Hitachi Virtual Storage Platform VSP G/F350, G/F370, G/F700, G/F900

88-01-0x

Product Overview

MK-97HM85011-00
May 2018

Contents

Preface

This guide provides an overview of the Hitachi Virtual Storage Platform G/F350, G/F370, G/F700, and G/F900 storage systems, including hardware components, general system specifications, software features, management interfaces, user documentation, and software management examples.

Product version

This document revision applies to the following product versions:

- VSP G/F350, G/F370, G/F700, G/F900: firmware version 88-01-0*x* or later
- Hitachi Storage Virtualization Operating System RF 8.1 or later
- Hitachi Storage Advisor: v2.3 or later
- Hitachi Command Suite: v8.6 or later
- Hitachi Infrastructure Analytics Advisor: v4.0 or later
- Hitachi Data Instance Director: v6.5 or later

Accessing product documentation

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

Getting help

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

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

Comments

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

Thank you!

Chapter 1: Introducing Hitachi Virtual Storage Platform G/F350, G/F370, G/F700, G/F900

Based on Hitachi's industry-leading storage technology, the Hitachi Virtual Storage Platform G350, G370, G700, G900 and the all-flash Hitachi Virtual Storage Platform F350, F370, F700, and F900 include a range of versatile, high-performance storage systems that deliver flash-accelerated scalability, simplified management, and advanced data protection.

Product summary

Key features

Key features of the Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 systems include:

- Up to 2.4M IOPS performance
- 100% data-availability guarantee
- AI optimized operations
- Cloud optimization
- Enhanced integration for VMware®, Windows®, and Oracle® environments
- Advanced active-active clustering, replication, and snapshots
- Active flash tiering and groundbreaking flash modules

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.

Storage Virtualization Operating System RF (SVOS RF)

Hitachi Storage Virtualization Operating System RF (SVOS RF), which includes an all new enhanced software stack, 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-compatibility for easy integration into IT environments, storage virtualization, and management capabilities provide the utmost agility and control, helping you build infrastructures that are continuously available, automated, and agile.

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

- External storage virtualization
- Thin provisioning and automated tiering
- Flash performance acceleration
- Deduplication and compression of data stored on internal flash drives
- Storage service-level controls
- Data-at-rest encryption
- Performance instrumentation across multiple storage platforms
- Centralized storage management
  - Simplified: Hitachi Storage Advisor
  - Advanced and powerful: Hitachi Command Suite, Command Control Interface
  - For organizations that have their own management toolset, we include standards-based application program interfaces (REST APIs) that centralize administrative operations on a preferred management application.

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 G350, G370, G700, G900 and VSP F350, F370, F700, F900, HAF with SVOS RF 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 better performance and response time.
- **Second- and third-generation flash modules:** The FMD DC2 and FMD HD 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 multilayer 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, these drives enable better performance than the SSDs.

Nondisruptive migration

Hitachi offers nondisruptive migration capabilities in two options. You can choose the service that Hitachi Vantara 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 the GSS migration center of excellence team is intended for more complex, large-scale, heterogeneous, and replicated data center environments.

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 G350, G370, G700, G900 and VSP F350, F370, F700, F900 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 VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900. Storage Provider enables policy-based storage management using VMware Storage Policy-based Management (SPBM) and VMware Virtual Volumes (VVols). 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 G350, G370, G700, G900 and VSP F350, F370, F700, F900 support VAMP to provide enhanced control of I/O path selection and failover.
- **vStorage API for Data Protection (VADP):** Hitachi VSP G350, G370, G700, G900 and VSP F350, F370, F700, F900 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 Vantara 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.

Chapter 2: Management architecture

System administration of the Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 storage systems involves initial setup and configuration of the storage system as well as ongoing configuration and administration activities, such as managing users, software licenses, and configuring the audit logging. The storage systems support several management tools for system administration, including software that is embedded on the controller of the storage system, software that is factory-installed on the (optional) service processor (SVP) of the storage system, and management software that is installed on a server.

System management architecture

The following figure illustrates the storage system management software architecture. It shows the access points that a system administrator can use to access tools that allow you to configure and manage the storage system.

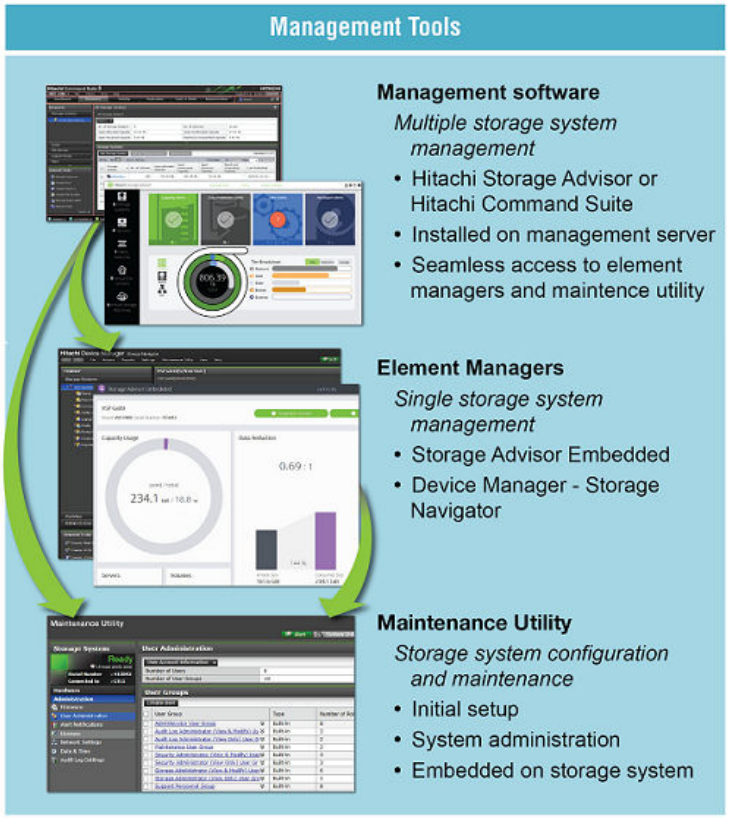

Chapter 3: About the hardware

The Hitachi Virtual Storage Platform G350, G370, G700, G900 and the all-flash Hitachi Virtual Storage Platform F350, F370, F700, and F900 are rack-mountable storage systems that incorporate state-of-the-art virtualization and advanced data-management and fault-tolerant technologies.

Hitachi Virtual Storage Platform G/F350, G/F370, G/F700, G/F900 hardware overview

The architecture of the Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 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 Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 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 Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 models allow a defective controller to be replaced.

VSP G350, G370, G700, G900 models support a variety of drives, including HDDs, SSDs for the G350 and G370. The G700 and G900 provide additional support for FMD HD drives including the second-generation Hitachi Accelerated Flash (HAF) flash module drives (FMD DC2) that are the foundation of the VSP F350, F370, F700, F900 models. The FMD DC2 drives deliver improved response times, effective flash capacity, and return on investment. The VSP G700 and VSP G900 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 F350, F370, F700, F900

The VSP F350, F370, F700, F900 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 all-flash arrays offers up to 2.4 million IOPS to meet the most demanding application requirements. The all-flash arrays 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 all-flash arrays with optional synchronous and asynchronous replication available for complete data protection. The all-flash arrays range in storage capacity from 1.4 PB (raw) up to 8.7 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.

Optional service processor (SVP)

Hitachi Virtual Storage Platform G350, G370, G700, G900 and Virtual Storage Platform F350, F370, F700, F900 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 Vantara or as a software application:

- The physical SVP is a 1U management server that runs Windows Embedded Standard 10.
- The SVP software application is installed on a customer-supplied server and runs on a customer-supplied version of Windows.

Hitachi VSP G350, G370, G700, G900 specifications

The following table provides the key specifications for the VSP G350, G370, G700, G900.

Feature	VSP G350	VSP G370	VSP G700	VSP G900
Maximum capacity (raw)	1.4 PB (SSD) 2.4 PB (HDD)	2.1 PB (SSD) 3.6 PB (HDD)	6.1 PB (FMD) 6.5 PB (SSD) 11.7 PB (HDD)	8.1 PB (FMD) 8.7 PB (SSD) 14.1 PB (HDD)

Feature	VSP G350	VSP G370	VSP G700	VSP G900
Solid-state drive options (raw, SFF)	7.6 TB, 3.8 TB, 1.9 TB, 480 GB			7.6 TB, 3.8 TB, 1.9 TB
Flash module drive options	-- --		14 TB, 7 TB, 3.5 TB	
SFF hard disk drive options	10K RPM: 600 GB, 1.2 TB, 2.4 TB			
LFF hard disk drive options	7.2K RPM: 6 TB, 10 TB			
Maximum number of drives	264	384	1,200	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 (without drives)	16 FC: 16 and 32 Gbps 8 iSCSI: 10 Gbps		64 FC: 16 and 32 Gbps 32 iSCSI: 10 Gbps	80 FC: 16 and 32 Gbps 40 iSCSI: 10 Gbps
Maximum cache	128 GB	256 GB	512 GB	1,024 GB
RAID levels	RAID-1+0, RAID-5, RAID-6			
Protocols	NFS, SMB, FTP, iSCSI, and HTTP to the cloud			

Hitachi VSP F350, F370, F700, F900 specifications

The following table provides the key specifications for the VSP F350, F370, F700, F900 models.

Feature	VSP F350	VSP F370	VSP F700	VSP F900
Maximum capacity (raw)	1.4 PB (SSD)	2.1 PB (SSD)	6.1 PB (FMD) 6.5 PB (SSD)	8.1 PB (FMD) 8.7 PB (SSD)

Feature	VSP F350	VSP F370	VSP F700	VSP F900
Flash packs	4 x 7.6 TB SSD 4 x 3.8 TB SSD 4 x 1.9 TB SSD 4 x 480 GB SSD		4 x 14 TB FMD 4 x 7 TB FMD 4 x 3.5 TB FMD 4 x 7.6 TB SSD 4 x 3.8 TB SSD 4 x 1.9 TB SSD 4 x 480 GB SSD	4 x 14 TB FMD 4 x 7 TB FMD 4 x 3.5 TB FMD 4 x 7.6 TB SSD 4 x 3.8 TB SSD 4 x 1.9 TB SSD
Maximum number of drives	192	288	432 FMD 864 SSD	576 FMD 1,152 SSD
Disk expansion trays	2U: 24 SFF (2.5")		2U: 24 SFF (2.5"); 12 FMD	
Block module height	2U		4U	
Host interfaces (without drives)	16 FC: 16 and 32 Gbps 8 iSCSI: 10 Gbps		48 FC: 16 and 32 Gbps 24 iSCSI: 10 Gbps	64 FC: 16 and 32 Gbps 32 iSCSI: 10 Gbps
Maximum cache	128 GB	256 GB	512 GB	1,024 GB
RAID levels	RAID-1+0, RAID-5, RAID-6			
Protocols	NFS, SMB, FTP, iSCSI, and HTTP to the cloud			

Chapter 4: Software components and features

Hitachi VSP G/F350, G/F370, G/F700, G/F900 storage systems are powered by Hitachi Storage Virtualization Operating System RF (SVOS RF) and supported by Hitachi storage management software to enable you to effectively manage and centralize your software-defined infrastructure.

Hitachi Storage Virtualization Operating System RF (SVOS RF)

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

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

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

Function	Description
Cache partitioning	Virtual Partition Manager supports up to 32 cache partitions.
Multipathing and failover	Dynamic Link Manager Advanced provides advanced SAN multipathing with centralized management.
Performance monitoring	Performance Monitor provides an intuitive, graphical interface to assist with performance configuration planning, workload balancing, and analyzing and optimizing storage system performance.
Storage system-based utilities	Storage system-based utilities include LUN manager, customized volume size, Data Retention Utility, quality of service controls, audit log, and volume shredder feature.
Standard management interface support	Management interface support includes SMI-S provider, SNMP agent, and REST.
Optimized storage for virtualized server infrastructure	A wide range of plugins and adapters are available to enhance virtual server infrastructure performance and administrator productivity. SVOS RF features integration of VMware applications (including VAAI, VASA, VAMP, VADP, SRA, VVOL) and Microsoft Windows applications (including VSS, ODX).

SVOS RF data reduction

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

Data reduction capabilities include:

- Pattern detection and removal

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

- Accelerated compression

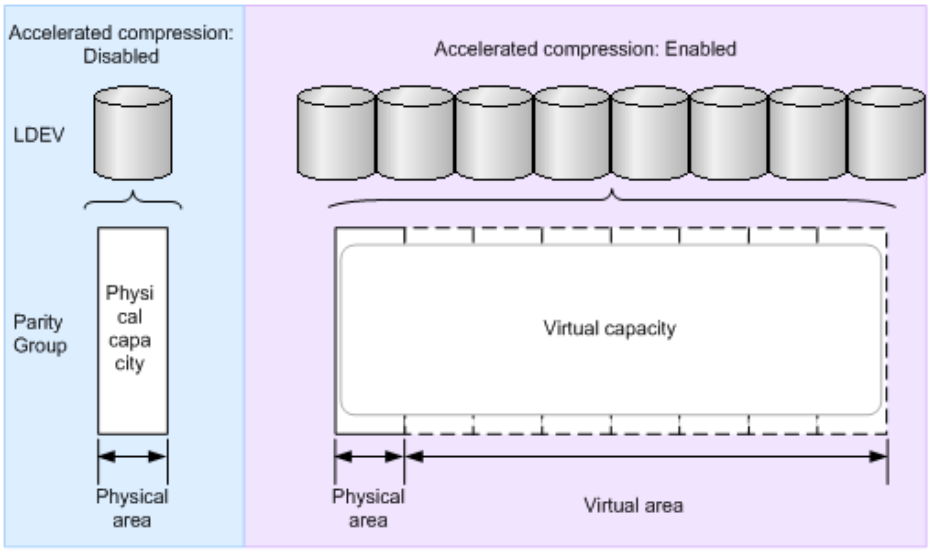

  The accelerated compression feature of Dynamic Provisioning delivers a data compression capability that enables you to realize more virtual capacity in a parity group than the actual usable capacity, providing improved storage optimization. You can enable accelerated compression at the parity-group level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

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

- Capacity saving

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

  - Deduplication

    The data deduplication function deletes duplicate copies of data written to different addresses in the same pool. The deduplication function is enabled on the desired DP-VOLs in the pool. When deduplication is enabled, data that has multiple copies between DP-VOLs assigned to that pool is removed.

  - Compression

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

Management of SVOS RF systems

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

- Hitachi Storage Advisor (HSA)

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

- Hitachi Command Suite (HCS)

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

- Command Control Interface (CCI)

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

- Hitachi Storage Advisor Embedded (HSAE)

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

Overview of Hitachi Storage Advisor

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

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

Some of the key Storage Advisor capabilities include:

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

Hitachi Command Suite

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

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

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

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

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

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

Advanced SAN multipathing

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

Dynamic Link Manager offers the following benefits:

Business continuity

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

Productivity and process

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

Command Control Interface

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

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

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

Single-system management and storage monitoring using Hitachi Storage Advisor Embedded

With Hitachi Storage Advisor Embedded (HSAE), storage configuration and management operations are easy. This is made possible by a simple and intuitive graphical user interface, high-performing REST API, and CLI embedded natively in the storage system. HSAE allows you to quickly configure and provision a single storage system and its resources. These resources include; volumes, pools and ports, where you can perform simple operations using any of the embedded interfaces (GUI, CLI, and REST API).

In addition to configuring and provisioning, HSAE also provides a dashboard that enables a quick and intuitive way to observe the health of your storage system. These include both component health alerts and key performance metrics reporting such as IOPS, response time, and data response speed. This performance metrics reporting capability also extends to a volume level capability.

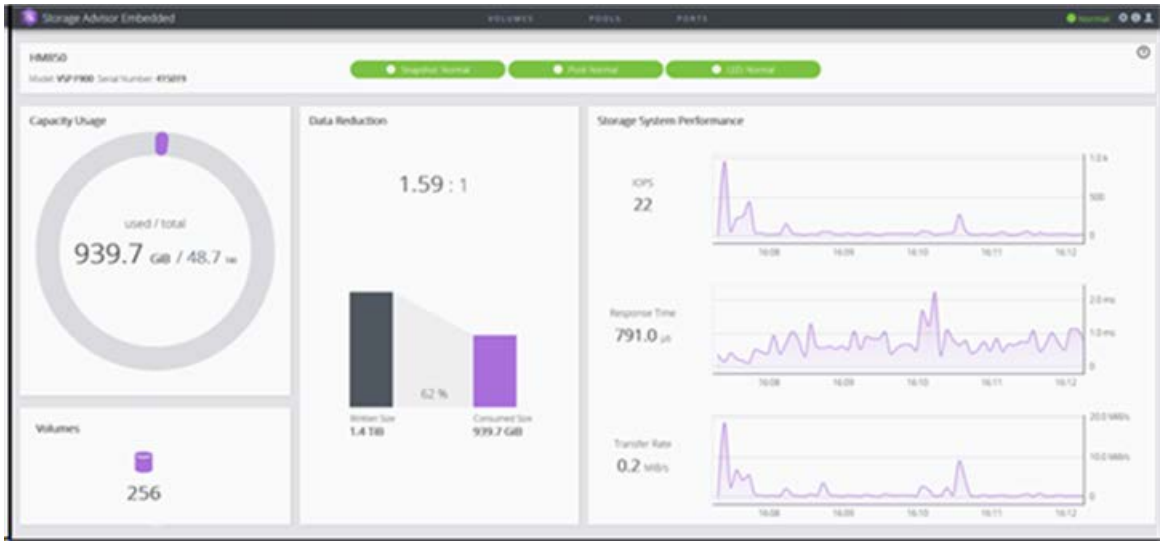

The Dashboard also displays the overall capacity of the storage system, its usage/available capacity and data reduction savings, including the amount of data reduction achieved by using the capacity saving feature (deduplication and compression).

Advanced global storage virtualization and software bundles

Optional SVOS RF features include best-in-class local and remote replication technologies as well as active-active metro clustering to provide rapid recovery from system and site-

level events that could disrupt access to data. SVOS RF business continuity solutions are designed for maximum flexibility, enabling organizations to build a recovery strategy that spans multiple data centers and delivers to their specific SLAs.

Optional software products and packages for SVOS RF systems include:

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

Hitachi Data Mobility software

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

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

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

High availability with global-active device

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

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

- Continuous I/O

  If a primary volume becomes unavailable, the host continues to transparently access the secondary volume.

- Clustered failover

  You do not need to perform storage system tasks such as suspension or resynchronization of a global-active device pair due to a host failure.

- Virtual machine integration

  If a virtual machine is creating a high load at one site, you can move the load to the other site, eliminating the need for data migration.

How global-active device works

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

The following figure shows an example GAD configuration.

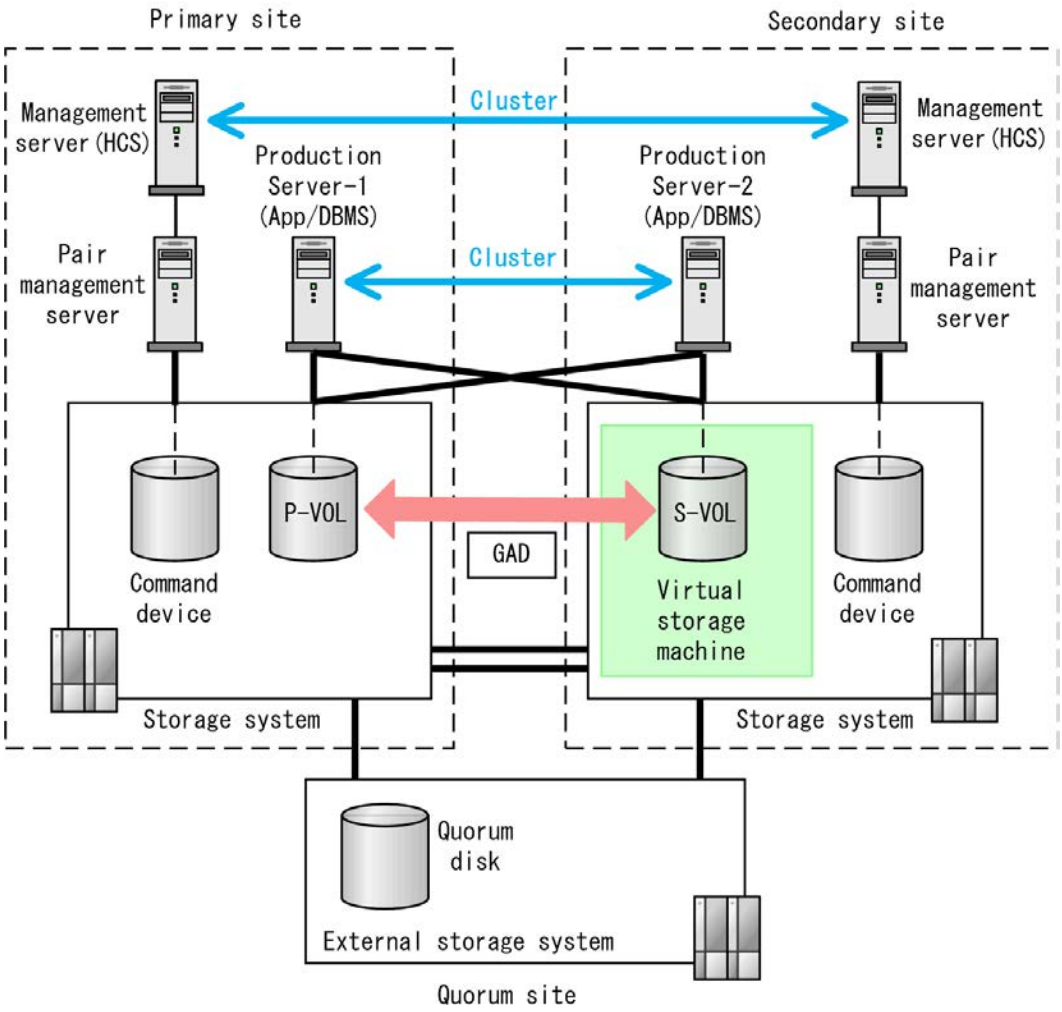

GAD pair volumes are monitored by a quorum disk (preferably located at a third site). The quorum disk acts as a heartbeat for the GAD pair, with the primary and secondary storage systems accessing the quorum disk periodically to check on the other storage system. In the event of a communication or hardware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

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

Global-active device storage system configurations

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

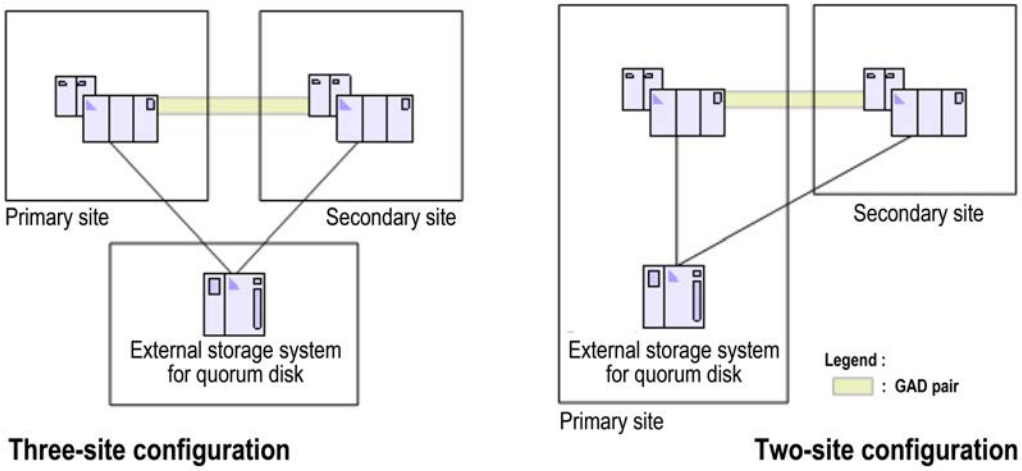

- In a three-site configuration (recommended), each storage system is located at a separate site. This configuration provides maximum protection against system failures and site failures.
- In a two-site configuration, both the primary storage system and the quorum storage system are located at the primary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

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

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

Combining global-active device and Universal Replicator

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

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

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

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

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

Setting preferred and nonpreferred paths using Asymmetric Logical Unit Access

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

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

Nondisruptive migration

One of the biggest challenges during technology refresh cycles is to eliminate downtime and service disruption when the data used by the host is copied to a new volume on the new storage system and the host is reconfigured to access the new volume. Nondisruptive migration makes it possible to relocate data from existing storage systems to newer storage systems without interrupting access by hosts. Data migration is accomplished using the global storage virtualization technology of the target storage systems. Resources on the source storage system are virtualized on the target storage system. From the perspective of the host, I/O requests continue to be serviced by the source storage system during the migration process.

The following storage system combinations are supported:

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

Nondisruptive migration offers these benefits:

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

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

How nondisruptive migration works

The following workflow summarizes the stages of the migration process.

1. Virtualization of source volumes

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

2. Switching of host I/O

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

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

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

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

3. Unallocation of source volumes

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

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

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

3. Re-creation of existing ShadowImage copy pairs

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

4. Data migration

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

Sample nondisruptive migration configuration

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

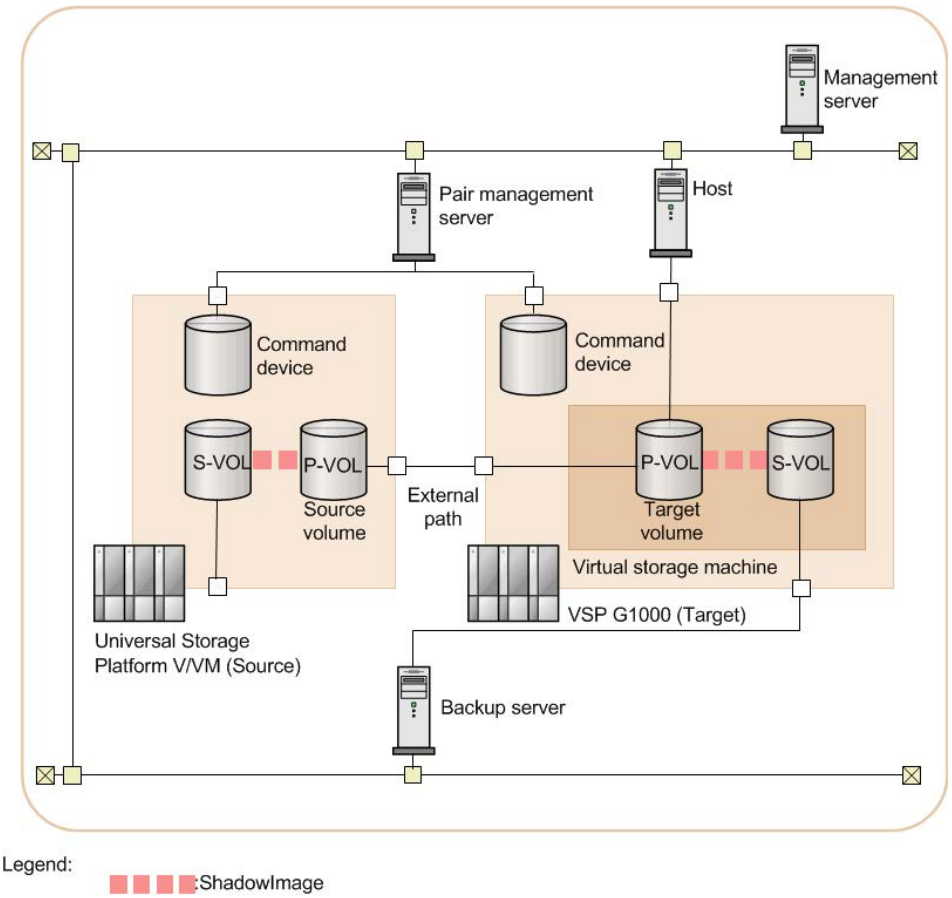

For a complete description of the nondisruptive migration feature, including requirements and setup, see the *Nondisruptive Migration User Guide* and the *Hitachi Command Suite User Guide*.

Hitachi Local Replication software

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

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

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

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

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

Hitachi Remote Replication software

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

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

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

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

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

Data-at-rest encryption

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

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

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

Management software

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

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

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

Hitachi Data Instance Director

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

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

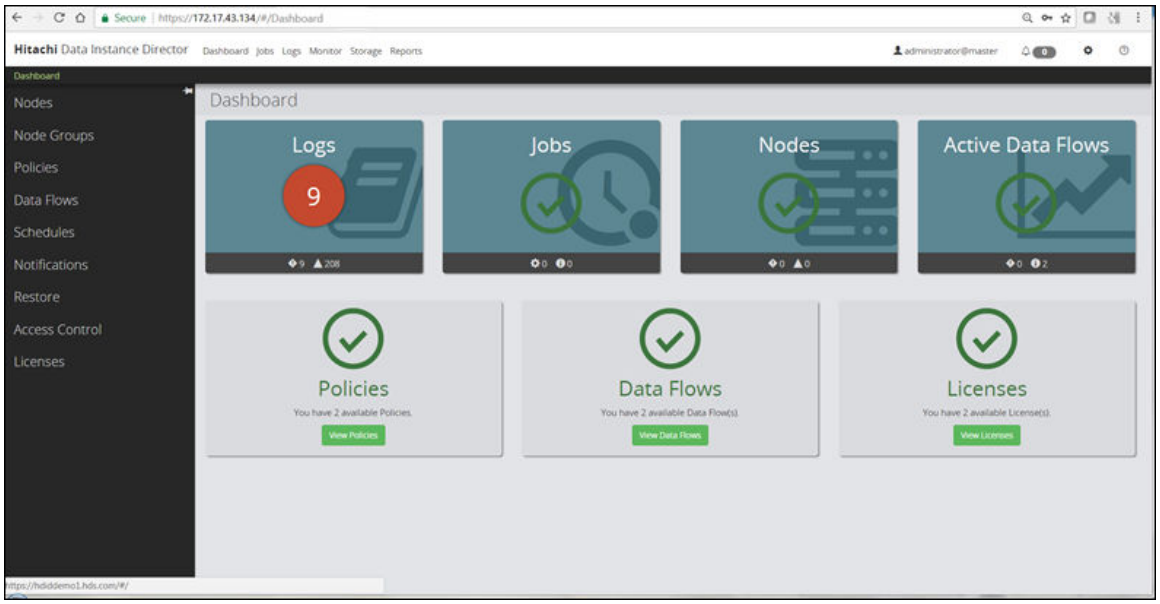

HDID provides the following benefits:

Operational recovery

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

- *Storage replication-based operational recovery*

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

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

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

- *Host-based operational recovery*

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

Disaster recovery

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

- *Storage-based disaster recovery*

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

Long-term retention

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

Unified management

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

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

Additional features of HDID include:

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

Hitachi Automation Director

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

Some of the key features of HAD are:

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

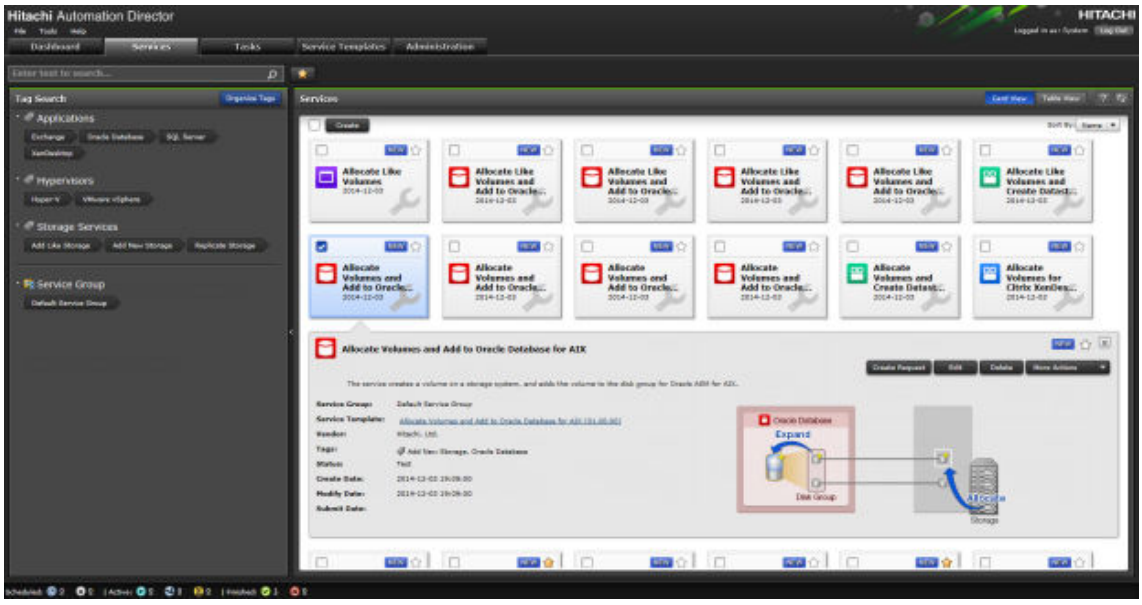

Figure 1 Select a service on the Services tab to review details and create a request for provisioning.

HAD offers the following benefits:

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

Hitachi Data Center Analytics

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

Hitachi Infrastructure Analytics Advisor

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

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

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

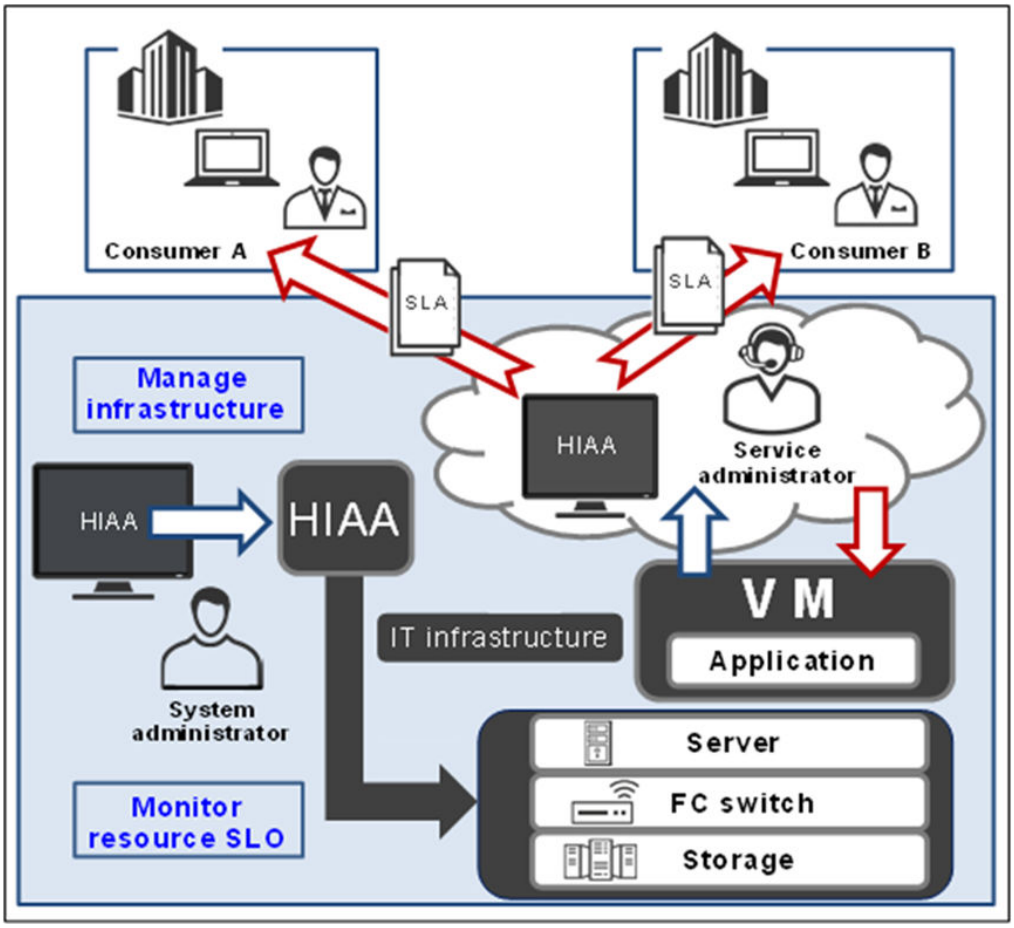

The system administrator uses HIAA to manage and monitor the IT infrastructure based on SLOs, which match the service-implementation guidelines that are negotiated under a service-level agreement (SLA) with consumers.

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

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

Chapter 5: Software solution examples

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

Delivering storage infrastructure as a service through automated workflows

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

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

Solution

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

- Implements intelligent automation workflows to streamline the storage provisioning process.
- Provides a catalog of predefined service templates and plugin components that incorporate Hitachi best practices in storage provisioning and that minimize human error.
- Provides customizable storage service templates requiring minimal input that administrative users can use to increase operational efficiency.

- Optimizes storage configurations for common business applications such as Oracle, Microsoft Exchange, Microsoft SQL Server and hypervisors such as Microsoft Hyper-V and VMware.
- Analyzes current storage pool capacity utilization and performance to automatically determine the optimized location for new storage capacity requests and to make storage provisioning more cost-efficient.

Management software

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

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

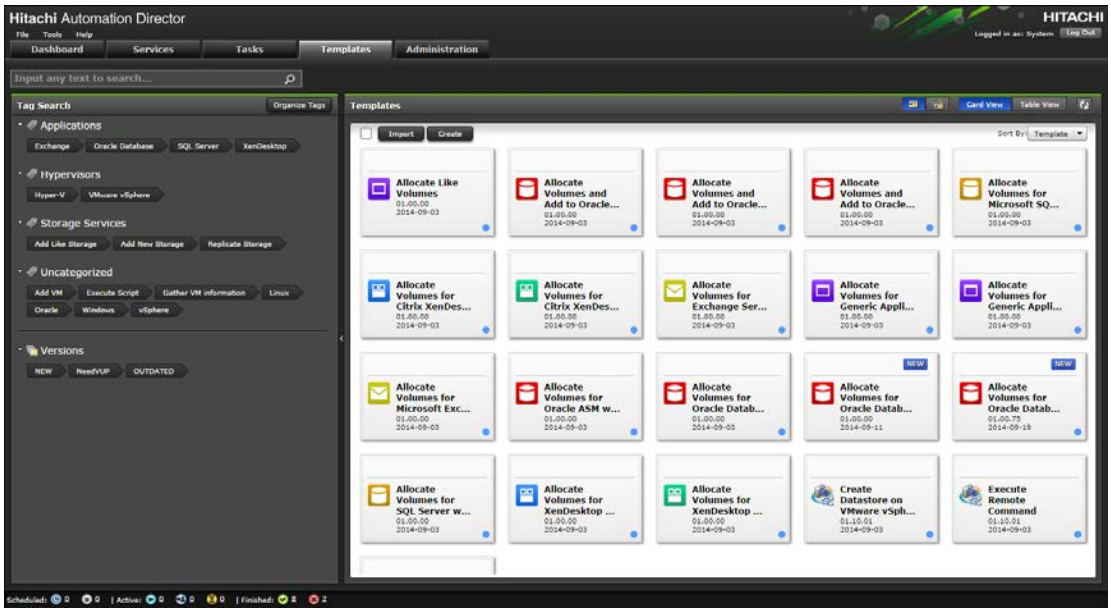

HAD includes a comprehensive tool, Service Builder, to create and modify existing workflows and plug-in components that automate the storage management tasks for a given operating environment.

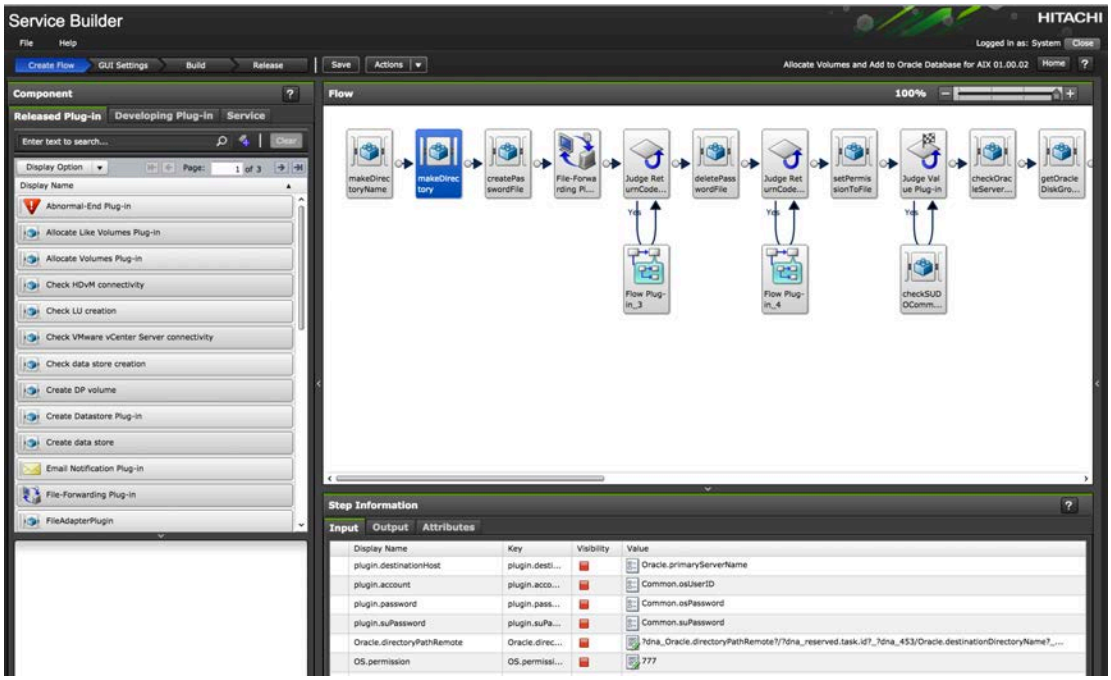

HAD supports all native block storage systems and third-party storage systems through virtualization technology.

Data protection for business-critical Oracle databases

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

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

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

Solution

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

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

Management software

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

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

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

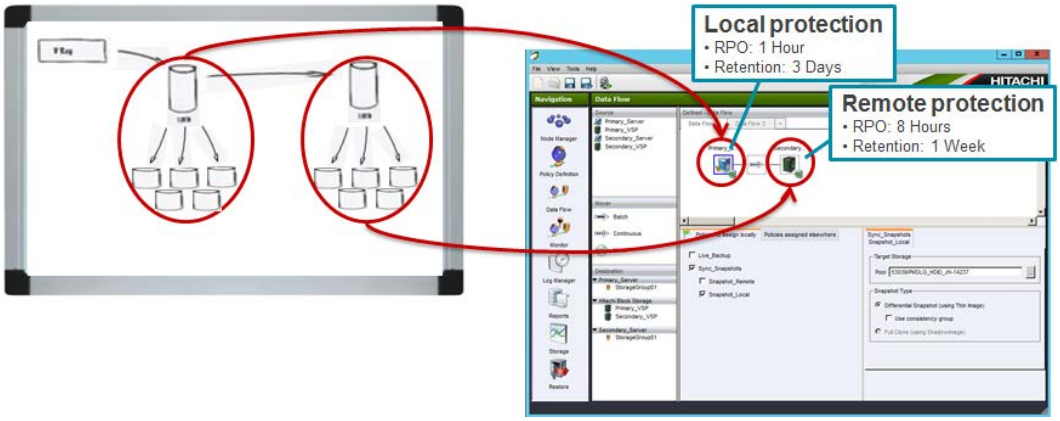

You can use different methods to back up data across multiple sites, as described in the following table and figure.

Method	Description
Identical snapshots and clones	Provide identical RPO and data retention regardless of location. Keeping identical backups provides identical recovery options and procedures during a site failover, which simplifies the entire restore process.

Method	Description
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.

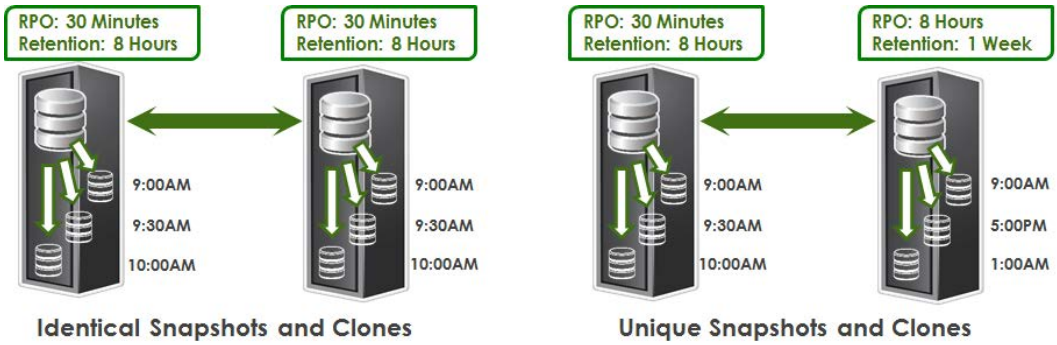

Using Infrastructure Analytics Advisor for data analysis: from deep dive to recovery planning

Hitachi Infrastructure Analytics Advisor is bundled with Hitachi Data Center Analytics in a software package called Performance Analytics.

Infrastructure Analytics Advisor provides an intuitive GUI for performance monitoring, management, and troubleshooting.

Data Center Analytics performs the data collection from monitored targets (such as storage systems, hosts, and switches) using software probes that support each device or environment. Data Center Analytics also provides historical trend analysis and extensive report generation capabilities.

The Performance Analytics solution provides end-to-end monitoring and troubleshooting capabilities for your infrastructure resources, from host to storage system. The basic workflow for Performance Analytics troubleshooting is called the MAPE loop:

- Monitor
- Analyze
- Plan
- Execute

When reviewing and evaluating reports and event information on the Infrastructure Analytics Advisor Dashboard, you can also perform a deep dive analysis by invoking the Data Center Analytics UI. The deep dive is part of the **Analyze** segment of the MAPE loop workflow.

The following workflow is an example of how to use this troubleshooting methodology as an infrastructure administrator, managing user resources (such as consumers, VMs, and volumes) and system resources (such as cache, ports, CPUs, and disks).

Viewing the dashboard

As an infrastructure administrator, you set up dynamic thresholds on the user resources you are monitoring. After seeing nine critical alerts on VM/Host resource gauge, you become interested in troubleshooting a threshold violation.

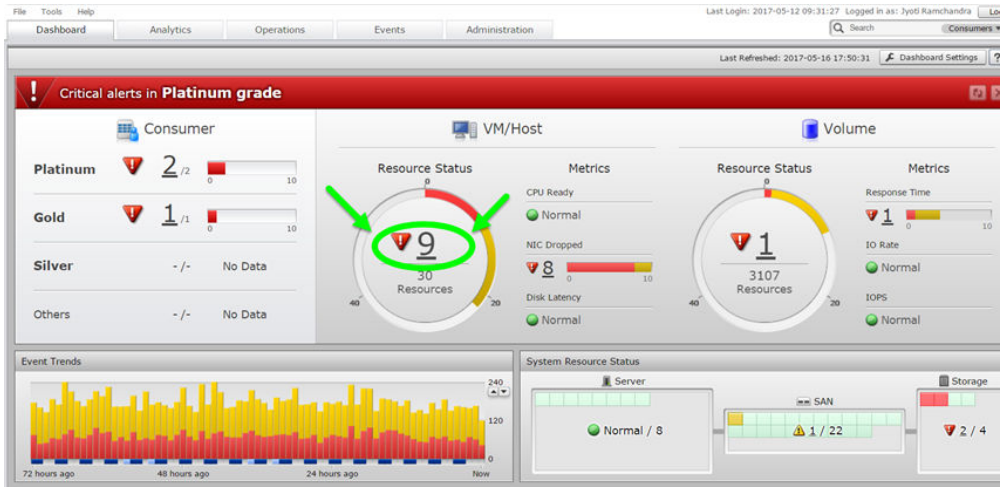

You browse the resources with critical alerts and select the target VM to analyze in the E2E View.

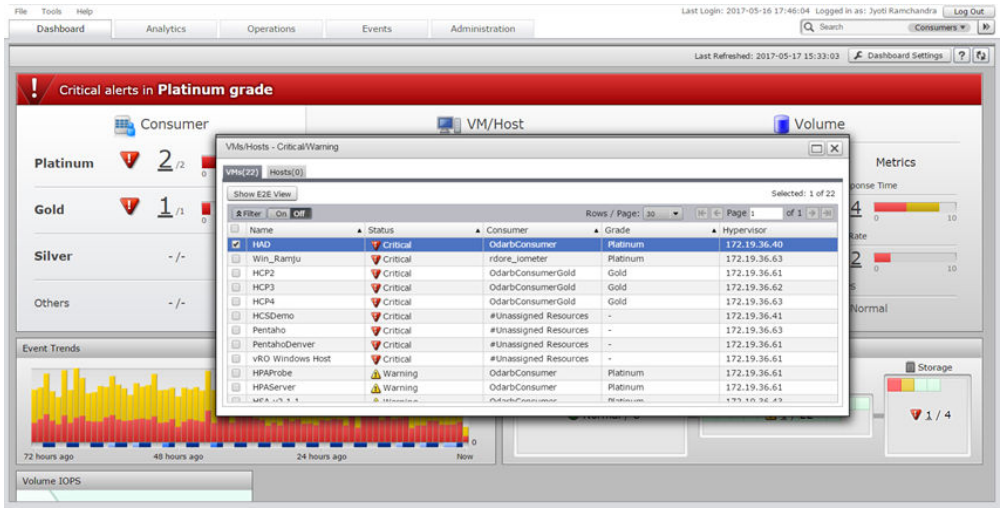

Using E2E or Sparkline views

The E2E view represents the topology of infrastructure resources: from host, to fabric switch, to storage system. The infrastructure administrator sets the base point of analysis on the target resource for analysis. This view allows you to see the relationship between resources.

To move deeper into the underlying resources, you can invoke the Sparkline view, which presents multiple charts that track performance by component. Use this view to correlate performance trends between user and system resources.

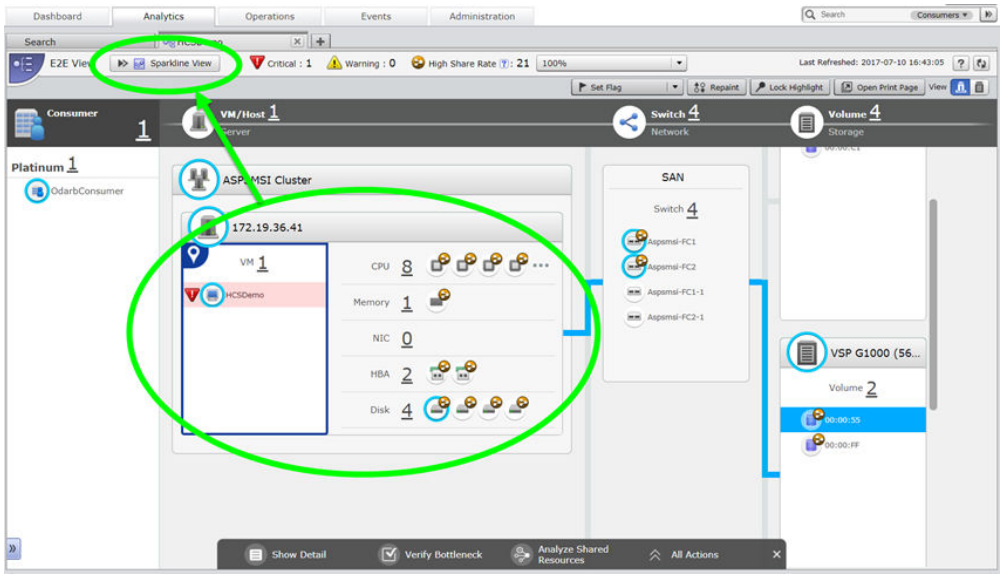

Using additional troubleshooting tools

Infrastructure Analytics Advisor offers multiple troubleshooting tools for isolating a bottleneck candidate and identifying its root cause. You can launch any of the following tools for further analysis:

- **Verify Bottleneck**: Use at the initial stage of analysis to compare performance charts of the base point of analysis with the bottlenecked candidate.
- **Identify Affected Resources**: Use to display the user resources that rely on the bottlenecked resource.
- **Analyze Shared Resources**: Use if you suspect that the root cause of the problem is resource contention, a noisy neighbor that disrupts the balance of resource usage. You compare performance charts of the bottleneck candidate to the resources using the bottleneck. After comparing performance across a number of resources with Analyze Shared Resources, you isolate the actual bottleneck.
- **Analyze Related Changes**: Use if **Analyze Shared Resources** does not reveal the actual bottleneck (noisy neighbor), or if you suspect that the root cause of the problem is a recent configuration change. In this view, you compare performance charts with configuration events. The bar graph portion of the chart represents the configuration changes made at a particular time. You can click on a bar to list those changes.

Performing a deep dive analysis

Regardless of which tool you use, once you have isolated the bottleneck candidate and validated its root cause, you can collect more information to understand its origin. For example, you have identified a storage system as the bottleneck. Subsequently, you want to understand how the problem affects other resources or vice versa. This phase of the troubleshooting analysis is called the deep dive. In a deep dive analysis, you can compare the data of various components from the resource tree, which displays all the resources and their components in your infrastructure, and run a customized report against that data.

To proceed with the deep dive for information, launch the Data Center Analytics UI, which provides detailed reports at the component level. You can launch this component-level view from the following windows in the Infrastructure Analytics Advisor UI during analysis:

- E2E view
- Sparkline view
- Performance tab of the Show detail window for a resource
- Analyze Shared Resources
- Analyze Related Changes

When analyzing system resources in Data Center Analytics, you can view performance charts based on various metrics to correlate components with resource performance. For example, you have validated the root cause of the storage system bottleneck, but you want to perform further analysis in Data Center Analytics.

The following figure examines the performance of the volume from the VM side. This report, **LDEV IOPS versus Response Time**, displays spikes at specific times, which you can then use as reference points for when the I/O activity was particularly intensive during otherwise typical workloads.

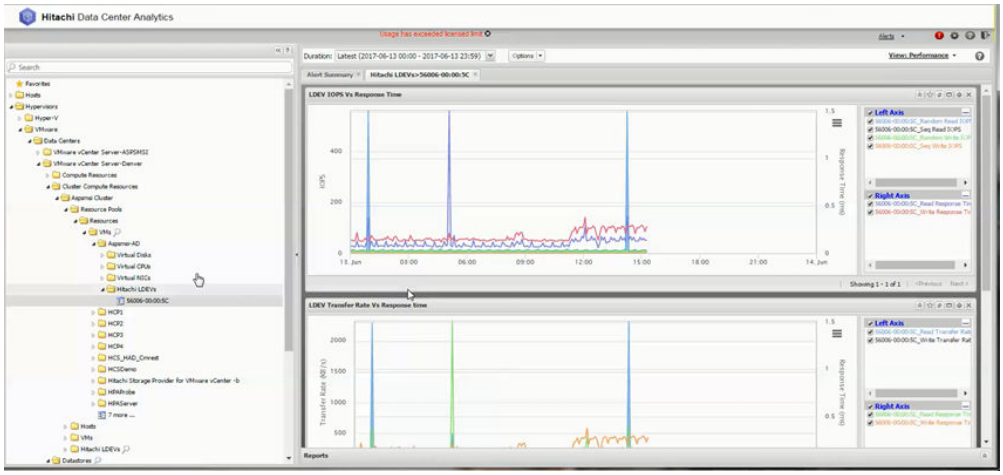

Digging deeper, you discover the storage systems and volumes associated with a particular VM. You cross-reference the resources in the VM performance chart and determine that the component with the performance that correlates to the VM is the cache on the storage side, or (specifically, CLPR, or Cache Write Pending Rate). This workload is typically intensive, but you realize that the times when the resource reached 100% correlate with the spikes in the **LDEV IOPS versus Response Time** report.

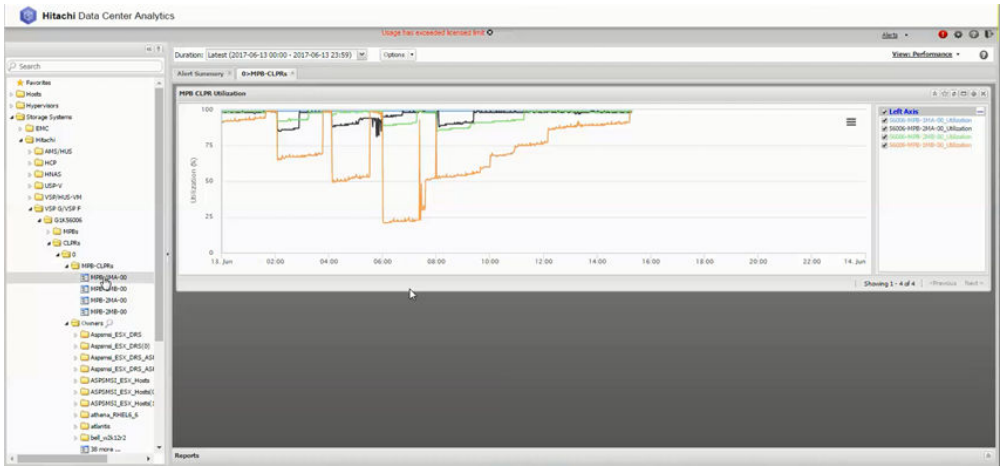

Often, the performance problem is a recurring trend; for example, when monitoring certain infrastructure resources, you notice spikes in I/O activity every weekday at 3 PM. When you create a customized report, you discover this trend has persisted for six months. (In theory, you can review performance from months to years.) This capability to review past performance adds a historical element to deep dive analysis.

Initiate recovery plan to solve the performance problem

After establishing the correlation between the two charts, you return to the Infrastructure Analytics Advisor UI to initiate a recovery plan. You can enter the key metric, date, and time of the problem occurrence, and the target value for the metric. In this case, the problem component is the CLPR; the key metric is IOPS. You can specify conditions, then review the recovery plan generated by Infrastructure Analytics Advisor before executing it.

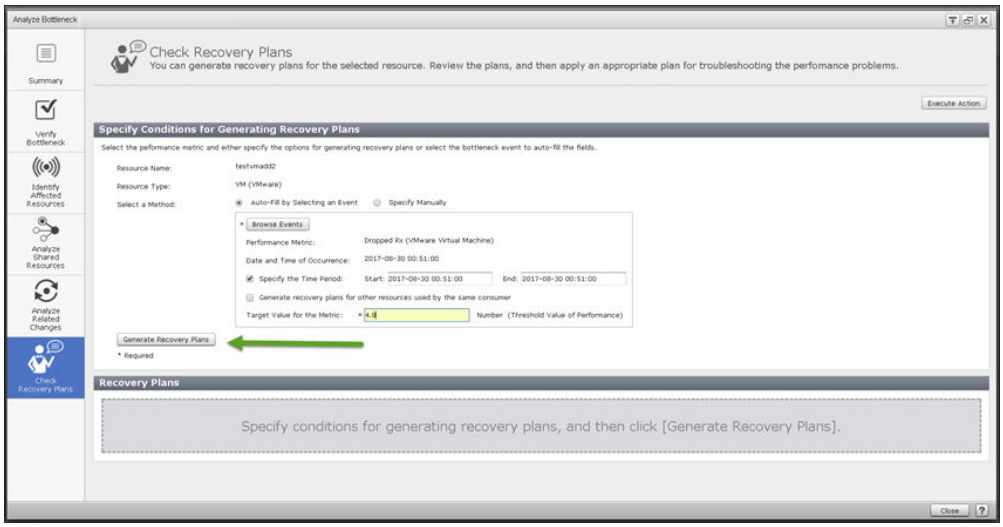

After you have successfully executed the recovery plan, you can adjust your thresholds with new metric settings to monitor the user resources (in this case, the VM and the affected volume). At this stage, you have completed the MAPE loop.

Hitachi Vantara Corporation

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

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