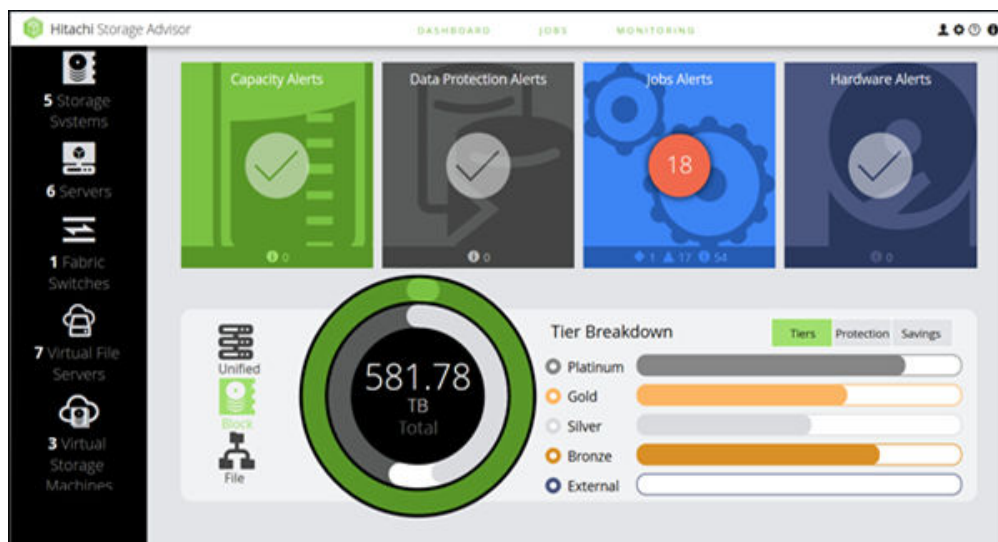
- Hitachi Storage Advisor (HSA)

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

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.

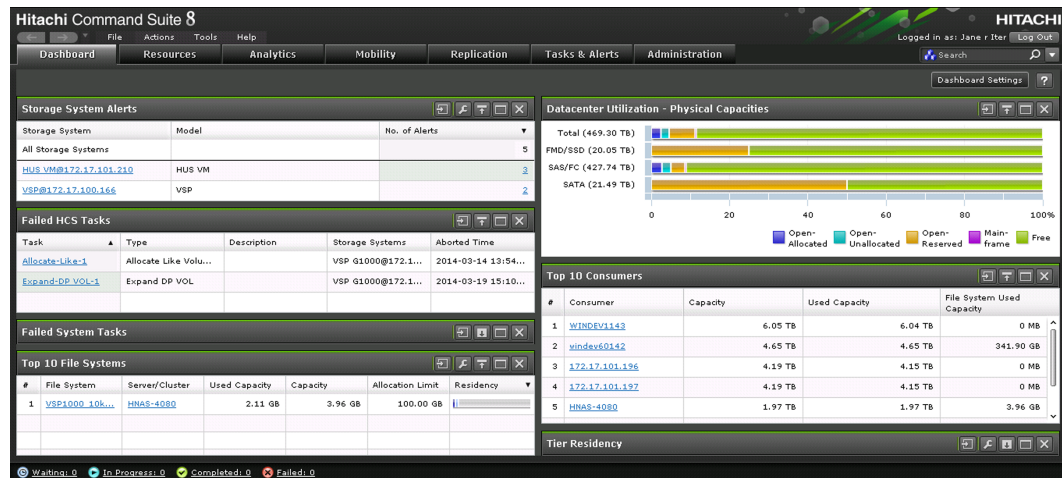
HCS comprises the following optional components, each of which is licensed separately:

- **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 the end-to-end storage provisioning process 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.

The following figure shows the main HCS dashboard.



Hitachi Command Suite offers the following benefits:

Central inventory management to properly manage growth

- Common administrative framework consolidates asset management across all virtualized storage resources for operational efficiency to increase storage return on investment
- Common management console to discover, configure, monitor, and report on all tiers and virtualized storage resources
- Dashboard highlights system-wide capacity usage, top consumers, and system alerts
- Logical group constructs to easily align storage resources with business applications
- Integrated management framework enables automation, mobility, service-level management, and data protection

Simplified storage provisioning for rapid deployment

The common management framework consolidates storage provisioning for both structured and unstructured data:

- Centrally configure storage pools for block, file, and object consumers
- Centrally manage data security, mobility, performance, and replication
- Simplified provisioning with contextual workflows
- Reduce operational expenses; manage more with less effort

Maximize business application performance

- Automatically align with business applications, define tiers, and set policies by application workload for maximum performance

- Automate optimal data placement to increase storage utilization by up to 50%

Automated lower cost and less management

- Automatically move inactive data to lower-cost storage
- Automatically move active data to highest-performing tier
- Define tiers and set policies to optimize cost

Ensure performance is running at peak efficiency with Hitachi Tuning Manager

Tuning Manager provides comprehensive storage system health monitoring and troubleshooting to deliver the operational efficiencies required to optimize shared Hitachi storage resources.

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 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)

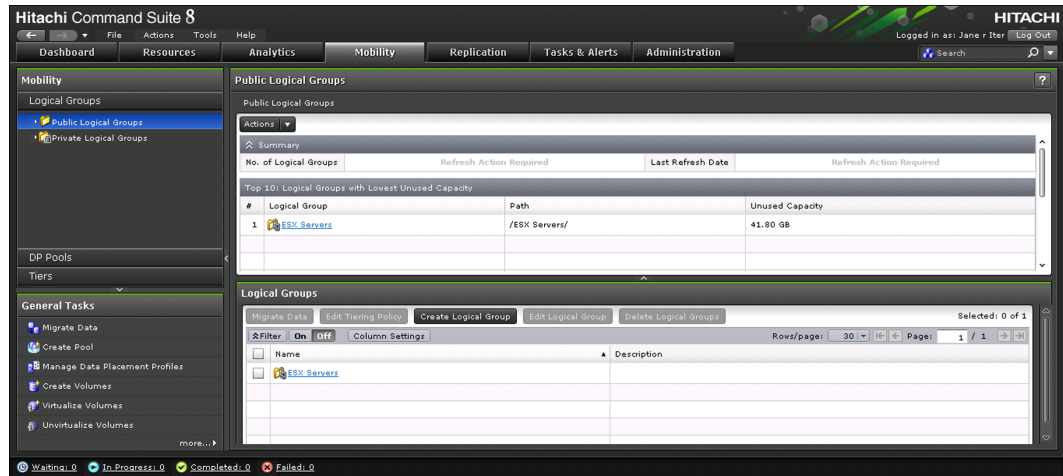
When using Universal Replicator, you can check copy performance visually and perform root cause analysis using the Replication tab of the Hitachi Command Suite GUI.

Data mobility

Hitachi Tiered Storage Manager provides the unique ability to migrate volumes between heterogeneous tiers of storage, without affecting

application access to data. The software supports up to 64 simultaneous migrations, using 64 processor threads, and it offers a single interface for all data movement between storage virtualized by the storage system.

The following figure shows the Mobility tab of HCS, which is used to access Hitachi Tiered Storage Manager.



Tiered Storage Manager offers the following benefits:

- Matches application price, performance, and availability needs to storage attributes
- Controls the automated behavior of Dynamic Tiering through the use of standard and custom policies and profiles. You can also proactively create and pool different classes of storage for maximum efficiency and long-term performance
- Manages storage resources according to the needs of specific business applications, while supporting the ability to migrate data nondisruptively

Hitachi Infrastructure Analytics Advisor

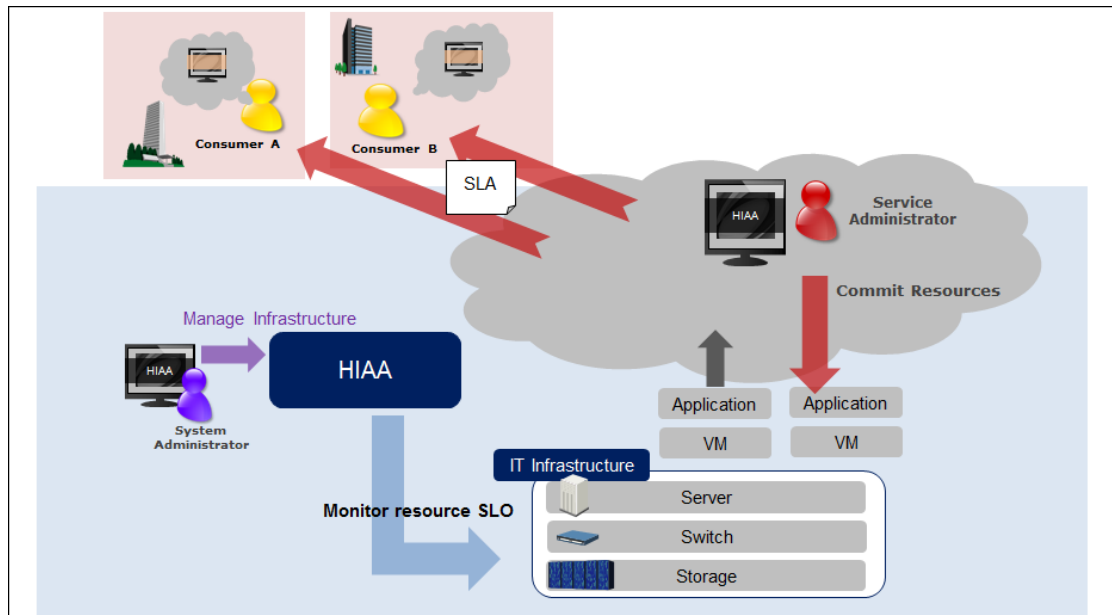
Hitachi Infrastructure Analytics Advisor is a data center management software for monitoring, reporting, and analyzing end-to-end performance from hosts through shared storage resources.

Product overview

With Infrastructure Analytics Advisor, 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.

Using Infrastructure Analytics Advisor, 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 describes how the Infrastructure Analytics Advisor ensures the performance of your storage environment based on real-time service level objectives (SLOs).



The system administrator uses Hitachi Infrastructure Analytics Advisor (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.

Infrastructure Analytics Advisor 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 Infrastructure Analytics Advisor to assign resources, such as VMs and storage capacity from registered storage systems, to consumer applications. The purpose of doing this is to manage critical SLO violations and to ensure that service performance meets the service level agreements.

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 Data Instance Director

Hitachi Data Instance Director (HDID) provides business-defined data protection, which simplifies the creation and management of complex, business-defined policies to meet service level objectives for availability.

HDID supports the Hitachi VSP G series storage systems, offering an orchestration layer for remote replication supporting Hitachi TrueCopy® and Hitachi Universal Replicator, local and remote snapshots and clones with Hitachi Thin Image and Hitachi ShadowImage®, continuous data protection, and incremental backup.

HDID provides the following benefits:

Operational recovery

HDID offers two approaches to meeting operational recovery requirements, depending on whether the data being protected is stored on Hitachi storage.

- *Storage-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 Hitachi 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® Exchange and Microsoft SQL Server® on Microsoft Windows® platforms, for Oracle database environments on Linux, AIX, and Solaris, and for SAP HANA environments. HDID is integrated with Hitachi Virtual Infrastructure Integrator (V2I) to provide storage-based protection of VMware vSphere® environments. Other applications can also be integrated using the simple scripting interface.

These 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 secondary purposes such as test and development, or backup to a target device such as a purpose-built backup appliance (PBBA) or tape. HDID administrators can also view and restore storage-based snapshots created in VMware environments by Hitachi Virtual Infrastructure Integrator.

- *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 Exchange and SQL Server. Batch mode backup is supported on Windows, Linux and IBM® AIX® systems.

Disaster recovery

HDID provides storage-based and software based choices for restoring operations at, or from, another location following a site level outage.

- *Storage-based disaster recovery*

HDID configures and automates Hitachi TrueCopy synchronous remote replication software and Hitachi Universal Replicator software 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.

- *Host-based disaster recovery*

The backup data stored locally by HDID can be asynchronously replicated, on a scheduled basis, to another location. It does not require specific storage for either the primary or disaster recovery copy.

Long-term retention

With HDID, moving Microsoft Exchange and Windows file data to Hitachi Content Platform (HCP) for archiving enables your administrators to reduce the amount of data in their production systems and meet corporate and regulatory data retention requirements.

Leave the archived file on the source system, delete it or leave a stub file as a pointer. HDID archives files as individual objects that can be easily viewed, retrieved or audited with standard HCP tools. No special software is needed to unpack or decode the archived files.

Unified management

One of the many benefits of Hitachi Data Instance Director is its single-footprint platform. It enables you to layer, combine and orchestrate backup, CDP, snapshots, replication and archive to achieve the specific service levels of data recovery and retention each application requires.

The unique graphical user interface (GUI) incorporates a powerful policy builder that resembles laying out business processes on a whiteboard. Easily create and change policies as needed, visualize data protection processes, and align them with 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 + incremental methods.
- To further reduce downtime, bare metal recovery images can be created using standard backup processes. The operating system volume and application volumes can be recovered in a single operation.
- HDID supports a range of storage repositories, including block, file, object, Microsoft Azure and tape storage.
- HDID scales seamlessly to manage hundreds of terabytes of data.

Hitachi Automation Director

Hitachi Automation Director is a software solution that provides tools to automate and simplify the end-to-end storage provisioning process 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 Automation Director 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 Automation Director 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.

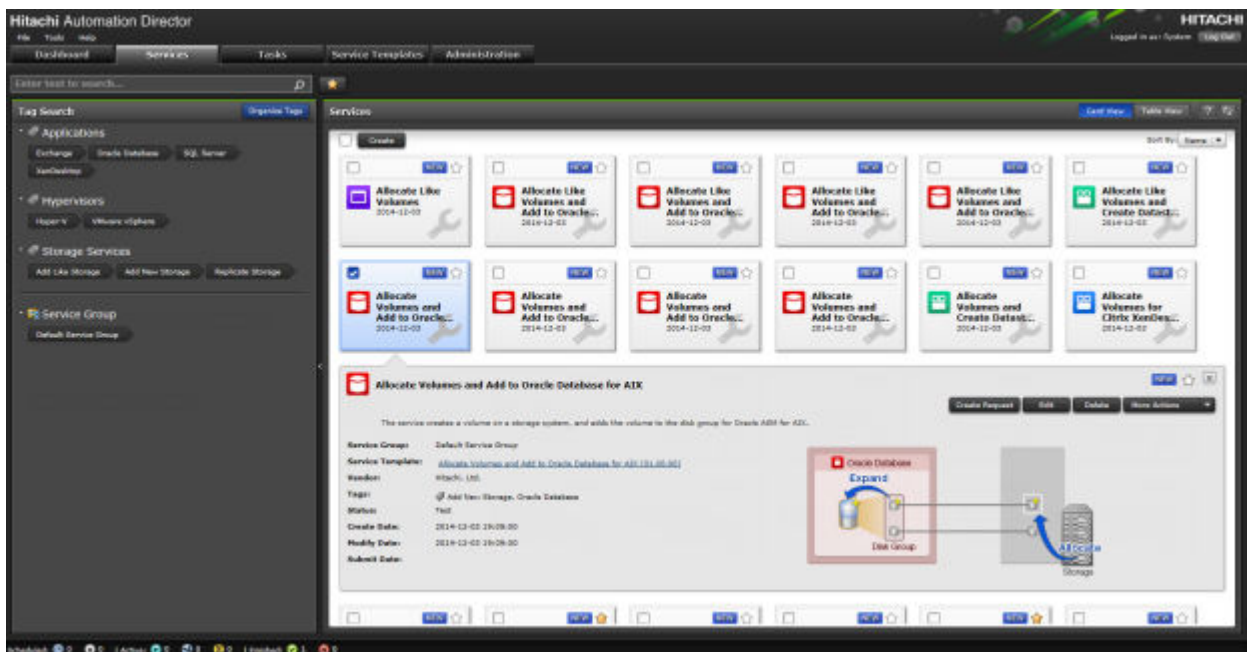


Figure 1 Select a service on the Services tab to review details and create a request for provisioning.

Hitachi Automation Director 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 pre-defined service templates, by using the Service Builder tool, to address an organization's changing needs.
- The REST API facilitates integration of Automation Director with relevant IT automation processes.

Software management 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](#)
- [Ensuring optimal storage performance and business application service levels with analytics](#)
- [Maximizing business application performance and availability with data mobility](#)
- [Delivering storage infrastructure as a service through automated workflows](#)
- [Data protection for business-critical Oracle databases](#)
- [End-to-end performance troubleshooting using Infrastructure Analytics Advisor](#)
- [Flexible reporting and analysis using Data Center Analytics](#)

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

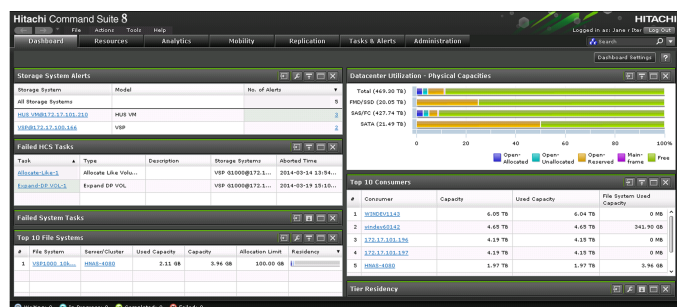

Today, financial institutions provide a wide array of services to their customers. These services must support both structured data (online and ATM transactions, such as withdrawing or depositing checks and cash) and unstructured data (such as email messages, SMS text messages, customer feedback, bank 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.

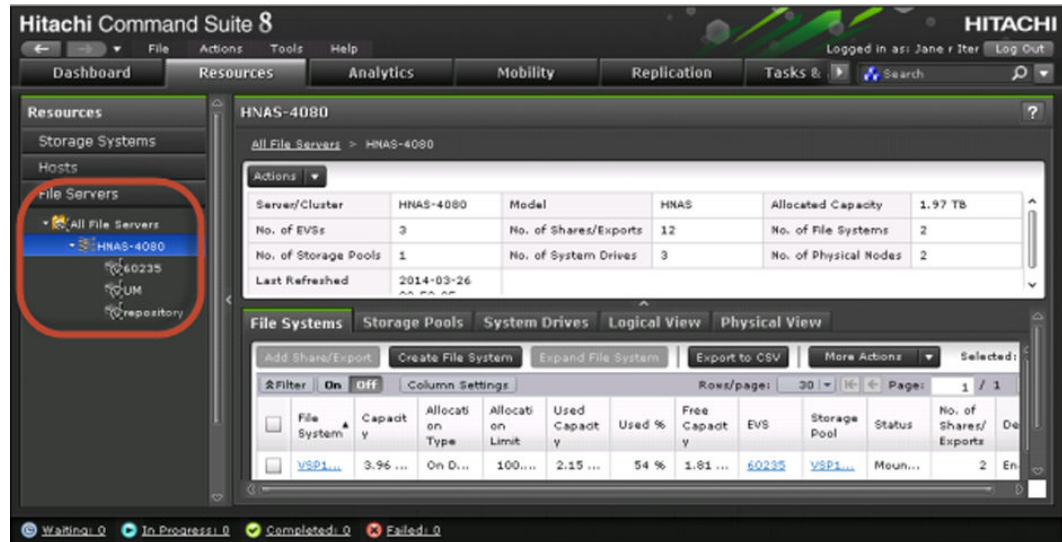
Overall, financial institutions require a platform with the breadth and flexibility to provide services wherever, whenever, and however customers need them.

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:

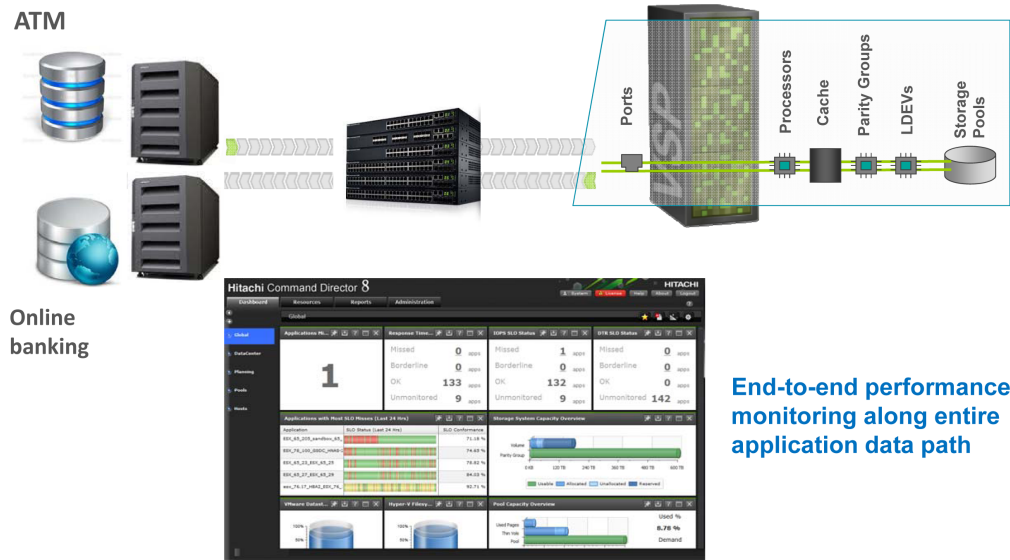
- Lack of performance baselines or benchmarks to analyze response time for online banking and ATM applications
- Insufficient root cause analysis (RCA) techniques that look deep into application performance problems, and ineffective existing techniques
- Absence of real-time monitoring capability and 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.



End-to-end performance monitoring along entire application data path

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

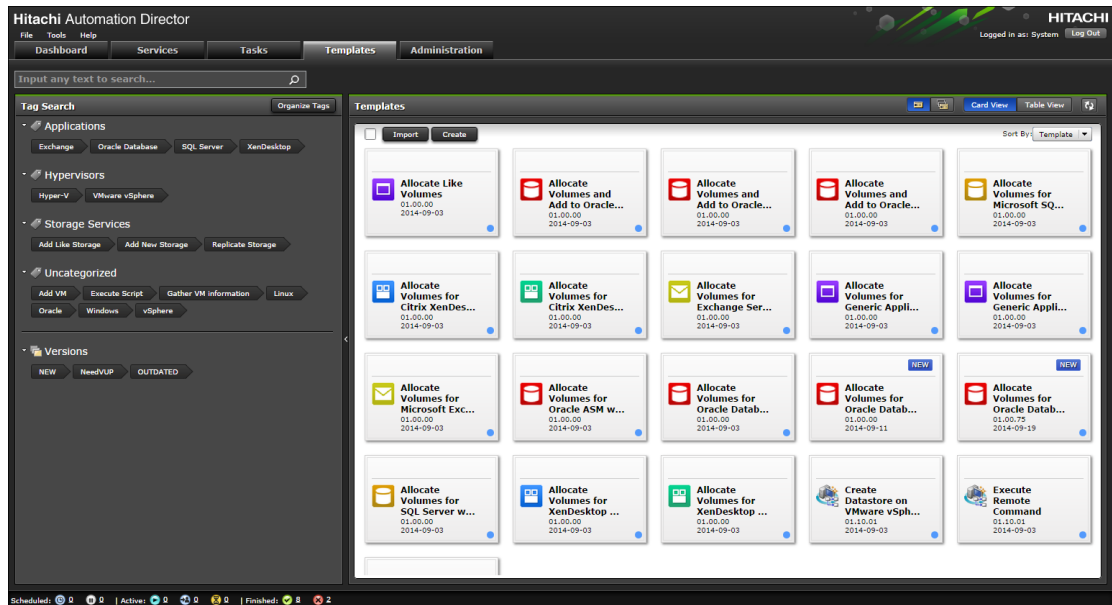
- Implements intelligent automation workflows to streamline the storage provisioning process.
- Provides a catalog of predefined service templates and plugin components that incorporate Hitachi best practices in storage provisioning and that minimize human error.
- Provides customizable storage service 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

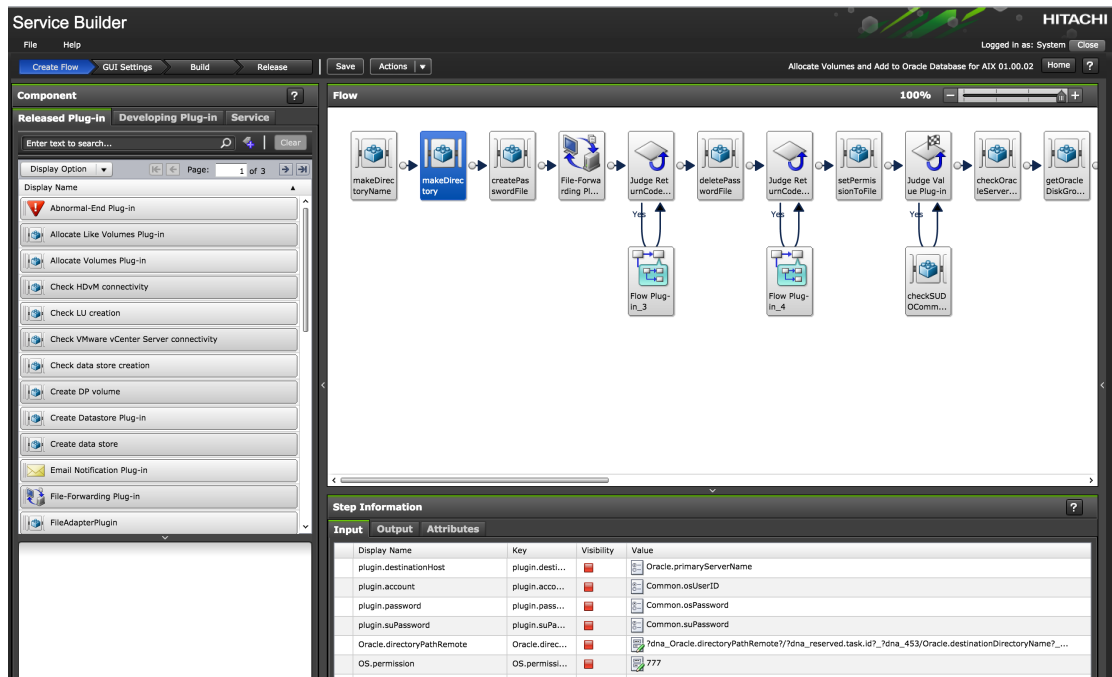
Hitachi Automation Director offers a web-based portal and includes a catalog of predefined workflows that are based on best practices for various applications. These workflows take into account infrastructure requirements for specific applications, including the appropriate storage tier. Capturing the 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 Automation Director intelligent engine matches the request with the appropriate infrastructure

based on performance and capacity analysis. Hitachi Automation Director 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.



Hitachi Automation Director 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.



Hitachi Automation Director supports all native block storage systems and 3rd-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:

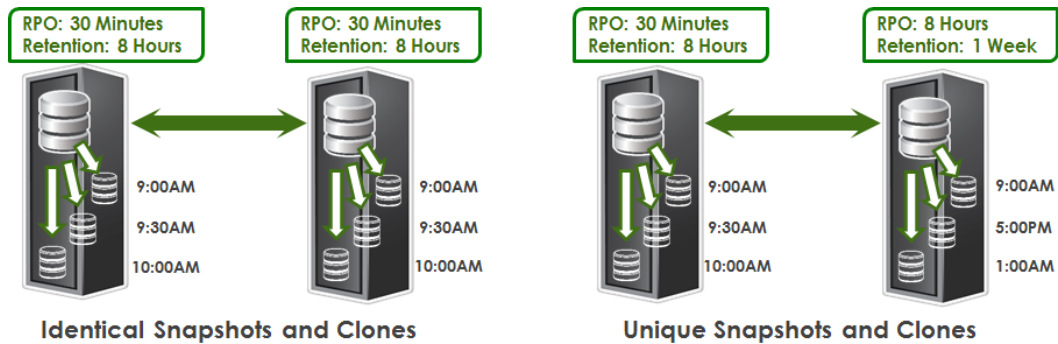
- *Service Level Agreement (SLA)-driven Policy* enables administrators to define the data classification (such as SQL Server or Oracle), data protection operations, and required SLAs (RPO, data retention).
- *Whiteboard-style Data Flow* enables the administrator to define the copy destinations and assign 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

Method	Description
	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.



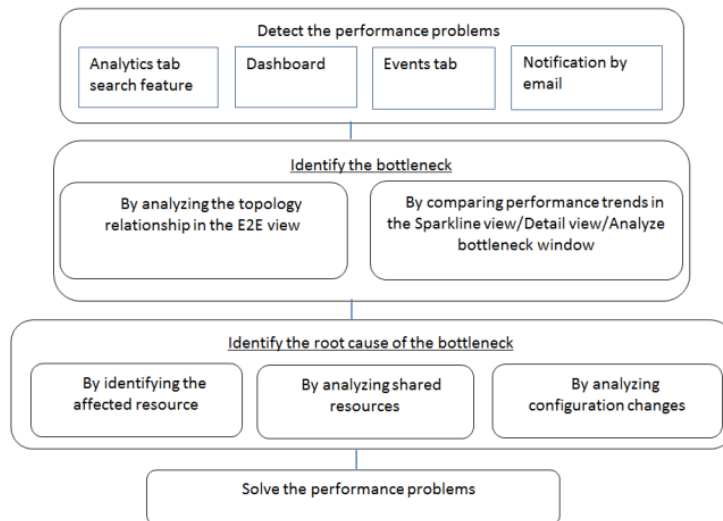
End-to-end performance troubleshooting using Infrastructure Analytics Advisor

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

This section describes an example of troubleshooting slow response times of a business critical application using Infrastructure Analytics Advisor.

The most significant metric to watch out for while monitoring the online transactions is the I/O rate. The application will be able to process large number of transactions when the I/O rates are higher. To maintain good response times in an OLTP environment which mostly generates random access I/O, the read I/O response times should be higher. For response time centric applications, such as OLTP applications, you must maintain low utilization values to ensure CPU availability and low Q-depth values to ensure no wait time.

The use case flow for troubleshooting high response times for an OLTP application using advanced analytics and troubleshooting features of Infrastructure Analytics Advisor is described as follows:



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 bottleneck on the server, network, or storage components. In the following example, the volumes, 00:00:03, 00:00:05, and 00:00:06 have been identified as the resources causing performance problems.

You can identify and analyze the component causing the bottleneck in any of the following views:

- E2E view
- Analyze bottleneck > Verify Bottleneck tab
- Sparkline view
- Detail view

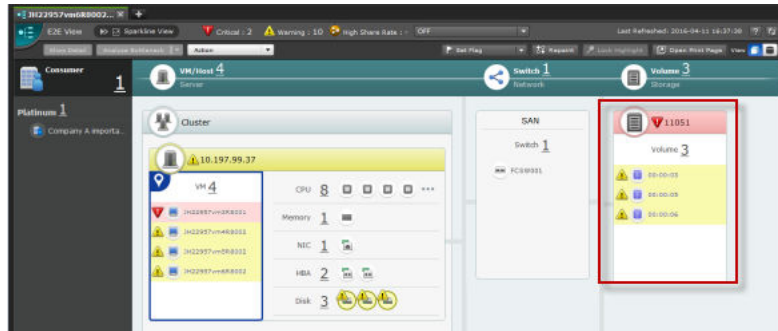
For details, see the *Infrastructure Analytics Advisor User Guide*.

- **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.

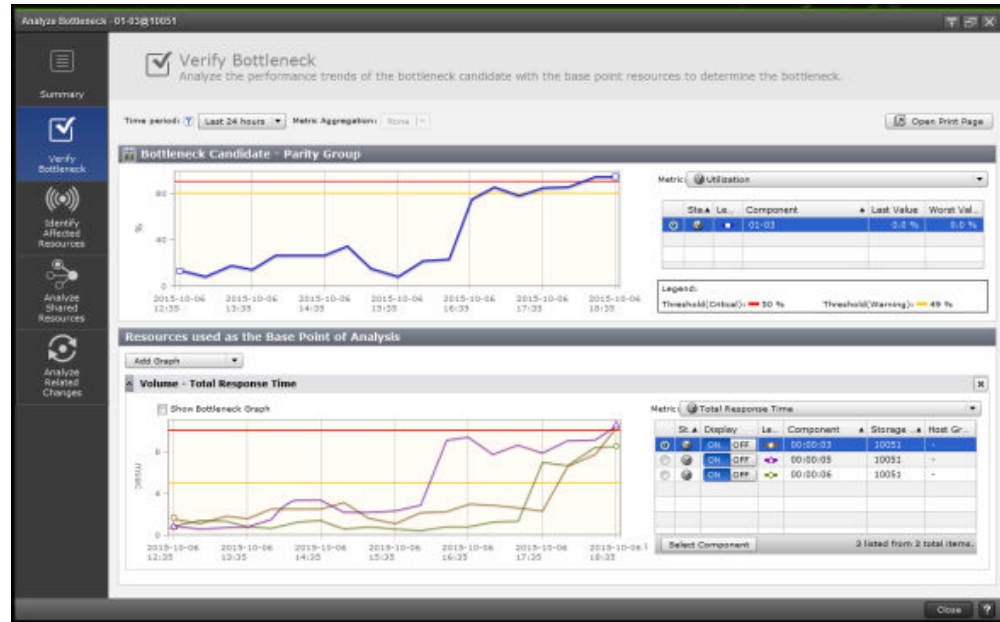
The following figure shows the E2E configuration related to the affected volumes, 00:00:03, 00:00:05, and 00:00:06:



You can change the base point of analysis to narrow down the topology associated with the affected volumes. Select the affected volume, right-click, and then select Change Base Point. The parity group is identified as the component causing the performance bottleneck.



- **Verify bottleneck window:** In the E2E view, right-click the parity group resource icon and then select Verify Bottleneck to launch the Verify Bottleneck window. In the Verify Bottleneck window, you can analyze the performance trends of the potential bottleneck candidate with the base point resources. If the performance charts display similar trend patterns in the same time period, you can assume that the selected resource is the bottleneck candidate.

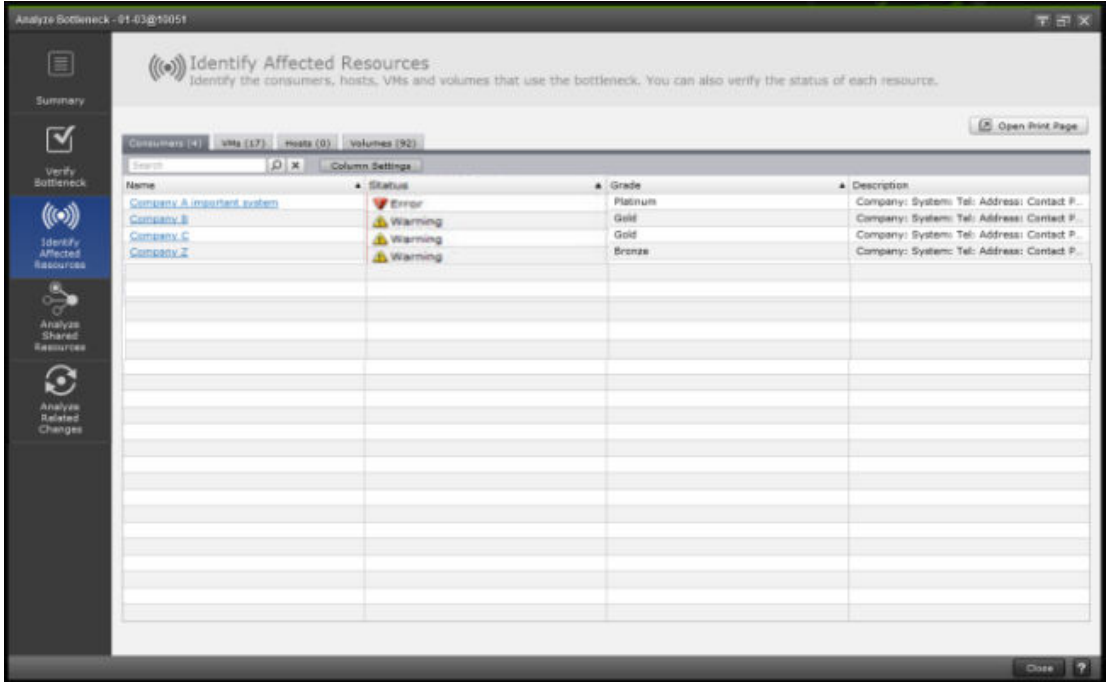


Analyze root cause of the bottleneck

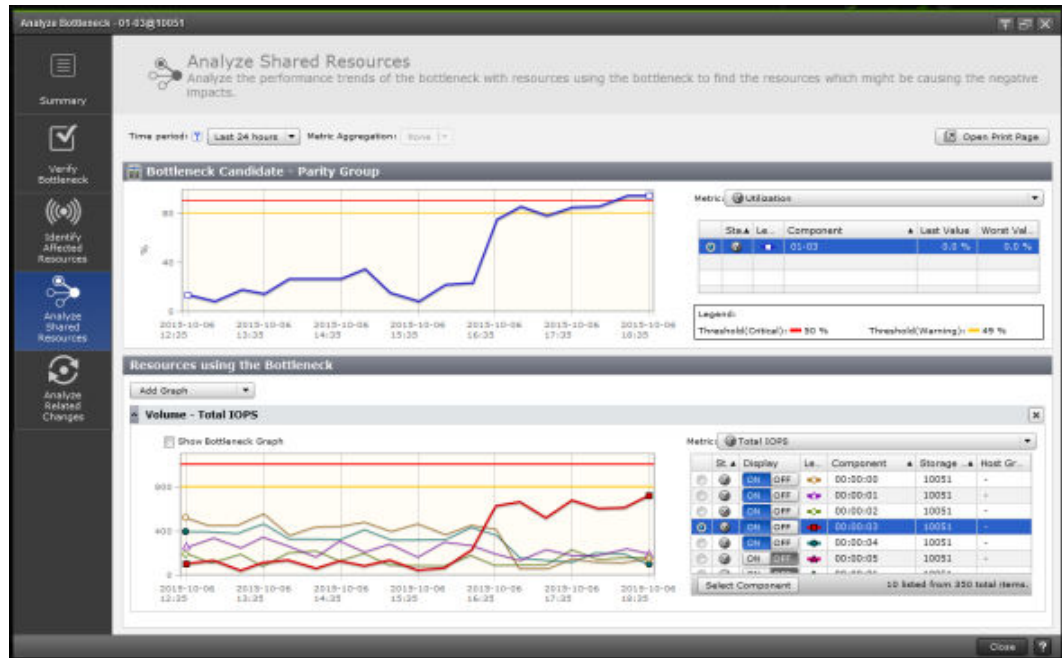
Infrastructure Analytics Advisor integrated troubleshooting aids provide guidance about how to find the root-cause of the performance problems. The root cause can be due to the resource contention issues in the shared infrastructure, or due to configuration changes in the environment.

- **Identify affected resources:** In the Analyze Bottleneck window, click the Identify affected resources tab. In this window, 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.

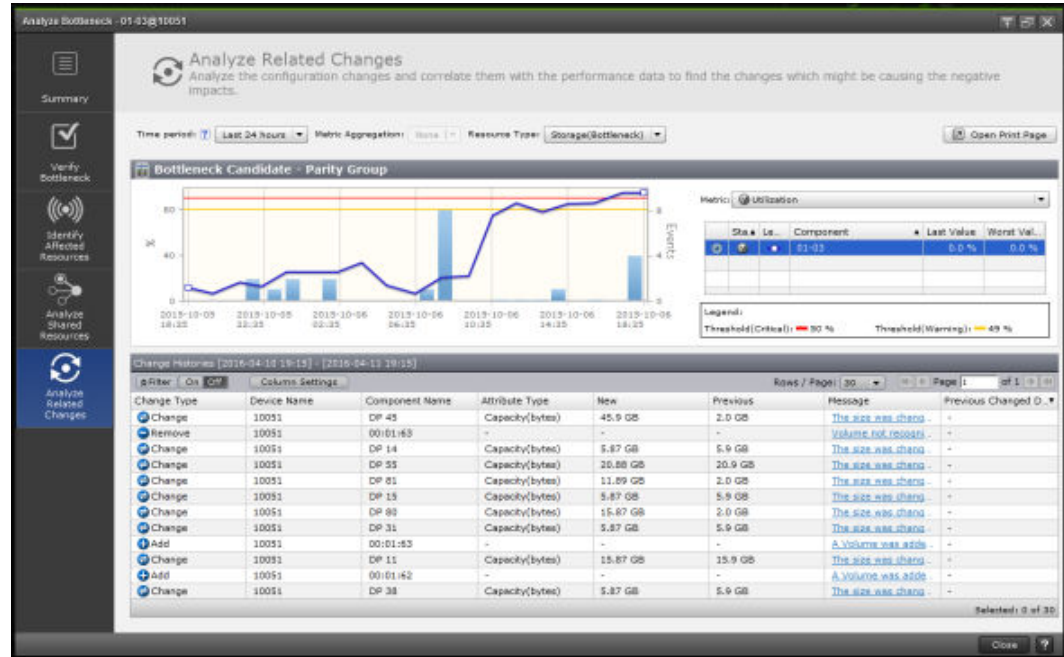


- Analyze shared resources:** In the Analyze Shared Resources window, compare the performance trends of the parity group with the Volumes and VMs that use the parity group. The performance trends of the parity group closely match with the trend patterns of one of the VMs, which leads to the confirmation that the VM is the resource in the shared infrastructure which is over utilizing the parity group. You can resolve the bottleneck caused by the shared resources by adopting efficient load balancing methodologies, which enables optimal utilization of the resources in the shared infrastructure.



- Analyze related changes:** The configuration changes can sometimes be the source of the performance problem in your environment. Infrastructure Analytics Advisor 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 the systems performance and behaviour. In the Analyze Related Changes window, a combination chart that contains two data series is displayed, the bars represent the change events and the line represents the performance of the bottleneck candidate. You can correlate the performance data of the parity group and the change events that occurred in the specified time period to determine the effects of the configuration changes. Based on the analysis you can confirm that there

were no configuration change events that caused the performance degradation in the parity group.



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 your 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:

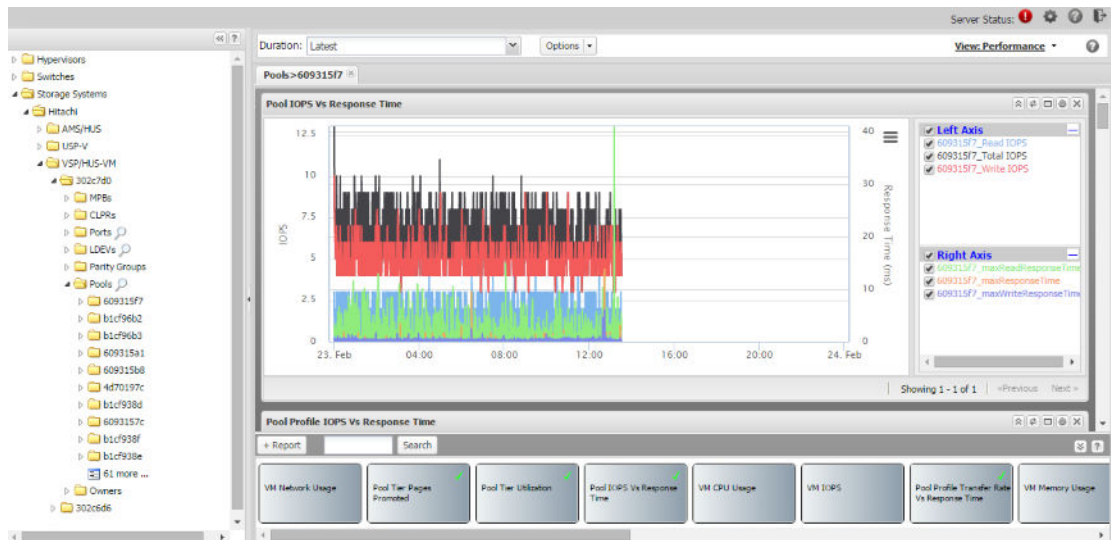
- **Predefined reports:** provide high-level details at the application level and also a granular report that shows component-level performance data.
- **Ad-hoc reports:** enable you to combine related and unrelated metrics of any monitored resource in one report to review the overall performance impact.

- **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.



User documentation for VSP Gx00 and VSP Fx00

The following tables list the user documentation for Hitachi Virtual Storage Platform G200, G400, G600, G800, Hitachi Virtual Storage Platform F400, F600, F800, Hitachi Storage Advisor (HSA), Hitachi Infrastructure Analytics Advisor (HIAA), optional Hitachi Command Suite components, and Hitachi NAS.

Table 1 Manuals for Hitachi VSP Gx00 and Fx00

Subject of guide	Documents
Release notes	<p><i>Hitachi Virtual Storage Platform Gx00 and Fx00 Release Notes</i>, RN-8304xx-Mxxx</p> <p><i>Command Control Interface Release Notes</i>, RN-90RD7194</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 v2.0 Release Notes</i>, RN-96HIAA000</p>
Hardware reference	<p><i>Hitachi Virtual Storage Platform G200 Hardware Reference Guide</i>, MK-94HM8020</p> <p><i>Hitachi Virtual Storage Platform G400, G600 Hardware Reference Guide</i>, MK-94HM8022</p> <p><i>Hitachi Virtual Storage Platform G800 Hardware Reference Guide</i>, MK-94HM8026</p> <p><i>Hitachi Virtual Storage Platform F400, F600 Hardware Reference Guide</i>, MK-94HM8045</p> <p><i>Hitachi Virtual Storage Platform F800 Hardware Reference Guide</i>, MK-94HM8046</p> <p><i>Service Processor Technical Reference for VSP Gx00 and VSP Fx00 Models</i>, FE-94HM8036</p>
Management software installation and configuration	<p><i>System Administrator Guide</i>, MK-94HM8016</p> <p><i>Command Control Interface Installation and Configuration Guide</i>, MK-90RD7008</p> <p><i>Hitachi Command Suite Installation and Configuration Guide</i>, MK-90HC173</p>

Subject of guide	Documents
	<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 Storage Advisor Getting Started Guide</i>, MK-94HSA001</p> <p><i>Hitachi Infrastructure Analytics Advisor Installation and Configuration Guide</i>, MK-96HIAA002</p>
External storage	<p><i>Hitachi Universal Volume Manager User Guide</i>, MK-92RD8024</p>
Host configuration	<p><i>Open-Systems Host Attachment Guide</i>, MK-90RD7037</p> <p><i>Provisioning Guide</i>, MK-94HM8014</p> <p><i>Hitachi Storage Provider for VMware vCenter (Block VASA Provider) Deployment Guide</i>, MK-90ADPTR010</p> <p><i>Hitachi Storage Provider for VMware vCenter (VASA Provider) Deployment Guide</i>, MK-92ADPTR104</p>
Provisioning and mobility	<p><i>Provisioning Guide</i>, MK-94HM8014</p> <ul style="list-style-type: none"> • Accelerated compression • Active flash • Capacity saving (deduplication and compression) • Data Retention Utility • Dynamic Provisioning • Dynamic Tiering • Global storage virtualization • LUN Manager • Resource Partition Manager • Virtual LUN <p><i>Hitachi Storage Advisor User Guide</i>, MK-94HSA004</p> <p><i>Hitachi Command Suite User Guide</i>, MK-90HC172</p>
Data protection	<p><i>Hitachi Thin Image User Guide</i>, MK-92RD8011</p> <p><i>Hitachi ShadowImage® User Guide</i>, MK-92RD8021</p> <p><i>Hitachi TrueCopy® User Guide</i>, MK-92RD8019</p> <p><i>Hitachi Universal Replicator User Guide</i>, MK-92RD8023</p> <p><i>Global-Active Device User Guide</i>, MK-92RD8072</p> <p><i>Hitachi Data Instance Director User Guide</i>, MK-93HDID001</p> <p><i>Hitachi Data Instance Director Quick Start Guide</i>, MK-93HDID002</p>
Analytics and performance	<p><i>Performance Guide</i>, MK-94HM8012</p> <ul style="list-style-type: none"> • Performance Monitor • Server Priority Manager • Virtual Partition Manager <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>

Subject of guide	Documents
	<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>
Security	<p><i>Hitachi Command Suite User Guide</i>, MK-90HC172</p> <p><i>Provisioning Guide (Data Retention Utility)</i>, MK-94HM8014</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-94HM8015</p> <p><i>Hitachi Command Suite Audit Log Reference Guide</i>, MK-92HC213</p> <p><i>Hitachi Audit Log User Guide</i>, MK-94HM8028</p>
Troubleshooting	<p><i>Hitachi Command Suite Messages</i>, MK-90HC178</p> <p><i>Hitachi Device Manager - Storage Navigator Messages</i>, MK-94HM8017</p> <p><i>Hitachi SNMP Agent User Guide</i>, MK-94HM8015</p>
Command line interface	<p><i>Command Control Interface User and Reference Guide</i>, MK-90RD7010</p> <p><i>Command Control Interface Command Reference</i>, MK-90RD7009</p>

Table 2 Command Suite optional component manuals

Subject of guide	Documents
Software installation	<p><i>Hitachi Command Suite Tuning Manager Installation Guide</i>, MK-96HC141</p> <p><i>Hitachi Command Suite Compute Systems Manager Installation and Configuration Guide</i>, MK-91HC195</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>
Management software setup	<p><i>Hitachi Command Suite Replication Manager Configuration Guide</i>, MK-90HC175</p> <p><i>Hitachi Command Suite Compute Systems Manager Installation and Configuration Guide</i>, MK-91HC195</p> <p><i>Hitachi Command Suite Tuning Manager Agent Administration Guide</i>, MK-92HC013</p>

Subject of guide	Documents
SAN multipathing	<p><i>Hitachi Command Suite Compute Systems Manager User Guide, MK-91HC194</i></p> <p><i>Hitachi Command Suite Dynamic Link Manager (for AIX®) User Guide, MK-92DLM111</i></p> <p><i>Hitachi Command Suite Dynamic Link Manager (for Linux®) User Guide, MK-92DLM113</i></p> <p><i>Hitachi Command Suite Dynamic Link Manager (for Solaris) User Guide, MK-92DLM114</i></p> <p><i>Hitachi Command Suite Dynamic Link Manager (for Windows®) User Guide, MK-92DLM129</i></p> <p><i>Hitachi Command Suite Dynamic Link Manager (for VMware®) User Guide, MK-92DLM130</i></p> <p><i>Hitachi Command Suite Global Link Manager Installation and Configuration Guide, MK-95HC107</i></p> <p><i>Hitachi Command Suite Global Link Manager Messages, MK-95HC108</i></p>
Data protection	<p><i>Hitachi Command Suite Replication Manager User Guide, MK-99HC166</i></p>
Analytics	<p><i>Hitachi Command Suite Tuning Manager Server Administration Guide, MK-92HC021</i></p> <p><i>Hitachi Command Suite Tuning Manager Agent Administration Guide, MK-92HC013</i></p> <p><i>Hitachi Command Suite Tuning Manager User Guide, MK-92HC022</i></p> <p><i>Hitachi Command Suite Tuning Manager Hardware Reports Reference, MK-95HC111</i></p> <p><i>Hitachi Command Suite Tuning Manager Operating System Reports Reference, MK-95HC112</i></p> <p><i>Hitachi Command Suite Tuning Manager Application Reports Reference, MK-95HC113</i></p> <p><i>Hitachi Command Suite Tuning Manager CLI Reference Guide, MK-95HC119</i></p> <p><i>Hitachi Command Suite Tuning Manager Getting Started Guide, MK-95HC120</i></p>
Automation	<p><i>Hitachi Command Suite Automation Director User Guide, MK-92HC205</i></p> <p><i>Hitachi Command Suite Automation Director Service Builder User Guide, MK-92HC222</i></p>
Troubleshooting	<p><i>Hitachi Command Suite Tuning Manager Messages, MK-95HC114</i></p> <p><i>Hitachi Command Suite Compute Systems Manager Messages, MK-91HC197</i></p> <p><i>Hitachi Command Suite Automation Director Messages, MK-92HC221</i></p>
Command line interface	<p><i>Hitachi Command Control Interface Installation and Configuration Guide, MK-90RD7008</i></p> <p><i>Hitachi Command Suite CLI Reference Guide, MK-90HC176</i></p> <p><i>Hitachi Command Suite Tiered Storage Manager CLI Reference Guide, MK-90HC177</i></p>

Subject of guide	Documents
	<i>Hitachi Command Suite Tuning Manager CLI Reference Guide</i> , MK-95HC119 <i>Hitachi Command Suite Compute Systems Manager CLI Reference Guide</i> , MK-91HC196
API	<i>Hitachi Command Suite Tuning Manager API Reference Guide</i> , MK-92HC218 <i>Hitachi Command Suite Automation Director API Reference Guide</i> , MK-92HC217

Table 3 Hitachi NAS manuals

Subject of guide	Documents
System installation and configuration	<i>Network Administration Guide</i> , MK-92HNAS008 <i>Virtual SMU Administration Guide</i> , MK-92HNAS074
System management	<i>Storage Systems User Administration Guide</i> , MK-92HNAS013 <i>System Access Guide</i> , MK-92HNAS014 <i>Server and Cluster Administration Guide</i> , MK-92HNAS010 <i>Storage Subsystem Administration Guide</i> , MK-92HNAS012
Data management	<i>Data Migrator Administration Guide</i> , MK-92HNAS005 <i>File Services Administration Guide</i> , MK-92HNAS006
Data protection	<i>Replication and Disaster Recovery Administration Guide</i> , MK-92HNAS009 <i>Backup Administration Guide</i> , MK-92HNAS007 <i>Snapshot Administration Guide</i> , MK-92HNAS011 <i>Antivirus Administration Guide</i> , MK-92HNAS004

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas

+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa

+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific

+852 3189 7900
hds.marketing.apac@hds.com

Contact Us

www.hds.com/en-us/contact.html



MK-94HM8013-04

March 2017