HITACHI
Inspire the Next

Hitachi Unified Storage VM
Product Overview Guide

FASTFIND LINKS

MK-92HM7003-05

Hitachi Data Systems

Contents

Preface

This manual provides instructions and information to use the Hitachi Unified Storage VM storage system.

Read this document carefully to understand how to use this product, and keep a copy for reference.

- ☐ Safety and environmental information
- ☐ Intended audience
- ☐ Release notes
- ☐ Product version
- ☐ Document revision level
- ☐ Changes in this revision
- ☐ Referenced documents
- ☐ Document conventions
- ☐ Convention for storage capacity values
- ☐ Accessing product documentation
- ☐ Getting help
- ☐ Comments

Safety and environmental information

Caution: Before operating or working on the Hitachi Unified Storage VM storage system, read the safety and environmental information in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

Intended audience

This document is intended for system administrators, Hitachi Data Systems representatives, and authorized service providers who install, configure, and operate the Hitachi Unified Storage VM storage system.

Readers of this document should be familiar with the following:

- Data processing and RAID storage systems and their basic functions.
- The Hitachi Unified Storage VM storage system and the *Hitachi Unified Storage VM Block Module Hardware User Guide*.
- The Storage Navigator software for the Hitachi Unified Storage VM storage system and the *Hitachi Storage Navigator User Guide*.

Release notes

The Hitachi Unified Storage VM Block Module Release Notes provide information about the HUS VM block module microcode (DKCMAIN and SVP). A second set of release notes provides information about the HUS VM file module NAS operating system (SU and SMU). Each of the release notes describe the new features, functions, and changes. Both release notes are available on the Hitachi Data Systems Portal: https://portal.hds.com

Product version

This document revision applies to Hitachi Unified Storage VM firmware version 73-03-0*x* and later.

Document revision level

Revision	Date	Description
MK-92HM7003-00	September 2012	Initial release
MK-92HM7003-01	December 2012	Supersedes and replaces MK-92HM7003-00
MK-92HM7003-02	March 2013	Supersedes and replaces MK-92HM7005-01
MK-92HM7003-03	May 2013	Supersedes and replaces MK-92HM7005-02
MK-92HM7003-04	October 2013	Supersedes and replaces MK-92HM7005-03
MK-92HM7003-05	November 2013	Supersedes and replaces MK-92HM7005-04

Changes in this revision

- Changed the name of this document from Getting Started Guide to Product Overview Guide
- Reorganized and rewrote the contents of this manual to fit the new manual name.
- Added File Module information to chapters 1, 2, and 3.
- Added 1.2 TB 10K SAS SFF and 3.2 GB FMD data drives to chapters 1 and 2.

Referenced documents

A complete list of both HUS VM block module and file module documentation is located in Hitachi Unified Storage VM user guides on page A-1.

Document conventions

Hitachi Data Systems user manuals use the following typographic conventions as needed to clarify information.

Convention	Description
Bold	Indicates the following: • Text in a window or dialog box, such as menus, menu options, buttons, and labels. Example: In the **Add Pair** dialog box, click **OK**. • Text appearing on screen or entered by the user. Example: The **-split** option. • The name of a directory, folder, or file. Example: The **CacheInfo.csv** file.
Italic	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy *source-file target-file* Angle brackets are also used to indicate variables.
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `# pairdisplay -g oradb`
`< >` angle brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: `# pairdisplay -g <group>` Italic is also used to indicate variables.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.
| vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a | b] indicates that you can choose a, b, or nothing. { a | b } indicates that you must choose either a or b.

Hitachi Data Systems user manuals use the following icons as needed to draw attention to information.

Icon	Meaning	Description
	Tip	Helpful information, guidelines, or suggestions for performing tasks more effectively.
	Important	Information that is essential to the completion of a task.

Icon	Meaning	Description
	Caution	Failure to take a specified action can result in adverse conditions or consequences such as damage to the software or hardware
	WARNING	Failure to take a specified action can result in severe conditions or consequences such as in loss of data or serious damage to hardware.
	ELECTRIC SHOCK HAZARD	Failure to take appropriate precautions such as not opening or touching hazardous areas of the equipment could result in injury or death.

Convention for storage capacity values

Physical and logical storage capacities of disk drives in Hitachi Data Systems storage products are calculated based on the following values:

Physical Disk Capacity	
1 KB = 1,000 bytes	1 TB = $1{,}000^4$ bytes
1 MB = $1{,}000^2$ bytes	1 PB = $1{,}000^5$ bytes
1 GB = 1,0003 bytes	1 EB = $1{,}000^6$ bytes

Logical storage capacity values (logical device capacity) are calculated based on the following values:

Logical Disk Capacity (1 block= 512 bytes)	
1 KB (kilobyte) = 1,024 bytes (2^{10})	1 TB (terabyte) = $1{,}024^4$ bytes
1 MB (megabyte) = $1{,}024^2$ bytes	1 PB (petabyte) = $1{,}024^5$ bytes
1 GB (gigabyte) = $1{,}024^3$ bytes	1 EB (exabyte) = $1{,}024^6$ bytes

Accessing product documentation

The Unified Storage VM user documentation is available on the Hitachi Data Systems Support Portal: https://Portal.HDS.com. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, log on to the Hitachi Data Systems support portal for contact information: https://Portal.HDS.com

Comments

Please send us your comments about 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.

Thank you!

1

System overview

This chapter provides an overview of the Hitachi Unified Storage VM storage system. Detailed information is located in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

NOTE: The Hitachi Unified Storage VM storage system may be installed only by Hitachi Data Systems service personnel or authorized service partners. However some high level installation information is included in this manual and in the *Hitachi Unified Storage VM Block Module Hardware User Guide* for reference.

- ☐ System features
- ☐ Block module storage system
- ☐ File module storage system

System features

The Hitachi Unified Storage VM storage system is an entry level enterprise storage platform combining storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a new storage virtualization system to provide central storage services to existing storage assets. Unified management delivers end to end central storage management of all virtualized internal and external storage. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, and virtual NAS functionality without compromising performance or scalability.

Using this system, you can deploy virtual server applications within a new storage virtualization framework, extend new services to your current storage investments, and more closely align IT with business objectives by enhancing capacity efficiency. Unified Storage VM storage systems provide the foundation for matching application service requirements to different classes of storage and for delivering critical services including:

- Virtualization of all storage assets
- Business continuity data protection services
- Simplified management of block, file and object data
- Host transparent data migration
- Dynamic storage tiering and thin provisioning
- High availability with scalable performance
- High reliability under all configurations and loads
- Non-disruptive service and upgrades
- Advanced file management services
- Secure multiple tenant partitioning
- Simplified, extremely easy-to-use data management
- A high degree of data protection for business
- Transparent data mobility for any number of different storage tiers
- The ability to reduce storage space through primary data de-duplication and TDM
- Industry-leading scalability
- Exceptional performance
- Low total cost-of-ownership

Block module storage system

This section briefly describes the Unified Storage VM block module storage system.

Overview

Hitachi Unified Storage VM block module storage systems contain new technology that was not available in previous Hitachi Data Systems storage systems. This technology is used to create a smaller, more efficient, cooler running controller than the VSP controller, while using the same high speed architecture.

Hitachi Unified Storage VM is based on a new block storage system powered by a storage virtualization controller connected to the data drives and the host servers. Its dual node, shared resource architecture creates a redundant configuration in which the storage system can continue operation should a component failure occur. Main components can be added, removed, and replaced without shutting down a device and while the storage system is in operation. The microcode can also be upgraded without shutting down the storage system. A service processor that monitors the operational status of the storage system is mounted in the controller chassis. Connecting the service process with a Hitachi service center enables remote maintenance.

The system can be configured starting with a single diskless system up to a large, three-rack system that includes up to 1,152 HDD drives including up to 128 SSD drives and up to 96 flash module drives, and a total of 256 GB cache. The system provides a highly granular upgrade path allowing the addition of file modules or data drives to the drive chassis in an existing system as storage needs increase, all mounted in standard 19-inch racks. A basic Unified Storage VM storage system consists of a controller chassis and either no drives or one or more drive boxes that contain the hard disk drives or solid state drives. Unified Storage VM supports multiple concurrent operating systems to create a heterogeneous system environment.

Features

This section describes the main features of the Unified Storage VM block module storage system.

Scalability

The Unified Storage VM storage system is highly scalable and can be configured in several ways as needed to meet customer requirements.

- The minimum configuration is a single rack containing one controller chassis in a diskless configuration.
- A single rack containing one controller chassis and up to seven dense drive boxes or 16 standard drive boxes, or a combination of the two.
- The maximum configuration is a system that contains one controller chassis and a combination of drive boxes containing a maximum of 1,152 drives, including 2-1/2 inch and 3-1/2 inch HDD or SSD drives,

and 5-1/4 inch flash module drives. The total storage space of this configuration is 3,383 TB. The *Hitachi Unified Storage VM Block Module Hardware User Guide* contains details about the data drives.

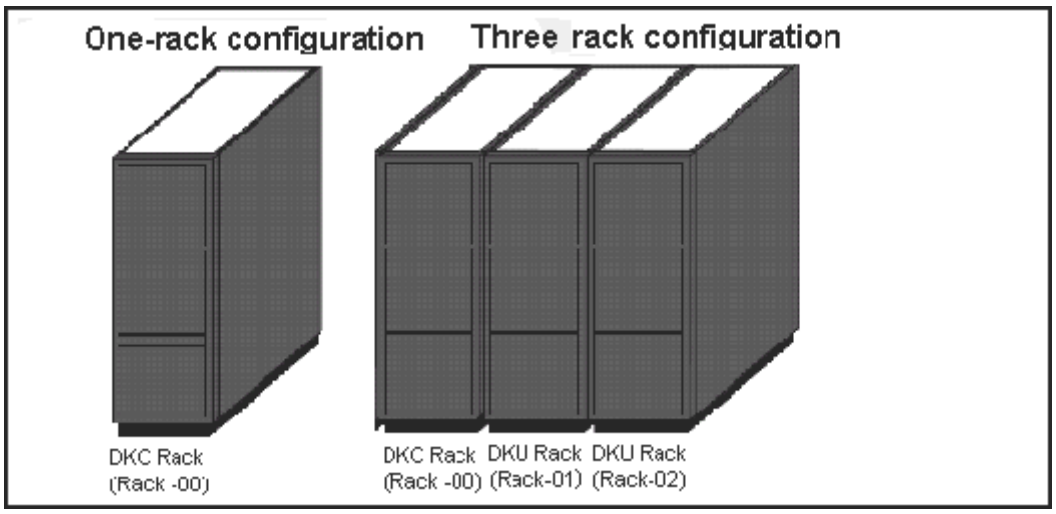

Figure 1-1 Example HUS VM storage system configurations

In addition to a varying number of data drives, the system can be configured with disk drives of different capacities and speeds, varying numbers of host I/O modules and back-end I/O modules, and varying cache capacities, as follows:

- Two to four host I/O modules (each is a pair of boards). This provides a total of 8 when all of the host I/O module slots are used and there are no back-end I/O modules installed, as in a diskless system. The maximum total number of host I/O modules and back-end I/O modules is 8.
- Two to four back-end I/O modules (each is a pair of boards). This provides a total of 8 when all of the back-end I/O module slots are used. In this case only two host I/O modules can be installed.
- Cache memory capacity: multiple configurations from 32 GB to 256 GB
- Disk drive capacities of 300 GB, 600 GB, 900 GB, and 3TB
- Flash drive capacities of 200 GB, 400 GB and 800 GB
- Maximum available channel ports: 40 with disks, 48 if diskless

High performance

The HUS VM block module storage system includes several features that provide very high system performance. These include:

- High speed disk drives that run at 7,200, 10,000, or 15,000 RPM
- SSD flash drives with ultra high speed response
- Flash module drives with ultra high speed response
- High speed data transfer between the back-end I/O module and HDDs at a rate of 6 GBps with the SAS interface

- High speed quad core CPUs that provide the same performance as the VSP system (three times the performance of a Universal Storage Platform V/VM storage system).

High capacity

The HUS VM block module storage system supports the following high-capacity features. See the *Hitachi Unified Storage VM Block Module Hardware User Guide* for details.

- HDD (disk) drives with capacities of 300 GB, 600 GB, 900 GB 1.2 TB, and 3TB.
- SSD (flash) drives with capacities of 200 GB, 400 GB, and 800 GB.
- Flash module drives with a capacity of 1.6 TB or 3.2 TB each.
- Controls up to 16,384 logical volumes and up to 1,152 disk drives, and provides a maximum physical disk capacity of approximately 3,383 TB per storage system

Connectivity

Hitachi Unified Storage VM

The HUS VM block module storage system supports most major operating systems, such as Microsoft Windows, Oracle Solaris, IBM AIX, Linux, HP-UX, Novell Netware, SUSU Linux, Red Hat Linux, And VMware ESX Server. For more complete information on the supported operating systems, go to: http://www.hds.com/products/interoperability/index.htm. The system supports the Fibre Channel host interface.

Storage Navigator

The required features for the Storage Navigator computer include operating system, available disk space, screen resolution, CD drive, network connection, USB port, CPU, memory, browser, Flash, and Java environment. These features are described in Chapter 1 of the *Hitachi Storage Navigator User Guide*.

High reliability

The HUS VM block module storage system includes the following features that make the system extremely reliable:

- The system supports RAID6 (6D+2P/14D+2P), RAID5 (3D+1P/7D+1P), and RAID1 (2D+2D/4D+4D). See the *Hitachi Unified Storage VM Block Module Hardware User Guide* for more information on RAID levels.
- All main system components are configured in redundant pairs. If one of the components in a pair fails, the other component performs the function alone until the failed component is replaced. Meanwhile, the storage system continues normal operation.

- The HUS VM block module storage system is designed so that it cannot lose data or configuration information if the power fails. This is explained in Battery backup operations in the Hitachi Unified Storage VM Block Module Hardware User Guide.

Non-disruptive service and upgrades

The HUS VM block module storage system is designed so that service and upgrades can be performed without interrupting normal operations. These features include:

- All main components can be “hot swapped” — added, removed, and replaced without any disruption while the storage system is in operation. The components include the Main blade, Cache memory module, Control PCB, Cache flash memory, power supplies, fans, disk drives, and flash drives. However, while additional cache memory is installed, the channel ports and drive ports of the affected cluster are blocked.
- A Service Processor mounted in the controller chassis monitors the running condition of the storage system. Connecting the SVP with a service center enables remote maintenance.
- The firmware (microcode) can be upgraded without disrupting the operation of the storage system. The firmware is stored in shared memory (part of the cache memory module) and transferred in a batch, reducing the number of transfers from the SVP to the controller chassis via the LAN. This increases the speed of replacing the firmware online because it works with two or more processors at the same time.

Economical and quiet

The three-speed fans in the controller chassis and drive boxes are thermostatically-controlled. Sensors in the units measure the temperature of the exhaust air and set the speed of the fans only as high as necessary to maintain the unit temperature within a preset range. When the system is not busy and generates less heat, the fan speed is reduced, saving energy and reducing the noise level of the system.

Power Saving Mode. When the storage system is in standby mode, the disk drives spin down and the controller and drive chassis use significantly less power. A system in standby mode uses approximately 70% of the power that it uses during normal operation.

File module storage system

This section describes the Unified Storage VM file module.

Overview

The Hitachi Unified Storage VM file module features a hybrid core architecture using the best properties of FPGA-based design to optimize data movement coupled with high performance, multi-core processors for efficient data management functions. Both classes of activity work at full speed without an impact on the other. They can handle a number of simultaneous workloads such as serving email to thousands of users and

hosting large scale OLTP applications, while maintaining high performance. They also provide high IOPS performance and utilize built-in 10 Gigabit and 1 Gigabit Ethernet for high throughput NAS and iSCSI networking connectivity. Up to 4 nodes in a single cluster will meet demands for scalable storage with greater access, capacity and performance. Storage can be added at any time to meet new application or business needs, or consolidate disparate storage all with a single point of management, without downtime. These systems offer a total usable capacity of from 4 to 8 PB under a single namespace, all easily managed from a central system management console.

The Hitachi File System is built around the object store, a collection of object structures referring to data on disks , and a set of rules. These rules govern the organizational layout and management of objects in the object store. The techniques behind creating, copying, moving, migrating, and deleting the objects in the object store make the Hitachi File System extremely effective and efficient.

Tiered file systems decrease file access times by placing file metadata on the highest performance storage, while file content is placed on storage with less performance. A policy-based automatic file migration feature helps organizations move data among storage and archive tiers whether internal to the data center, or externally from remote or branch offices. The cluster namespace functionality provides a single namespace with a directory structure that is independent of where data actually resides in physical storage. Virtual servers are used to group server resources to match the needs of application or organizational requirements. In a clustered environment, file systems can be quickly relocated among physical servers for load balancing, and virtual servers automatically migrate in a cluster failover scenario.

Features

This section describes the key features and benefits of the Unified Storage VM storage system file module.

Scalability

The Unified Storage VM storage system file module is highly scalable and can be configured in several ways as needed to meet customer requirements, without impacting performance.

The Hitachi File System can support billions of files in a single file system and millions of files in a single directory, while keeping directory search times to a minimum and sustaining overall system performance. Combined with cluster namespace, the Hitachi File System can support many petabytes in a single unified namespace and present it all as a single file system accessible to many concurrent hosts, through a single mount point if desired.

The basic configuration is a single rack containing one NAS server and up to 8U of disk storage, as shown in the following figure.

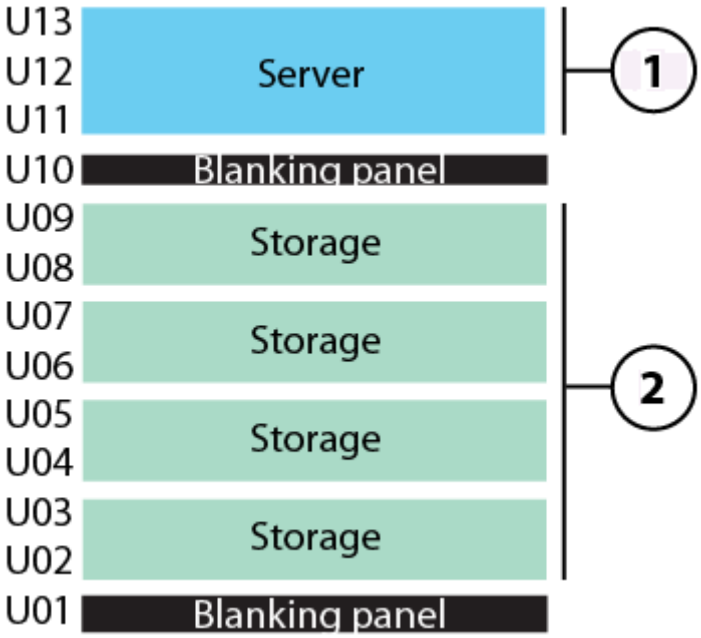

Figure 1-2 Basic Unified Storage VM file module configuration

Item	Description	Item	Description
1	Management block / server (equivalent to block module controller chassis)	2	Storage block (data drives)

The system can contain a single file module (Hitachi NAS Platform), or several file modules operating as a cluster. Clusters using more than two Hitachi NAS Platforms typically include 2 x 10Gbps Ethernet switches (one is required, but two are recommended for redundancy). For instructions on how to implement a cluster configuration, refer to the Server and Cluster Administration Guide.

Figure 1-3 Example Unified Storage VM file module configurations on page 1-9 shows two examples of file module system configurations including a rack containing up to two NAS servers, two system management units (SMUs), two fibre channel switches, and up to 32U of storage per rack, depending on the number of servers and system management units installed. The hardware section following the figure describes the individual system components.

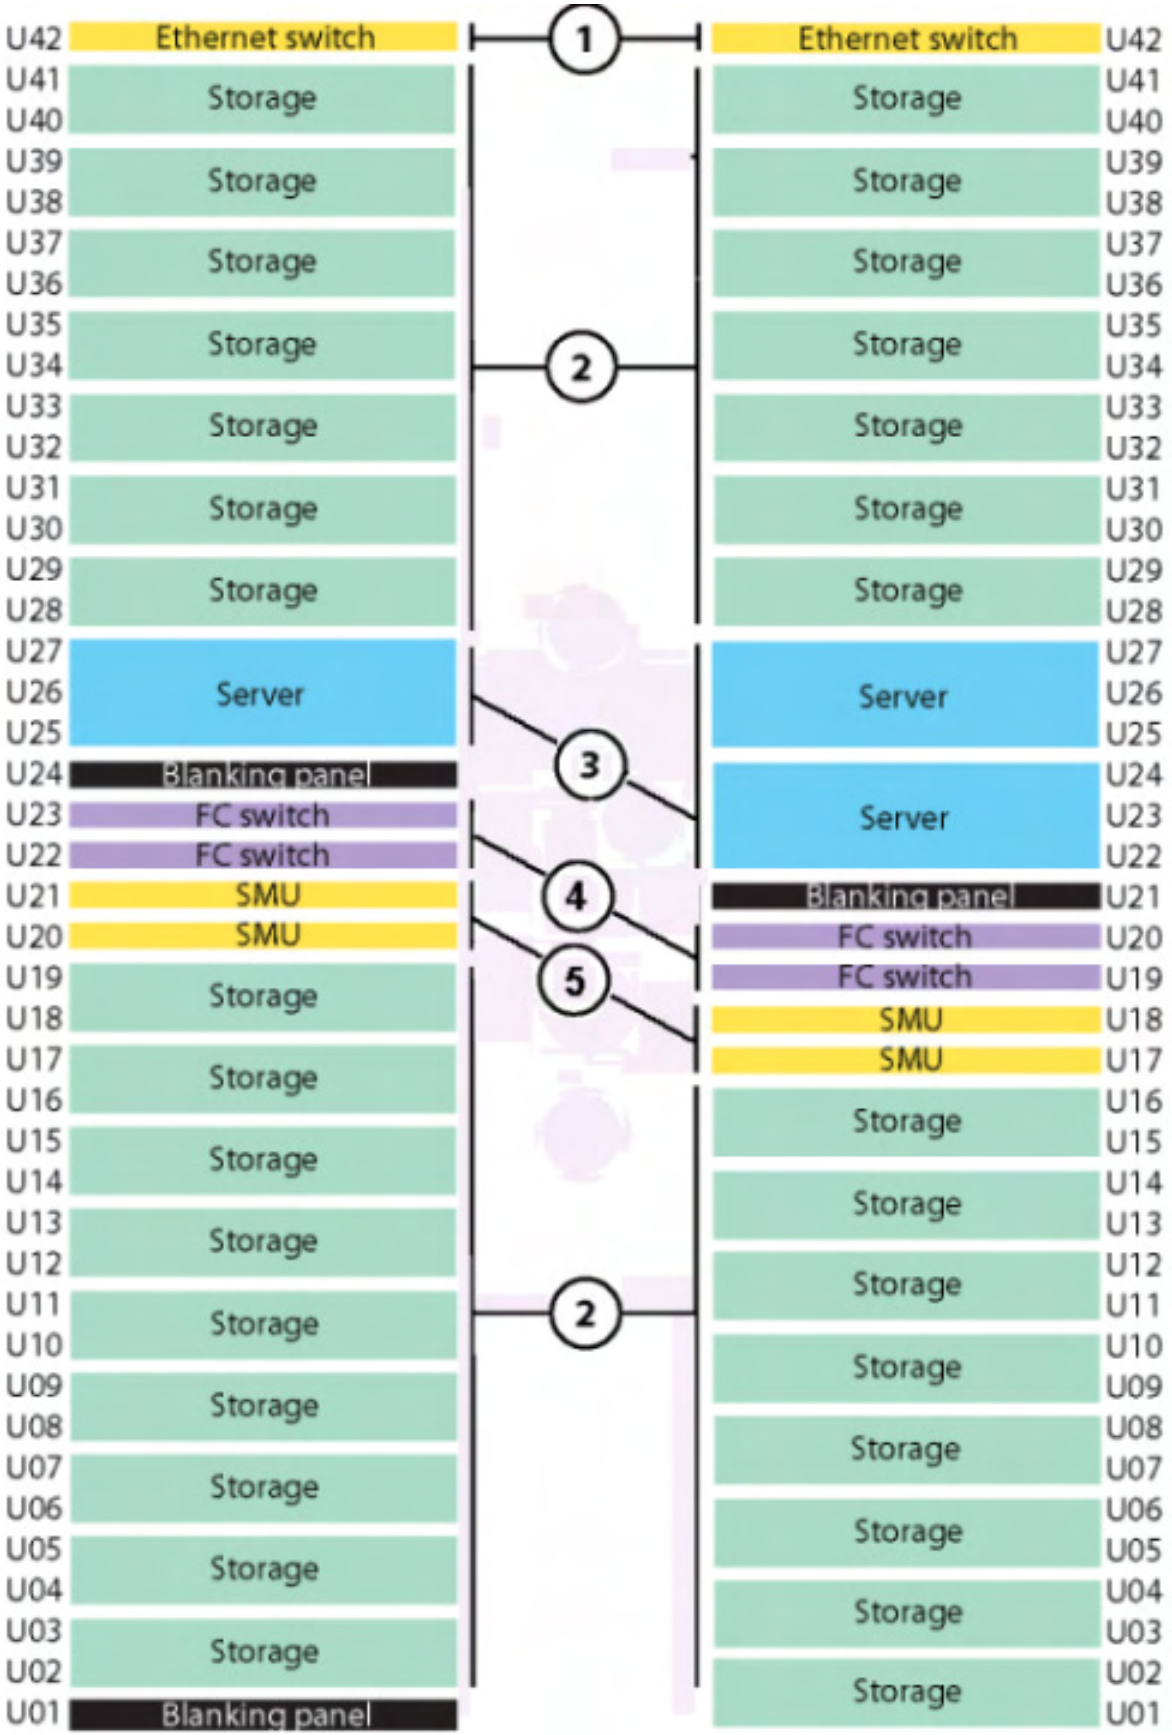

Figure 1-3 Example Unified Storage VM file module configurations

Item	Description	Item	Description
1	Management block (ethernet switch)	2	Storage block (data drives)
3	Management block (server/ controller chassis)	4	Fibre channel switch
5	Management block (system management unit)		

Primary data de-duplication

The Hitachi File System has built-in data deduplication capability that reduces storage consumption, and is particularly effective when used for server virtualization or VDI. Unlike other NAS products, the Hitachi File System achieves data reduction with minimal impact on overall system performance, so dedupe can be enabled even when the system is under load. Combined with Hitachi flash technology, the Hitachi File System can deliver superior performance for a fraction of the cost.

Consolidation and virtualization

Scalability enables consolidation, particularly of older hardware and "storage islands". Creating a unified, large-scale storage solution allows storage administrators to combine the functions of what were separately implemented file servers, reaping cost-savings and management benefits of a consolidated platform. Any successful consolidation strategy has to support virtual server and virtual desktop (VDI) environments. The Hitachi File System will operate very effectively in virtualization technology, such as VMware, providing functionality like space-efficient writeable clones, which enables fast backup and recovery of VM images.

Storage virtualization

Virtualization is about making more efficient use of high performance NAS server. The more powerful the server is, the better suited it is to virtualize a larger number of less capable, under-utilized devices. HNAS can virtualize 3rd party NAS devices and relocate data in the background, allowing the customer to preserve existing assets until they can be taken out of service. Our implementation of virtual servers allows groups to retain "ownership" of their virtual entity within a single physical server. Thin provisioning makes it possible for multiple virtual servers to share a single pool of storage devices.

The Hitachi File System offers benefits beyond sustained, predictable, consistent file server performance. Because of its architecture, the Hitachi File System can adjust to the users' workflow and data sets. Not all data has the same "value" to user workflow. When data is first created it may be extremely valuable and must therefore reside on storage built for high performance. As the data ages, application and eventually archival requirements tend to dominate, imposing further conditions on the storage where the data now resides. Yet, once an application knows where the data resides, data relocations are difficult. The Hitachi File System eliminates this difficulty by providing a mechanism to migrate transparently across multiple

tiers of storage, including devices optimized for archival. Users do not necessarily need to be aware of the actual data location, and applications need not be rewritten either

High flexibility and wide applicability

The Hitachi File system flexible architecture enables wide applicability to changing workloads, data sets, and access patterns. It provides high scalability, fast metadata processing performance, and excellent data movement speed under a wide range of host loads, usage patterns, and data types. Fine-grained parallelism, off-loading of specific file system operations to FPGAs, and data pipelining all contribute to the Hitachi File System handling of both throughput and metadata processing.

Separation of server and storage functions

The Hitachi File System delivers performance with relatively small storage systems. The file system also allows for performance to increase granularly as more disks are added. Typically this benefit is immediately apparent, even before "restriping" the data across both old and new spindles, as writes spread automatically and immediately. As a result, customers may start small and scale performance by adding storage when needed. Additional file servers are not necessarily required for additional performance. As performance requirements grow, customers may also take advantage of clustering technology within the Hitachi File System to add more servers while maintaining a single namespace. This provides easy management of large pools of data, but the Hitachi File System still offers true separation of function between storage and servers. Each may be scaled independently to meet your needs; there is no requirement to purchase one to get the other as with many competing NAS product

High namespace scalability

Scaling beyond a single NAS server is essential for high performance storage solutions. Many parallel file system implementations rely on clustering multiple servers together for greater aggregate performance. Individual servers are much more powerful than traditional CPU-based architectures, meaning fewer servers are needed in a given cluster to achieve the same level of performance. The difference with the Hitachi File System is the scale. The Hitachi File System also makes it possible to create a single, unified namespace across the entire cluster of file servers. This makes it appear as a single file system to all network hosts and is known as cluster namespace or CNS. CNS satisfies the most common scalability requirements, allowing network hosts to access data on any server in the cluster, regardless of physical location. The Hitachi File System moves data seamlessly between multiple cluster nodes with minimal impact to performance.

Advanced multitier storage mechanisms

The Hitachi File System provides policy-driven data migration mechanisms for transparent data migration between many storage tiers. Data has an assigned value (by age, data type, owner, etc.) so transparently relocating data to an applicable storage "tier" without requiring users and applications to be pointed to the new location is a key feature of the Hitachi File System. This ability to extend the Hitachi File System to external devices, including other Hitachi products (such as HCP), private or public cloud, and even third party, or foreign, file systems accessible from the servers via NFSv3. Such data migration mechanisms also allow for repurposing of existing storage devices as external storage tiers, lowering total costs and offering easier platform transition.

Effective data reduction

The Hitachi File System can reduce storage requirements by various mechanisms, thus making storage more cost effective. The Hitachi File System can de-duplicate primary data with minimal performance impact. In addition, it can move data that is not actively used to the private or public cloud, thus greatly reducing the amount of data that is retained in expensive storage. Using HNAS in conjunction with Hitachi Content Platform (HCP) or Hitachi Capacity Optimizer (HCO) makes it possible to "backup less", which results in significant saving in operational costs.

2

Hardware overview

This chapter provides a brief description of the hardware used in the Hitachi Unified Storage VM storage system. Detailed information is located in the *Hitachi Unified Storage VM Block Module Hardware User Guide.*

NOTE: The Hitachi Unified Storage VM storage system may be installed only by Hitachi Data Systems service personnel or authorized service partners. However some high level installation information is included in this manual and in the *Hitachi Unified Storage VM Block Module Hardware User Guide* for reference.

- ☐ Block module storage system
- ☐ File module storage system

Block module storage system

This section briefly describes the Unified Storage VM block module storage system.

The HUS VM block module storage system can be configured in many ways, starting with a one-rack, diskless system to a multi-rack system that includes the following:

- A 5U controller chassis
- Up to 48 standard 2U 24-drive boxes
- Up to 24 4U 48-drive dense drive boxes
- Up to eight 2U 12-drive flash module drive boxes

The system supports a maximum of 1,152 SFF, LFF, and FMD data drives or a combination of the two. The system supports up to 128 SSD drives and up to 96 flash module drives.

The HUS VM block module storage system controller chassis contains the control logic, processors, memory, and interfaces to the drive boxes and the host servers.

The HUS VM block module storage system supports multiple cache configurations, with a maximum of 256 GB. The system provides a highly granular upgrade path, allowing the addition of disk drives to the drive chassis, and both main (cache) and processor blades as needed to increase system capacity and performance as storage needs increase.

All system components are mounted in either a standard Hitachi Data Systems 19-inch, 42U rack, or a customer-supplied rack that meets the rack specifications listed in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

The following sections provide descriptions and illustrations of the HUS VM block module storage system and its components.

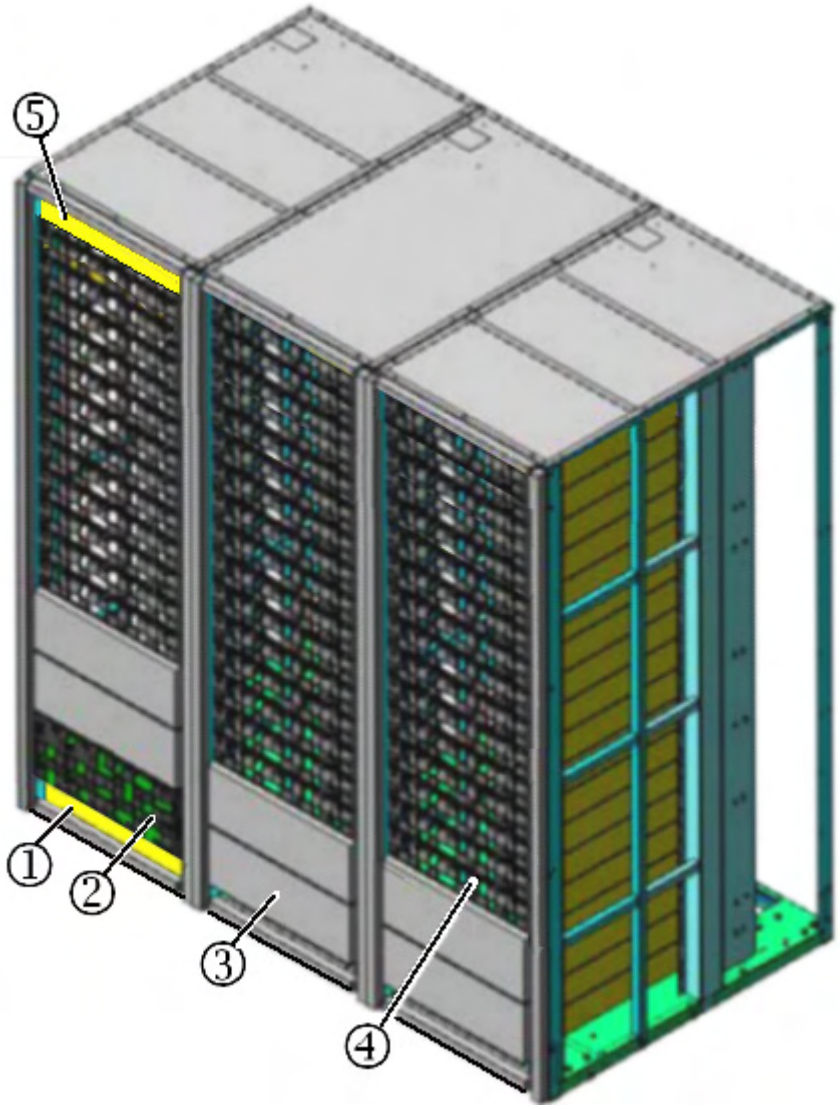

Figure 2-1 HUS VM block module storage system

Item	Description	Item	Description
1	Blank space for installation lifter (2U)	2	Controller chassis
3	DBX (dense) drive box (48 LFF drives)	4	DBS drive box (24 SFF drives) or DBL drive box (12 LFF drives)
5	Blank space (1U)		

The following illustration shows a basic illustration of a controller, a standard drive box, and a high-density drive box. They are all described briefly in this chapter and in detail in the *Hitachi Unified Storage VM Block Module Hardware User Guide.*

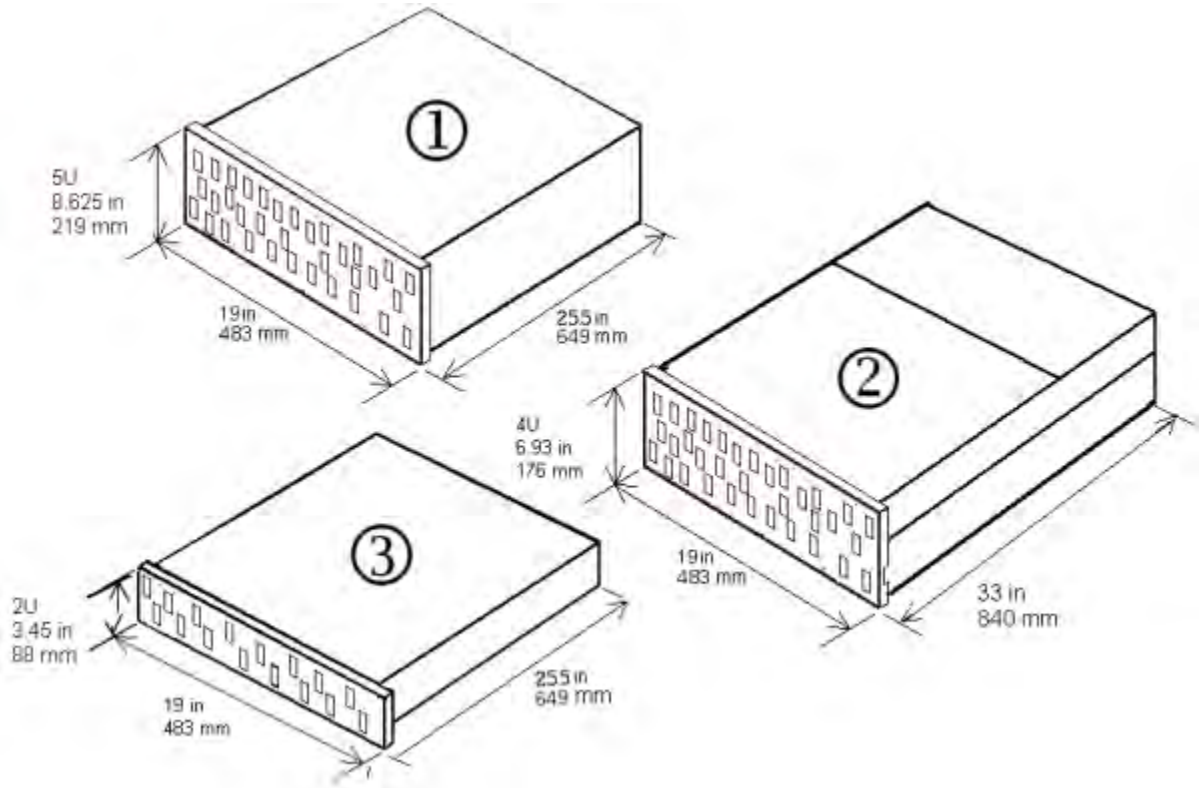

Figure 2-2 HUS VM block module storage system components

Item	Description	Item	Description
1	Controller Chassis	2	Dense Drive Box
3	Standard 2U SFF, LFF, or FMD Drive Box		

Controller chassis

The controller chassis (factory designation DKC) includes the logical components, memory, disk drive interfaces, and host interfaces. It can be expanded with a high degree of granularity to a system offering up to twice the number of processors, cache capacity, host interfaces and disk storage capacity. It is mounted at the bottom of the rack because it is heavier than the drive boxes.

The following illustration shows the front and rear views of the controller chassis.

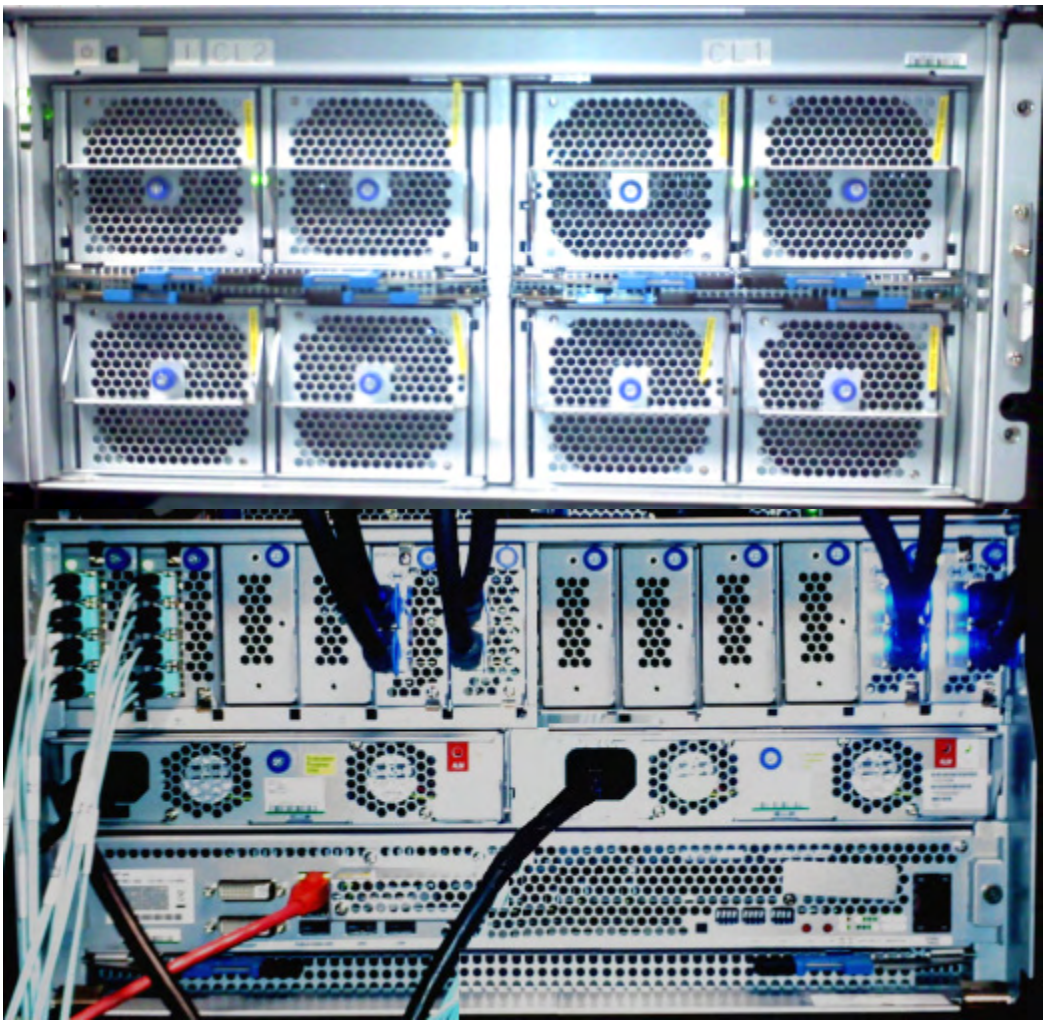

Figure 2-3 Controller chassis, front (above) and rear (below)

Drive boxes

The Unified Storage VM storage system can be configured in many ways, starting with a one-rack, diskless system to a large, three-rack system that includes a controller chassis and up to 48 standard drive boxes or up to 24 dense drive boxes. Standard and dense drive boxes can be intermixed in a Unified Storage VM system to meet specific storage requirements. More detail about the drive boxes is provided in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

- **DBS (D**rive **B**ox, **S**mall form factor drives). This drive box can contain up to 24 vertically mounted. 2-1/2-inch HDD or SDD drives.
- **DBL (D**rive **B**ox, **L**arge form factor drives). This drive box can contain up to 12 horizontally mounted. 3-1/2-inch HDD or SDD drives.
- **DBF (D**rive **B**ox, **FMD** drives). This drive box can contain up to 12 horizontally mounted. 5-1/4-inch Flash module drives.

- The **DBX** dense drive box is a 4U chassis that contains two identical 24-drive units. Each unit contains a maximum of 24 vertically-mounted (from the top of the drive box) LFF HDD or SSD drives. Up to 24 dense drive boxes can be mounted in a Unified Storage VM system. The standard and dense drive boxes are shown in the following figure.

Figure 2-4 Standard drive boxes (front bezels removed)

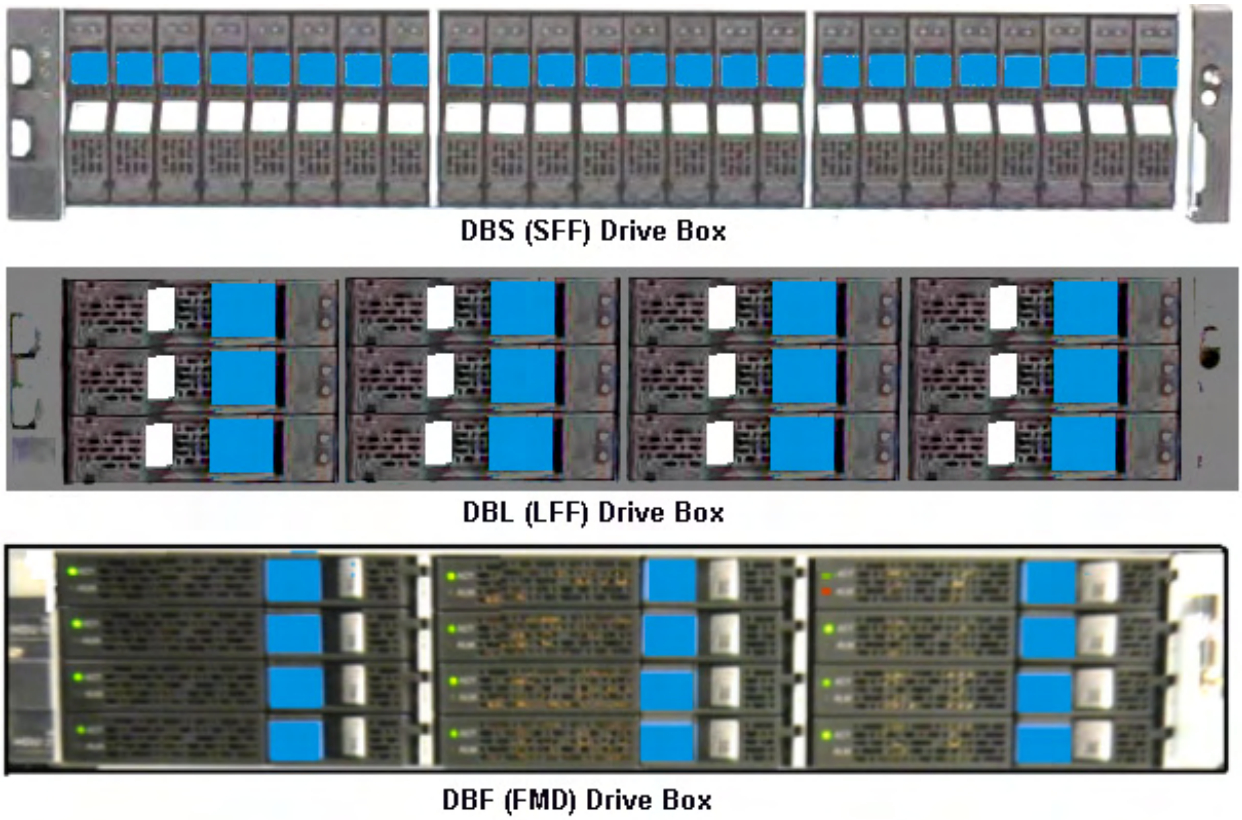

Figure 2-5 Dense drive box (front bezel removed)

Specifications

The following tables provide general specifications of the Unified Storage VM system. Additional specifications are located in Appendix A, Specifications, in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

Table 2-1 HUS VM system specifications

Item	Specifications	
System	Maximum Storage Capacity	3,383TB (3TB-SAS HDD used)
	Number of Disk Drives	Min: 4 (disk-in model) 0 (diskless model) Max: 1,152 (2-1.2-in and 3-1/2 in)
	Maximum Number of Flash Drives	128[1]
	Maximum Number of Flash Module Drives	96
	Maximum Number of Spare Drives	64
	Maximum Number of LDEVs / volumes	16,384
	Supported RAID Levels	RAID1, RAID 5, RAID 6
	RAID Group Configuration	RAID1: 2D+2D, 4D+4D RAID5: 3D+1P, 7D+1P RAID6: 6D+2P,14D+2P
	Internal Path	Architecture: Hierarchical Star Net
		Maximum Bandwidth: Cache Path = 128 GB/s Control Path = 64 GB/s
	Back-end Path	SAS 6G: 32 (2WL*6)
Memory	Cache memory capacity	32 GB, 64 GB, 96 GB, 128 GB
	Cache flash memory capacity	160 GB
Number of ports per installation unit	FC 2/4/8 GB	80 (96 *1)/16,8
Device I/F	Controller chassis-drive chassis Interface	SAS/Dual Port
	Data transfer rate	Max. 6 GBps
	Maximum number of HDD per SAS I/F	144
	Number of DKB PCB	4
Channel I/F	Open-systems	2/4/8 GBps Fibre Shortwave: 8UFC/16UFC
	Data Transfer Rate (MB/s, Fibre Channel)	200/400/800
Supported drives	See Table 2-2 Drive Specifications on page 2-8.	
Management Processor Cores	Quantity	16 cores

Item	Specifications	
MP configuration Minimum/maximum	host I/O modules	2 min, 4 with drives, 8 with no drives max[2]
	back-end I/O modules	0 or 2 / 4[3]
	Cache	2 / 8
	Switches /CSW	2 / 4
Power requirements (see Appendix A, Specifications, in the *Hitachi Unified Storage VM Block Module Hardware User Guide* for details)	Single phase storage system components. Single or three phase PDU input.	60Hz : 200V to 240V 50Hz : 200V to 240V
Acoustic level	Operating, Controller Chassis	60dB (24ºC or less), 62dB (32ºC)
	Operating, DBS / DBX	54dB (24ºC or less), 60dB (32ºC)
	Standby, Controller Chassis	60dB (24ºC or less), 62dB (32ºC)
	Standby, DBS / DBX	54dB
Notes: 1. When 300 GB is mounted 2. All FED configuration, no BEDs (diskless system) 3. Zero BEDs in a diskless configuration, two BEDs min if drives are installed		

Table 2-2 Drive Specifications

Drive Type	Size (inches)[1]	Drive Capacity	Speed (RPM)
HDD (SAS)	3-1/2	3 TB	7,200
	2-1/2	600 GB, 900 GB, 1.2 TB	10,000
n/a	SSD (Flash) MLC SAS 2	2-1/2	200 GB, 400 GB
Flash Module (MLC)	5-1/4	1.6 TB, 3.2 TB	n/a
Minimum Number of Drives			
Four (two in upper half, two in lower half). Drives must be added four at a time to create RAID groups, unless they are spare drives.			
Maximum Number of Drives			
Drive Type (inches)	**Drive Box**	Max per drive box	Max Per system
HDD, 3-1/2	DBX (dense)	48	1,152
HDD, 2-1/2	DBS	24	1.152
HDD, 3-1/2	DBL	12	1.152
FMD, 5-1/4	DBF	12	96
SSD, 2-1/2	128[2]	24 (DBS) or 48 (DBX)	128[3]
Spare drives[4]			64

Drive Type	Size (inches)[1]	Drive Capacity	Speed (RPM)
Notes. 1. the dense drive box uses only 3-1/2 in. drives. The DBL (LFF) drive box uses 3-1/2 in. drives. The DBS (SFF) drive box uses 2-1/2 in. drives.. The DBF drive box uses only 5-1/4 in. flash module drives. 2. SSD drives can be mounted all in one drive box or spread out among all of the drive boxes in the storage system. 3. Recommended maximum number. 4. Recommended number of spare drives: one spare drive per set of 32 HDDs and one per set of 32 SSDs.			

File module storage system

The HUS VM file module (also called the Hitachi NAS system) consists of the following hardware units.

- One or more servers which contain the control logic, processors, memory, and interfaces to the HNAS data drives
- One or more system management units

All system components are mounted in either a standard Hitachi Data Systems 19-inch, 42U rack, or a customer-supplied rack that meets the rack specifications listed in the *Hitachi Unified Storage VM Block Module Hardware User Guide*.

The following sections provide descriptions and illustrations of the HUS VM block module storage system and its components.

File module server

The server includes the logical components, memory, fibre channel interfaces, and Ethernet interfaces. The server contains dual mirrored data drives to store file module system data, a battery pack that, in the event of a power outage, preserves data that has not been written to disk for up to 72 hours, and fans to keep the server cool and operating efficiently. Typically, the file module is placed in the middle or upper portion of a rack, because fully populated drive boxes are significantly heavier than the complete file module chassis.

The following illustration shows the front and rear views of the file module server.

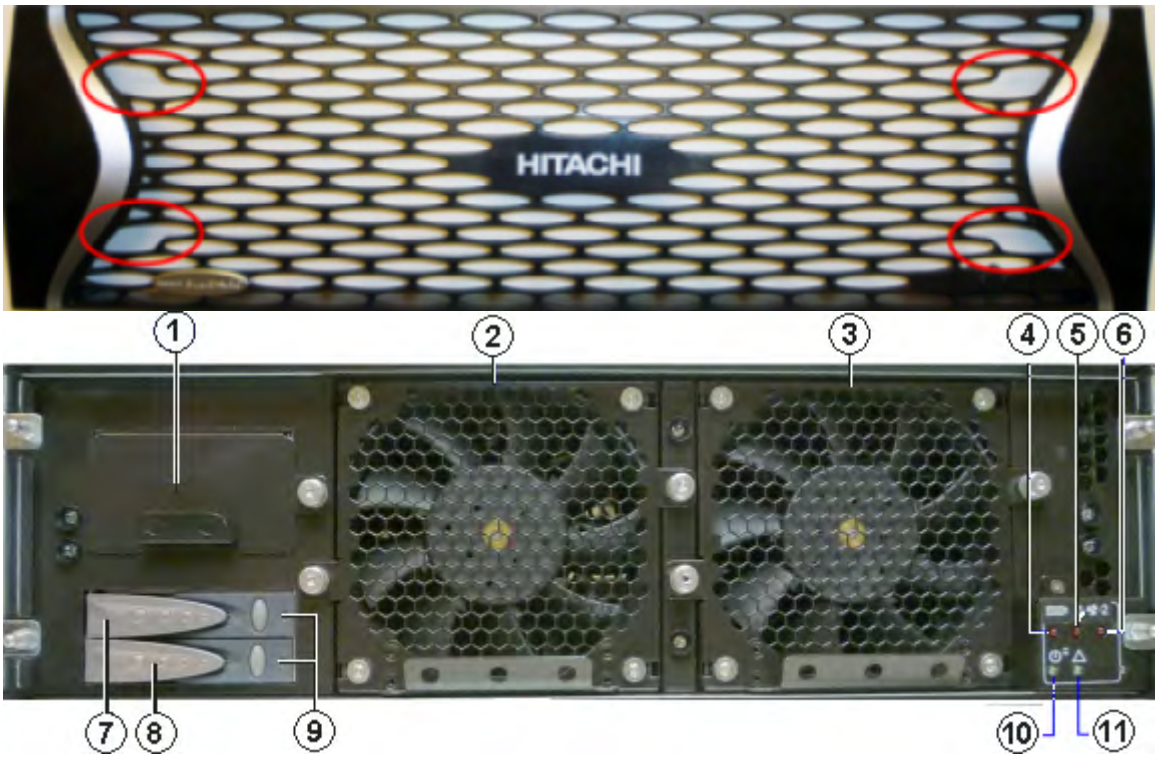

Figure 2-6 File module server, front view

Item	Description	Item	Description	Item	Description
1	NVRAM backup battery	2	Fan 1	3	Fan 2
4	NVRAM battery backup pack status LED	5	Fan 1 status LED	6	Fan 2 status LED
7	Hot-swappable power supplies	8	Keyboard (purple) and mouse (green) ports	9	USB ports
10	Serial port (RS-232)	11	Monitor port		

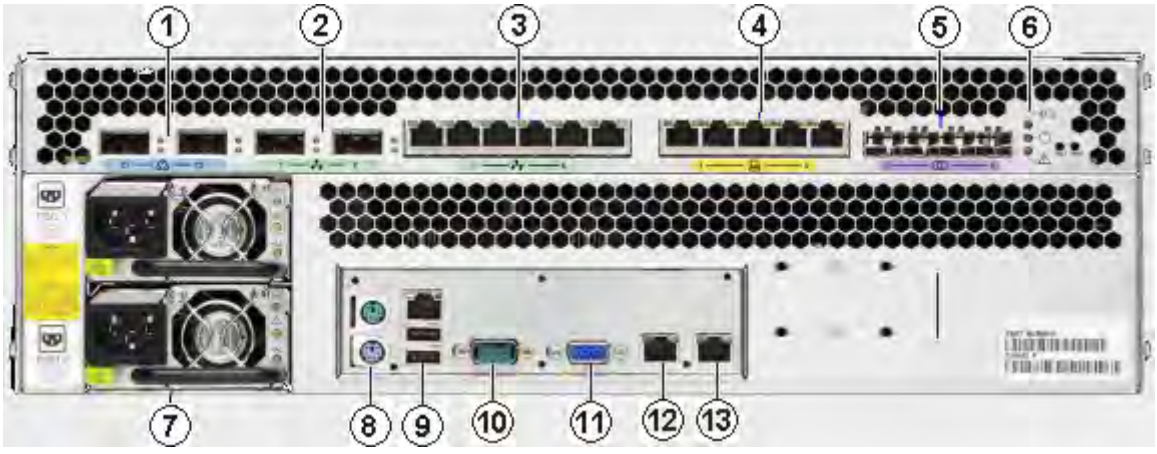

Figure 2-7 File module server, rear view

Item	Description	Item	Description	Item	Description
1	10GbE cluster interconnect ports	2	10GbE public (data) network ports	3	GE ports
4	10/100 Ethernet ports	5	1/2/4Gbps FC ports	6	NVRAM, power, and server status LEDs, and Power and Reset buttons
7	Hot-swappable power supplies	8	Keyboard (purple) and mouse (green) ports	9	USB ports
10	Serial port (RS-232)	11	Monitor port	12	Ethernet management port 0 (Eth0)
13	Ethernet management port 1 (Eth1)				

System management unit

The system management unit (SMU) in the file module (HNAS) system is equivalent to the service processor in the block module system. It provides administration and monitoring tools, supports data migration and replication, and acts as a quorum device in a cluster configuration. Although integral to the system, the SMU does not move data between the network client and the servers.

Note: The server includes an embedded SMU, which can manage a single file module server. However, embedded SMUs are not active when an external SMU is attached. All cluster configurations require an external SMU.

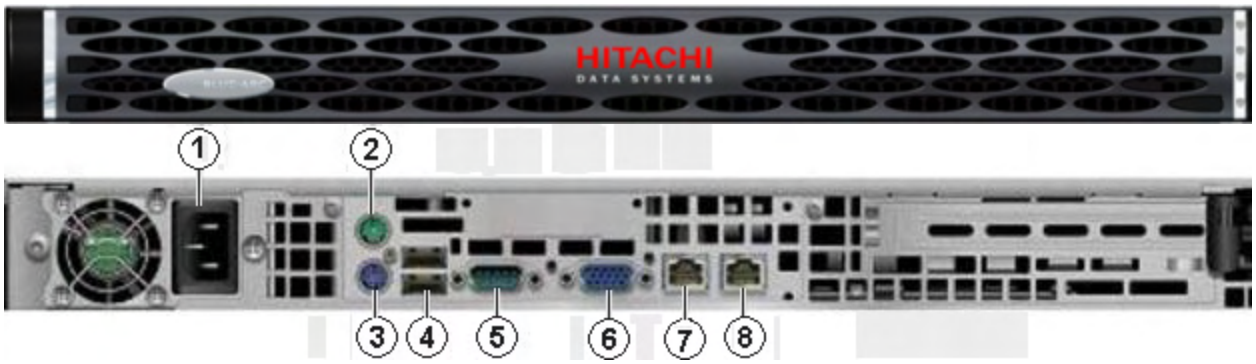

Figure 2-8 System management server

Item	Description	Item	Description	Item	Description
1	Power supply	2	Mouse port	3	Keyboard port
4	USB ports (2)	5	Serial connection	6	Video connection
7	Eth0 port (public network)	8	Eth1 port (private network)		

Fibre channel (FC) switches

The HUS VM file module supports Fibre Channel switches that connect servers to storage subsystems. Contact Hitachi Data Systems Support Center for information about which Fibre Channel switches are supported.

External fast ethernet (10/100) or GigE switches

A standalone HUS VM file module can operate without an external Ethernet switch, as long as it uses an internal SMU and there are less than three RAID subsystems attached.

A standalone file module server requires an external Ethernet switch if there are three or more RAID subsystems attached or if there are two RAID subsystems attached and an external SMU is used. All cluster configurations require an external Ethernet switch.

10 GbE (10 Gigabit Ethernet) switches (cluster configurations only)

The HUS VM file module server supports 10 GbE switches that connect multiple storage servers configured as a cluster. Contact your Hitachi Data Systems representative for information about the 10 GbE switches that have been qualified for use with the file module system, and to find out about the availability of those switches.

Currently, the only 10 GbE switch that has been qualified is the Brocade TurboIron 24X, a standalone 10Gbps Managed Layer 2 Ethernet switch with eight ports. It is available through Hitachi Data Systems. Hitachi Data Systems requires dual 10 GbE switches for redundancy. If one switch fails, the cluster nodes remain connected through the second switch.

3

Software Overview

This chapter provides a brief description of the software used in the Hitachi Unified Storage VM storage system.

NOTE: The Hitachi Unified Storage VM storage system includes basic software installed at the factory. Users can install additional software and licenses as needed, or can ask their Hitachi Data Systems authorized service partner to do it for them.

- ☐ Block module software features and functions
- ☐ File module storage system software features and functions

Block module software features and functions

The Unified Storage VM storage system provides advanced software features and functions that increase data accessibility and deliver enterprise-wide coverage of online data copy/relocation, data access/protection, and storage resource management. Hitachi Data Systems software products and solutions provide a full set of industry-leading copy, availability, resource management, and exchange software to support business continuity, database backup and restore, application testing, and data mining.

Virtualization

The following table describes the virtualization software that is available on the Unified Storage VM storage system.

Table 3-1 Virtualization Features and Functions

Feature	Description
Hitachi Virtual Partition Manager	Provides logical partitioning of the cache which allows you to divide the cache into multiple virtual cache memories to reduce I/O contention.
Hitachi Cache Residency Manager	Supports the virtualization of external storage systems. Users can connect other storage systems to the Unified Storage VM storage system and access the data on the external storage system via virtual devices created on the Unified Storage VM storage system. Functions such as TrueCopy and Cache Residency can be performed on external data through the virtual devices.

Performance Management

The following table describes the performance management software that is available on the Unified Storage VM storage system.

Table 3-2 Performance Management Features and Functions

Feature	Description
Hitachi Cache Residency Manager	Cache Residency Manager locks and unlocks data into the cache to optimize access to the most frequently used data. It makes data from specific logical units resident in a cache, making all data accesses become cache hits. When the function is applied to a logic unit, frequently accessed, throughput increases because all reads become cache hits.
Hitachi Performance Monitor	Performs detailed monitoring of the storage system and volume activity. This is a short tem function and does not provide historical data

Provisioning

The following table describes the provisioning software that is available on the Unified Storage VM storage system.

Table 3-3 Provisioning Features and Functions for Open Systems

Feature	Description
Dynamic Tiering	Provides automated support for a multi-tiered Dynamic Provisioning pool. The most accessed data within the pool is dynamically relocated onto the faster tiers in the pool. Data that is most referenced has improved performance due to the inclusion of fast storage such as SSD while controlling the overall storage cost by incorporating lower costing storage such as SATA.
Hitachi LUN Manager	The LUN Manager feature configures the fibre-channel ports and devices (logical units) for operational environments (for example, arbitrated-loop and fabric topologies, host failover support).
Hitachi LUN Expansion	The LUN Expansion feature expands the size of a logical unit (volume) to which an open-system host computer accesses by combining multiple logical units (volumes) internally.
Hitachi Dynamic Provisioning software	The Dynamic Provisioning feature virtualizes some or all of the system's physical storage. This simplifies administration and addition of storage, eliminates application service interruptions, and reduces costs. It also improves the capacity and efficiency of disk drives by assigning physical capacity on demand at the time of the write command receipt without assigning the physical capacity to logical units.
Hitachi Virtual LVI	Converts single volumes (logical volume images or logical units) into multiple smaller volumes to improve data access performance.
Hitachi Data Retention Utility	Protects data in logical units / volumes / LDEVs from I/O operations illegally performed by host systems. Users can assign an access attribute to each volume to restrict read and/or write operations, preventing unauthorized access to data.

Data Replication

The following table describes the Data Replication software that is available on the Unified Storage VM storage system.

Table 3-4 Data Replication Features and Functions

Feature	Description
Hitachi TrueCopy® Remote Replication software	Performs remote copy operations between storage systems at different locations.
Hitachi ShadowImage® In-System Replication software	Creates internal copies of volumes for purposes such as application testing and offline backup. Can be used in conjunction with True Copy or Universal Replicator to maintain multiple copies of data at primary and secondary sites.
Hitachi Copy-on-Write Snapshot software and Hitachi Thin Image (open systems only)	Snapshot creates a virtual, point-in-time copy of a data volume. Since only changed data blocks are stored in the Snapshot storage pool, storage capacity is substantially less than the source volume. This results in significant savings compared with full cloning methods. With Copy-on-Write Snapshot, you create virtual copies of a data volume in the Virtual Storage Platform. Thin Image can perform the cost-effective duplication by storing only differential data between primary volumes and secondary volumes of VVOLs.

Feature	Description
Hitachi Universal Replicator software	This feature provides a RAID storage-based hardware solution for disaster recovery which enables fast and accurate system recovery, particularly for large amounts of data which span multiple volumes. Using Universal Replicator, you can configure and manage highly reliable data replication systems using journal volumes to reduce chances of suspension of copy operations.

Security

The following table describes the security software that is available on the Unified Storage VM storage system.

Table 3-5 Security Features and Functions

Feature	Description
Encryption License Key	This feature implements encryption for open-systems data using the encrypting back-end I/O module. It includes enhanced key support up to 32 separate encryption keys allows encryption to be used as access control for multi-tenant environments. It also provides enhanced data security for the AES-XTS mode of operations.
External Authentication and Authorization	Storage management users of Unified Storage VM systems can be authenticated and authorized for storage management operations using existing customer infrastructure such as Microsoft Active Directory, LDAP, and RADIUS-based systems.
Role Based Access Control (RBAC)	Provides greater granularity and access control for Unified Storage VM storage administration. This new RBAC model separates storage, security, and maintenance functions within the array. Storage management users can receive their “role” assignments based on their group memberships in external authorization sources such as Microsoft Active Directory and LDAP. This RBAC model will also align with the RBAC implementation in HCS 7.
Resource Groups	Resource Groups allow for additional granularity and flexibility of the management of storage resources.

System Maintenance

The following tables describes the system maintenance software that is available on the Unified Storage VM storage system.

Table 3-6 System Maintenance Features and Functions

Feature	Description
Audit Log Function	The Audit Log function monitors all operations performed using Storage Navigator (and the SVP), generates a syslog, and outputs the syslog to the Storage Navigator computer.
Hitachi SNMP Agent	Provides support for SNMP monitoring and management. Includes Hitachi specific MIBs and enables SNMP-based reporting on status and alerts. SNMP agent on the SVP gathers usage and error information and transfers the information to the SNMP manager on the host.

Host Server software

The following tables describes the Host Server software that is available on the Unified Storage VM storage system.

Table 3-7 Host Server-based Features and Functions

Feature	Description
Hitachi Command Control Interface software	On open-systems, performs various functions, including data replication and data protection operations by issuing commands from the host to the Hitachi Data Systems storage systems. The CCI software supports scripting and provides failover and mutual hot standby functionality in cooperation with host failover products.
Dataset Replication	Operates with the ShadowImage feature. Rewrites the OS management information (VTOC, VVDS, and VTOCIX) and dataset name and creates a user catalog for a ShadowImage target volume after a split operation. Provides the prepare, volume divide, volume unify, and volume backup functions to enable use of a ShadowImage target volume.

File module storage system software features and functions

This section describes the Unified Storage VM storage system file module software.

Intelligent tiering

Intelligent tiering allows customers to build scalable and flexible storage solutions that offer cost-effective performance with simplified and consolidated storage management. Using the various tiers of storage available, customers can keep data on line longer without relying exclusively on tape technologies, minimizing the impact of backup, replication, or disaster-recovery requirements as the strategy requires. Intelligent tiering gives data a longer disk lifecycle if desired, which can improve data access times for hosts and users.

Data migrator

Data migrator is an embedded feature of the Hitachi File System. All file system functions (snapshots, replication, quotas, etc.) work seamlessly as if the data were still on the original storage tier. Data integrity is also maintained during the migration or recall. But the data appears as if it has not moved at all. Users and applications see the data as if it still exists in the original location. The Hitachi File System contains the metadata that indicates where the data actually resides.

Cross-volume links

Cross-volume links (CVL) and external cross-volume links (XVL) are complementary technologies that extend the reach of data migrator. A cross-volume link is a zero-length file on a source file system (the primary file system) which "points" at a corresponding file on a target file system (the secondary file system). The pointer is stored in the onode of the

primary file. A flag in the onode is used to indicate it is a cross-volume link rather than a regular file, and an extended onode contains the information required to access the migrated file. All of the metadata required for directory level operations (including owner, access mode and ACLs) are maintained on the primary file system, so operations such as "ls -l" or "chmod" do not require access to the secondary file system. Similarly, the information needed for quota tracking is maintained on the primary file system, so quotas reported will include migrated files on the secondary file system as well.

Dynamic caching

Dynamic caching reserves space on a storage tier for caching of "hot" files accessed through NFS. The space reserved is actually an entire file system unto itself, and as such can be as large as any other file system in the namespace. By definition, any file that is recently accessed may have a copy also located in the dynamic cache. If the cache is created in a high-performance tier of storage, this copy guarantees that any hot files are automatically on the highest performance disk tier, which may actually be an SSD or a hybrid SDRAM/SSD tier. Having the cache prevents the need for reverse data migration.

Cluster read caching

Cluster read caching is dynamic caching applied to a single server or to a cluster of servers. If applied to a single server, the feature is called local read caching. When used with a cluster of servers, each server maintains its own dynamic cache, but is aware of the files accessed by all the other servers in the cluster. Copies of hot files from anywhere in the cluster therefore make their way to every cache on every server, which can result in dramatic aggregate read performance improvements. This is because every server can respond to any read request for a given set of hot files. In this way, dynamic caching works with the data migrator to simultaneously provide policy-based data movement in both the forward and reverse senses.

Data relocation

The file module system allows you to relocate date in three ways:

- **EVS migration** is typically used for load balancing among servers or cluster nodes, and vacating a server for scheduled maintenance. It makes it possible to relocate a virtual server within a cluster, or to a server outside of the cluster that shares access to the same storage devices. EVS migration has minimal impact on network hosts and, once it has completed, those hosts may access the data using the same path names that were in use prior to the relocation.
- **File system relocation**: Any file system accessed via cluster namespace can be relocated to another server within the cluster. File system relocation has minimal impact on network hosts; once completed the data is accessed using the same path names that were in use prior to the relocation. File system relocation is typically used to load balance within the unified namespace.

- **Data relocation:** Data may be relocated from any given file system to another using a mechanism referred to as transfer of primary access (TPA). TPA makes it possible to relocate individual directories as well as entire file systems. TPA does, however, involve a small amount of downtime, and data is no longer accessible using the same path names that were in use prior to the relocation. TPA is generally used to better organize file systems and/or directories within them.

Data Protection

The Hitachi File System also offers features for online data recovery, data replication, mirroring, backup, disaster recovery, and complete system monitoring.

Hitachi NAS solutions support high availability (HA) clustering of servers in a 2-node active-active configuration or an n-way (more than 2 nodes) cluster configuration. Clustered servers provide NVRAM mirroring for enhanced data protection, automated file system failover, and higher levels of performance as additional servers are added to the cluster.

In addition, the Hitachi File System provides additional mechanisms for data protection. Three of the more important mechanisms are snapshots, data replication, and data backup. Snapshots are generally described as point-in-time copies of the file system, and are a very convenient way to give end-users a way to "roll back" to a previous point in time to recover their own data. There are several data replication options within the Hitachi NAS system software; these may be described as either file, object or block-based, and synchronous or asynchronous.

High availability

The Hitachi File System high availability (HA) design is the concept of enterprise virtual servers (EVS). Virtual servers are logical entities that reside on a physical server, similar to operating system virtualization techniques such as VMware, Microsoft Hyper-V, or Citrix XenServer. An EVS does not have physical interfaces, but instead has virtual interfaces that map to the physical interfaces of the server. As a result of the separation between the physical and logical interfaces of the EVS and the physical server, an EVS can be migrated from one physical server to another, transparently. In a failover situation for an HA cluster, automated EVS migration process takes place without system shutdown, and in most cases occurs quickly enough that hosts using stateless protocols (such as NFSv3) will not require unmounting and remounting of NFS exports.

Snapshots

Snapshots allow the storage administrator to capture a point-in-time image of the file system and the point-in-time image is a read-only view of the file system. Using point-in-time images (snapshots), the storage administrator can:

- Allow end-users to retrieve files that have been deleted without administrator intervention

- Perform backups of the file system from a snapshot instead of using the live file system

Snapshots are rule-based, with the flexibility to define when they are taken based on business policies. Rules-based snapshots provide entity management, a more useful configuration compared to simpler volume-based snapshot management. For example, hourly snapshot rules are managed as one entity, daily or weekly rules are managed as a separate entity, and monthly rules are managed as a 3rd separate entity. There is also an implied hierarchy to snapshot rules: an hourly snapshot will not overwrite a daily or weekly snapshot, etc. An hourly snapshot will overwrite only another hourly snapshot, a daily or weekly snapshot will only overwrite a snapshot taken because of a daily or weekly rule, etc.

Snapshots can be set up in four ways:

- **Automatically**, via a prescribed rule
- **Manually**, using the servers' GUI or CLI management interfaces
- **Scripted**, using a script, the storage administrator may automate the beginning of the snapshot process, in a manner similar to the rules-based method described above
- **Event based,** using scripts and the remote scripting tool, the storage administrator may automate the taking of a snapshot based on trigger events generated by the server

Space efficient clones

File clones provide the ability to instantly create space-efficient, writeable copies of single files. This is a key feature for virtualization, as it allows VMware administrators to quickly deploy new virtual machines without consuming additional disk space.

When a file clone is created, a hidden file-level and read-only snapshot is created which preserves an image of the source file at that point in time. This snapshot allows the source file and clone to share common data blocks. If additional clones of the source file are created, new snapshots are created if necessary.

The key benefits of file clones include:

- **Space efficiency**: File clones are initially created through the use of pointers to blocks. No data is duplicated, and only new file writes are saved to the file system.
- **Speed**: File clones are created in seconds, regardless of the size of the file being cloned.
- **Flexibility**: File clones are created per file, allowing for a VMware virtual machine focused view of file cloning.

Specific features of file clones include writable snapshots, clone of a clone (or cascading clones), deep copy, and unlimited number of file clones.

Replication features

Replication is the process of storing data in redundant sources, as a method to ensure data consistency, and to improve reliability and accessibility of the entire system. Replication differs from backup in that replication aims to have the data in two or more places at once (theoretically, identical copies of the data in all locations at the same time), while backup aims to have two or more copies of the data at different points in time. Despite these differences, there are many common design components in replication and backup, as both are designed with data movement in mind. The Hitachi File System provides several robust mechanisms for data movement, many of which are useful for replication scenarios.

- **Accelerated data copy** (ADC) is a file-based, asynchronous method of data replication. ADC allows the storage administrator to define a policy-based data migration, or a mass data migration to occur either among or between servers.
- **Incremental data replication** (IDR) is an optional replication feature of the Hitachi File System. Replication occurs at the file level, and only includes files that have changed since the last scheduled replication. Multiple schedules may be defined on a per-EVS basis with support for pre- and post-scripting, enabling automated functions to occur prior to and immediately after the IDR schedule.
-

Object replication

The Hitachi File System is object-based, which means core file system structures and user data are stored as objects, rather than files or blocks. Object replication takes advantage of this structure to replicate these blocks natively, regardless of which file or directory that they may belong to, which negates the need to assemble all of the objects associated with a file before transfer, making the overall transfer more efficient.

The key benefits of object-based replication include:

- **Replication performance -** Full replication performance will generally improve 2-3x depending on the structure of the file system. The greatest performance improvements will be seen in the incremental replication, especially for customers that have dense file systems (millions of small files) or those that have a high rate of change in the file system.
- **Improved disaster recovery** - Object replication enables the ability to quickly failover in the event of a disaster.
- **Disaster recovery** (DR) or backup - Customers using replication for disaster recovery and business continuance will see major benefits from object replication and the disaster recovery features that it enables.
- **Dense file systems** - file systems with millions of small files will realize a major benefit for the time to start and complete incremental replication. Even if the change rate is low, object replication will start much faster due to the time difference required for change discovery, and complete faster with its ability to fill the pipeline.

Block replication and SyncDR

Depending on the disk technology being used, Hitachi NAS system software also supports specific block-based replication options available at the Hitachi storage controller level (synchronous or asynchronous). Dark fiber connections between controller pairs are required for controller mirroring; depending on class of hardware and features selected, replication distances can range from 500m to 10km. Through the use of advanced fiber-optic wavelength-division multiplexing technologies, this range can be considerably extended to 100km or more.

Backup features

Backup services seek to achieve a copy of the data at a specific point in time, usually for off-line preservation. As with replication, the Hitachi File System leverages snapshots to allow the storage administrator to perform backups while continuing to serve data to hosts with the live file system.

The Hitachi NAS server supports LAN-free backup of data using either Fibre Channel or Ethernet networks and dedicated connections with dedicated bandwidth. This separation allows the live file system to continue serving data to hosts unimpeded by the backup operations.

4

Setup and configuration overview

This chapter provides an ordered overview of the tasks that you and your Hitachi Data Systems authorized service provider will need to complete to plan and prepare the installation site, and to install and configure the system.

This chapter does not contain the actual instructions to set up and configure the site and the system. It includes pointers to the Unified Storage VM user guides that provide the instructions for these tasks. Note that only Hitachi Data Systems employees or authorized service providers are allowed to install and configure the hardware.

- ☐ Required setup tasks
- ☐ Optional setup tasks

Required setup tasks

The following tables show ordered lists of the tasks that must be completed to set up a Unified Storage VM storage system.

General task list

The following table provides a high-level list of installation and configuration tasks. The tables following it provide more detailed task lists. See the information following this table or the related user guides for detailed task lists and instructions.

Task	Block Module Manual	File Module Manual
1. Prepare the site	1. Site preparation tasks on page 4-3	*Hitachi Unified Storage File Module System Installation Guide* (technical support only)
2. Turn the storage system power on	2. Startup tasks on page 4-4	• *Hitachi Unified Storage File Module Hardware Reference* • *Hitachi Unified Storage File Module System Installation Guide (technical support only)*
3. Log in to the system	3. Basic system admin tasks on page 4-4	*Hitachi Unified Storage File Module Hardware Reference*
4. Log in to the system	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide*	*Hitachi Unified Storage File Module Storage Systems User Administration Guide*
5. Set up user accounts, enable software licenses	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide*	*Hitachi Unified Storage File Module Storage Systems User Administration Guide*
6. Provision and configure the system	4. Provisioning and configuration tasks on page 4-5	*Hitachi Unified Storage File Module Storage Subsystem Administration Guide*
7. Set up virtualization	5. Storage virtualization tasks on page 4-6	• *Hitachi Unified Storage File Module File Services Administration Guide* • *Hitachi Unified Storage File Module Storage Subsystem Administration Guide*
8. Set up replication	6. Data replication and disaster recovery tasks on page 4-6	• *Hitachi Unified Storage File Module Network Administration Guide* • *Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide* • *Hitachi Unified Storage File Module Snapshot Administration Guide*

Task	Block Module Manual	File Module Manual
9. Set up monitoring, logs, and alerts	• *Hitachi Unified Storage VM Block Module Hitachi Audit Log User Guide* • *Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide*	• *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* • *Hitachi Unified Storage File Module Server and Cluster Administration Guide*
10. Tune system for performance	*Hitachi Unified Storage VM Block Module Performance Guide*	*Hitachi Unified Storage File Module Server and Cluster Administration Guide*
11. Troubleshoot and maintain the system	• *Hitachi Unified Storage VM Block Module Hitachi Storage Navigator Messages* • *Hitachi Unified Storage VM Block Module Hardware User Guide*	• *Hitachi Unified Storage File Module Antivirus Administration Guide* • *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* • *Hitachi Unified Storage File Module Server and Cluster Administration Guide* • *Hitachi Unified Storage File Module Hardware Reference*

1. Site preparation tasks

Site preparation tasks	Block Module	File Module
1. Review the entire installation procedure before beginning work. 2. Familiarize yourself with the system hardware 3. Plan the hardware installation 4. Prepare the facility as needed to meet the power and environmental requirements. 5. Prepare the location where the system will be installed 6. Assist service personnel as needed to install the system. 7. Connect the system to a host (block module only) 8. Connect the system to the corporate network (file module only)	*Hitachi Unified Storage VM Block Module Hardware User Guide*	• *Hitachi Unified Storage File Module System Installation Guide* (technical support only)

2. Startup tasks

Storage Management Task	Block Module	File Module
1. Verify that power to the rack is correct and turned on 2. Verify that all PDU breakers and switches are ON. 3. Turn on system power 4. Verify that the ready LED is ON.	*Hitachi Unified Storage VM Block Module Hardware User Guide*	*Hitachi Unified Storage File Module Hardware Reference*

3. Basic system admin tasks

Storage Management Task	Block Module	File Module
1. Log in 2. Change admin password 3. Set up user accounts 4. Enable license keys	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide*	• *Hitachi Unified Storage File Module System Access Guide* • *Hitachi Unified Storage File Module Storage System User Administration Guide* • *Hitachi Unified Storage File Module Server and Cluster Administration Guide*

4. Provisioning and configuration tasks

Storage Management Task	Block Module	File Module
Block Module Provision the Unified Storage VM system for use with open systems. 1. Verify that the system on which you are installing the software meets the minimum requirements for the software release. 2. Prepare the fibre-channel HBAs for the installation. 3. Verify recognition of the new devices. 4. Verify device files and the driver. 5. Create open-systems volumes 6. Partition disk devices, create file systems, and set device parameters. 7. Create mount directories, mount and verify the file systems, and set and verify auto-mount parameters. 8. Configure hosts and ports. 9. Define paths to volumes. 10. Assign access attribute to volumes. 11. Implement dynamic allocation of disk space on an as-needed basis.	• *Hitachi Unified Storage VM Block Module Provisioning Guide* • *Hitachi Unified Storage VM Hitachi System Operations Using Spreadsheets*	• *n/a*
File Module 1. Perform basic system configuration 2. Create system drives 3. Create storage pools 4. Create file systems 5. Create mount points (CIFS shares or NFS exports)	• *n/a*	• *Hitachi Unified Storage File Module System Installation Guide* (technical support only) • *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* • *Hitachi Unified Storage File Module File Services Administration Guide*

5. Storage virtualization tasks

Storage Management Task	Block Module	File Module
1. Connect external storage systems to the HUS VM storage system for storage virtualization. 2. Set up virtual volumes.	• *Hitachi Unified Storage VM Block Module Hitachi Universal Volume Manager User Guide* • *Hitachi Unified Storage VM Block Module Performance Guide* • *Hitachi Unified Storage VM Hitachi System Operations Using Spreadsheets*	• *Hitachi Unified Storage File Module File Services Administration Guide*

6. Data replication and disaster recovery tasks

Storage Management Task	Block Module	File Module
Perform synchronous remote replication and disaster recovery.	*Hitachi Unified Storage VM Block Module Hitachi TrueCopy® User Guide*	• *Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide* • *Hitachi Unified Storage File Module Snapshot Administration Guide*
Perform asynchronous remote replication and disaster recovery.	*Hitachi Unified Storage VM Block Module Hitachi Universal Replicator User Guide*	• *Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide* • *Hitachi Unified Storage File Module Snapshot Administration Guide*
Perform local replication.	*Hitachi Unified Storage VM Block Module Hitachi ShadowImage® User Guide*, *Hitachi Unified Storage VM Block Module Thin Image User Guide*	• *Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide*
Install Command Control Interface (CCI).	*Hitachi Command Control Interface Installation and Configuration Guide*	n/a
Use CCI commands to perform data replication.	*Hitachi Command Control Interface User and Reference Guide*	n/a

Optional setup tasks

The following tables list the optional setup tasks that you can use to customize your storage system.

1. Provisioning and configuration

Storage Management Task	Block Module	File Module
Use SNMP-based reporting on status and alerts for the storage system	*Hitachi Unified Storage VM Block Module Hitachi SNMP Agent User Guide*	*Hitachi Unified Storage File Module Server and Cluster Administration Guide*
(Security) Delete data from volumes to prevent it from being restored	*Hitachi Unified Storage VM Block Module Hitachi Volume Shredder User Guide*	n/a

2. Performance monitoring and tuning tasks

Storage Management Task	Block Module	File Module
Monitor performance statistics of disk drives, volumes, and other storage components	*Hitachi Unified Storage VM Block Module Performance Guide*	• *Hitachi Unified Storage File Module Server and Cluster Administration Guide*
Improve data access speed using cache		n/a
Create and managing virtual cache partitions		n/a

3. Troubleshooting and other tasks

Storage Management Task	Block Module	File Module
Review information about GUI messages	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator Messages*	n/a
Review the log of operations to investigate causes of errors and avoid potential errors	*Hitachi Unified Storage VM Block Module Hitachi Audit Log User Guide*	*Hitachi NAS Platform Troubleshooting Guide* (Technical Support only)
Review syntax of CCI commands	*Hitachi Command Control Interface Command Reference*	Hitachi NAS Platform Command Line Reference (accessed through HNAS Web Manager)

A

Hitachi Unified Storage VM user guides

This section provides a list of all Hitachi Unified Storage VM user documentation, by topic. The Hitachi Unified Storage VM system includes both the block module and the file module systems.

- Unified Storage VM user guides by topic
- Unified Storage VM user guides by number

Unified Storage VM user guides by topic

The following table provides a list of all Unified Storage VM user documentation, by topic.

Type of manual / Topic	Block Module Manual Name	File Module Manual Name
Introductory	*Hitachi Unified Storage VM Product Overview Guide* MK-92HM7003	Included in *Hitachi Unified Storage VM Product Overview Guide* MK-92HM7003
Hardware	*Hitachi Unified Storage VM Block Module Hardware User Guide* MK-92HM7005	*Hitachi Unified Storage File Module Hardware Reference* MK-92USF001
System Configuration and management	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide* MK-92HM7016	• *Hitachi Unified Storage File Module System Installation Guide* (technical support only) • *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* MK-92USF011 • *Hitachi Unified Storage File Module File Services Administration Guide* MK-92USF004
Storage Management	• *Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide* MK-92HM7016 • *Hitachi Unified Storage VM Hitachi System Operations Using Spreadsheets* MK-92HM7015	• *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* MK-92USF006 • *Hitachi Unified Storage File Module Server and Cluster Administration Guide* MK-92USF007 • *Hitachi Unified Storage File Module NDMP Backup Administration Guide* MK-92USF012
Security	*Hitachi Unified Storage VM Block Module Hitachi Volume Shredder User Guide* MK-92HM7021	*Hitachi Unified Storage File Module Antivirus Administration Guide* MK-92USF010
Provisioning	• *Hitachi Unified Storage VM Block Module Provisioning Guide* MK-92HM7012 Covers tiered storage, dynamic provisioning, LUN manager, LUSE, virtual LUN, and data retention. • *Hitachi Unified Storage VM Hitachi System Operations Using Spreadsheets* MK-92HM7015	• *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* MK-92USF006 • *Hitachi Unified Storage File Module Server and Cluster Administration Guide MK-92USF007* • *Hitachi Unified Storage File Module System Access Guide* MK-92USF002

Type of manual / Topic	Block Module Manual Name	File Module Manual Name
Host Configuration	• *Configuration Guide for IBM® AIX® Host Attachment* MK-96RD636 • Configuration Guide for HP-UX® Host Attachment MK-96RD638 • *Configuration Guide for Microsoft® Windows® Host Attachment* MK-96RD639 • *Configuration Guide for Red Hat® Linux® Host Attachment* MK-96RD640 • *Configuration Guide for SuSE® Linux® Host Attachment* MK-96RD650 • *Configuration Guide for VMware™ ESX Server Host Attachment* MK-98RD6716 • Configuration Guide for Citrix® XenServer® Host Attachment MK-90RD6766	n/a
Virtualization	*Hitachi Unified Storage VM Block Module Hitachi Un Manager User Guide* MK-92HM7020	*Hitachi Unified Storage File Module File Services Administration Guide* MK-92USF7004
Replication	• *Hitachi Unified Storage VM Block Module Hitachi ShadowImage® User Guide* MK-92HM7013 • *Hitachi Unified Storage VM Block Module Hitachi TrueCopy® User Guide* MK-92HM7018 • *Hitachi Unified Storage VM Block Module Hitachi Universal Replicator User Guide* MK-92HM7019 • *Hitachi Unified Storage VM Block Module Thin Image User Guide* MK-92HM7010	• *Hitachi Unified Storage File Module Network Administration Guide* MK-92USF003 • *Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide* MK-92USF009 • *Hitachi Unified Storage File Module Snapshot Administration Guide* MK-92USF008
Command line entry	• *Hitachi Command Control Interface Command Reference* MK-90RD7009 • *Hitachi Command Control Interface User and Reference Guide* MK-90RD7010 • *Hitachi Command Control Interface Installation and Configuration Guide* MK-90RD7011	*Hitachi NAS Platform Command Line Reference* (accessed through HNAS Web Manager)
Performance	*Hitachi Unified Storage VM Block Module Performance Guide* MK-92HM7011	• *Hitachi Unified Storage File Module Server and Cluster Administration Guide MK-92USF7029* • *Hitachi Unified Storage File Module Storage Subsystem Administration Guide* MK-92USF7007

Type of manual / Topic	Block Module Manual Name	File Module Manual Name
System Maintenance	*Hitachi Unified Storage VM Block Module Hitachi SNMP Agent User Guide* MK-92HM7014	*Hitachi Unified Storage File Module Storage Subsystem Administration Guide* MK-92USF7006
	Hitachi Unified Storage VM Block Module Hitachi Audit Log User Guide MK-92HM7009	*Hitachi Unified Storage File Module NDMP Backup Administration Guide* MK-92USF012
Backup Operations	n/a	*Hitachi Unified Storage File Module NDMP Backup Administration Guide* MK-92USF012
Troubleshooting	• *Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide* MK-92HM7016 • *Hitachi Unified Storage VM Block Module Hitachi Storage Navigator Messages* MK-92HM7017	*Hitachi NAS Platform Troubleshooting Guide* (technical support only)

Unified Storage VM user guides by number

This section contains lists of the block module and file module user guides by part number.

Part number	Block module user guides
MK-92HM7003	*Hitachi Unified Storage VM Product Overview Guide*
MK-92HM7005	*Hitachi Unified Storage VM Block Module Hardware User Guide*
MK-92HM7006	*Hitachi Command Control Interface Command Reference*
MK-92HM7007	*Hitachi Command Control Interface User and Reference Guide*
MK-92HM7008	*Hitachi Command Control Interface Installation and Configuration Guide*
MK-92HM7009	*Hitachi Unified Storage VM Block Module Hitachi Audit Log User Guide*
MK-90RD7010	*Hitachi Unified Storage VM Block Module Thin Image User Guide*
MK-90RD7011	*Hitachi Unified Storage VM Block Module Performance Guide*
MK-92HM7012	*Hitachi Unified Storage VM Block Module Provisioning Guide*
MK-92HM7013	*Hitachi Unified Storage VM Block Module Hitachi ShadowImage® User Guide*
MK-92HM7014	*Hitachi Unified Storage VM Block Module Hitachi SNMP Agent User Guide*
MK-92HM7015	*Hitachi Unified Storage VM Hitachi System Operations Using Spreadsheets*
MK-92HM7016	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator User Guide*
MK-92HM7017	*Hitachi Unified Storage VM Block Module Hitachi Storage Navigator Messages*
MK-92HM7018	*Hitachi Unified Storage VM Block Module Hitachi TrueCopy® User Guide*
MK-92HM7019	*Hitachi Unified Storage VM Block Module Hitachi Universal Replicator User Guide*
MK-92HM7020	*Hitachi Unified Storage VM Block Module Hitachi Universal Volume Manager User Guide*
MK-92HM7021	*Hitachi Unified Storage VM Block Module Hitachi Volume Shredder User Guide*

Part number	File module user guides
MK-92USF001	*Hitachi Unified Storage File Module Hardware Reference*
MK-92USF002	*Hitachi Unified Storage File Module System Access Guide*
MK-92USF003	*Hitachi Unified Storage File Module Network Administration Guide*
MK-92USF004	*Hitachi Unified Storage File Module File Services Administration Guide*
MK-92USF005	*Hitachi Unified Storage File Module Data Migrator Administration Guide*
MK-92USF006	*Hitachi Unified Storage File Module Storage Subsystem Administration Guide*
MK-92USF007	*Hitachi Unified Storage File Module Server and Cluster Administration Guide*
MK-92USF008	*Hitachi Unified Storage File Module Snapshot Administration Guide*
MK-92USF009	*Hitachi Unified Storage File Module Replication and Disaster Recovery Administration Guide*
MK-92USF010	*Hitachi Unified Storage File Module Antivirus Administration Guide*
MK-92USF011	*Hitachi Unified Storage File Module Storage Systems User Administration Guide*
MK-92USF012	*Hitachi Unified Storage File Module NDMP Backup Administration Guide*
MK-92USF015	*Hitachi Unified Storage (HUS) File Module NAS Operating System SU 11.2 and SMU 11.2 Release Notes* RN-92USF018

Glossary

This glossary defines the special terms used in this document.

A

array

See disk array.

B

back-end I/O module

The hardware component that controls the transfer of data between the drives and cache. A back-end I/O module feature consists of a pair of boards. A back-end I/O module is also referred to as a disk adapter (DKA).

C

CHA

See channel adapter.

controller chassis

The hardware assembly that contains the logic and processing components of the Unified Storage VM storage system, including the front-end directors, virtual storage directors, cache memory, switches, and back-end directors. The Unified Storage VM storage system can be configured with one or two control chassis.

#	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

D

disk array

Disk array, or just array, is a complete storage system, including the control and logic devices, drives, connecting cables, and racks.

drive box

The hardware component of the Unified Storage VM that houses disk drives and/or flash drives.

dynamic provisioning

An approach to managing storage. Instead of “reserving” a fixed amount of storage, it removes capacity from the available pool when data is actually written to disk. Also called thin provisioning.

F

flash drive

A data drive that uses a solid-state memory device instead of a rotating hard disk drive to store information.

free capacity

The amount of storage space (in bytes) that is available for use by the host system(s).

H

host I/O module

The hardware component that processes channel commands from hosts and manages host access to cache. A host I/O module is also referred to as a channel adapter (CHA).

J

JRE

Java Runtime Environment

JVM

Java Virtual Machine

K

kVA

kilovolt-ampere

L

LDEV

logical device

license key

A specific set of characters that unlocks an application and allows it to be used.

logical device (LDEV)

An individual logical data volume (on multiple drives in a RAID configuration) in the storage system. An LDEV may or may not contain any data and may or may not be defined to any hosts. Each LDEV has a unique identifier or “address” within the storage system composed of the logical disk controller (LDKC) number, control unit (CU) number, and LDEV number. The LDEV IDs within a storage system do not change. An LDEV formatted for use by open-system hosts is called a logical unit (LU).

logical volume

See *volume*.

LUN

logical unit number. Sometimes used in place of LU, logical unit.

LUSE

LUN Size Expansion

M

mirror

In Universal Replicator, each pair relationship in and between journal groups is called a “mirror”. Each pair is assigned a mirror ID when it is created. The mirror ID identifies individual pair relationships between journal groups.

modify mode

The mode of operation of Storage Navigator that allows changes to the storage system configuration. See also *view mode*.

P

pair

Two logical volumes in a replication relationship in which one volume contains original data to be copied and the other volume contains the

copy of the original data. The copy operations can be synchronous or asynchronous, and the pair volumes can be located in the same storage system (in-system replication) or in different storage systems (remote replication).

pair status

Indicates the condition of a copy pair. A pair must have a specific status for specific operations. When an operation completes, the status of the pair changes to the new status.

PDB

power distribution box

PDP

power distribution panel

PDU

power distribution unit

pool

A set of volumes that are reserved for storing Copy-on-Write Snapshot data or Dynamic Provisioning write data.

R

RAID

redundant array of independent disks. A disk array in which part of the physical storage space is used to store user data and parity information, and another part is used to store a duplicate set of user data and parity information. This redundant configuration prevents data loss in case a disk drive within the RAID configuration fails, and enables regeneration of user data in the event that one of the array's member disks or the access path to it fails.

RAID group

A set of RAID disks that have the same capacity and are treated as one group for data storage and recovery. A RAID group contains both user data and parity information. This allows user data to be accessed in the event that one or more of the drives within the RAID group are not available. The RAID level of a RAID group determines the number of data drives and parity drives and how the data is “striped” across the drives. For RAID1, user data is duplicated within the RAID group, so there is no parity data for RAID1 RAID groups.

A RAID group can also be called an array group or a parity group.

RAID level

The type of RAID implementation. RAID levels include RAID0, RAID1, RAID2, RAID3, RAID4, RAID5 and RAID6.

S

SAS

serial-attached SCSI

SATA

serial Advanced Technology Attachment

service information message (SIM)

SIMs are generated by a storage system when it detects an error or service requirement. SIMs are reported to hosts and displayed on Storage Navigator.

service processor (SVP)

The computer in a VSP storage system that hosts the Storage Navigator software and is used to configure and maintain the storage system.

SIM

service information message

SOM

system option mode

SSD

solid state drive. See *flash drive.*

SVP

See *service processor (SVP)*.

system disk

The volume from which an open-systems host boots.

system option mode (SOM)

Additional operational parameters for the RAID storage systems that enable the storage system to be tailored to unique customer operating requirements. SOMs are set on the service processor.

V

volume

A logical device (LDEV), or a set of concatenated LDEVs in the case of LUSE, that has been defined to one or more hosts as a single data storage unit (LU).

Index

Symbols

Numer-

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

Hitachi Data Systems

MK-92HM7003-05