





HITACHI VIRTUAL STORAGE PLATFORM Architecture Guide



Udland International		1			
Merknader Remarks	Route nr. Flight no.	Tid Time	Utg. Gate	Riktning Destination	Merknader Remarks
Go to gate	SK4001	12:00	A1	Stockholm	
Go to gate	SK4002	12:15	A2	Stockholm	
Go to gate	SK4003	12:30	A3	Stockholm	
Go to gate	SK4004	12:45	A4	Stockholm	
Go to gate	SK4005	13:00	A5	Stockholm	
Go to gate	SK4006	13:15	A6	Stockholm	
Go to gate	SK4007	13:30	A7	Stockholm	
Go to gate	SK4008	13:45	A8	Stockholm	
Go to gate	SK4009	14:00	A9	Stockholm	
Go to gate	SK4010	14:15	A10	Stockholm	
Go to gate	SK4011	14:30	A11	Stockholm	
Go to gate	SK4012	14:45	A12	Stockholm	
Go to gate	SK4013	15:00	A13	Stockholm	
Go to gate	SK4014	15:15	A14	Stockholm	
Go to gate	SK4015	15:30	A15	Stockholm	
Go to gate	SK4016	15:45	A16	Stockholm	
Go to gate	SK4017	16:00	A17	Stockholm	
Go to gate	SK4018	16:15	A18	Stockholm	
Go to gate	SK4019	16:30	A19	Stockholm	
Go to gate	SK4020	16:45	A20	Stockholm	
Go to gate	SK4021	17:00	A21	Stockholm	
Go to gate	SK4022	17:15	A22	Stockholm	
Go to gate	SK4023	17:30	A23	Stockholm	
Go to gate	SK4024	17:45	A24	Stockholm	
Go to gate	SK4025	18:00	A25	Stockholm	
Go to gate	SK4026	18:15	A26	Stockholm	
Go to gate	SK4027	18:30	A27	Stockholm	
Go to gate	SK4028	18:45	A28	Stockholm	
Go to gate	SK4029	19:00	A29	Stockholm	
Go to gate	SK4030	19:15	A30	Stockholm	
Go to gate	SK4031	19:30	A31	Stockholm	
Go to gate	SK4032	19:45	A32	Stockholm	
Go to gate	SK4033	20:00	A33	Stockholm	
Go to gate	SK4034	20:15	A34	Stockholm	
Go to gate	SK4035	20:30	A35	Stockholm	
Go to gate	SK4036	20:45	A36	Stockholm	
Go to gate	SK4037	21:00	A37	Stockholm	
Go to gate	SK4038	21:15	A38	Stockholm	
Go to gate	SK4039	21:30	A39	Stockholm	
Go to gate	SK4040	21:45	A40	Stockholm	
Go to gate	SK4041	22:00	A41	Stockholm	
Go to gate	SK4042	22:15	A42	Stockholm	
Go to gate	SK4043	22:30	A43	Stockholm	
Go to gate	SK4044	22:45	A44	Stockholm	
Go to gate	SK4045	23:00	A45	Stockholm	
Go to gate	SK4046	23:15	A46	Stockholm	
Go to gate	SK4047	23:30	A47	Stockholm	
Go to gate	SK4048	23:45	A48	Stockholm	
Go to gate	SK4049	00:00	A49	Stockholm	
Go to gate	SK4050	00:15	A50	Stockholm	
Go to gate	SK4051	00:30	A51	Stockholm	
Go to gate	SK4052	00:45	A52	Stockholm	
Go to gate	SK4053	01:00	A53	Stockholm	
Go to gate	SK4054	01:15	A54	Stockholm	
Go to gate	SK4055	01:30	A55	Stockholm	
Go to gate	SK4056	01:45	A56	Stockholm	
Go to gate	SK4057	02:00	A57	Stockholm	
Go to gate	SK4058	02:15	A58	Stockholm	
Go to gate	SK4059	02:30	A59	Stockholm	
Go to gate	SK4060	02:45	A60	Stockholm	
Go to gate	SK4061	03:00	A61	Stockholm	
Go to gate	SK4062	03:15	A62	Stockholm	
Go to gate	SK4063	03:30	A63	Stockholm	
Go to gate	SK4064	03:45	A64	Stockholm	
Go to gate	SK4065	04:00	A65	Stockholm	
Go to gate	SK4066	04:15	A66	Stockholm	
Go to gate	SK4067	04:30	A67	Stockholm	
Go to gate	SK4068	04:45	A68	Stockholm	
Go to gate	SK4069	05:00	A69	Stockholm	
Go to gate	SK4070	05:15	A70	Stockholm	
Go to gate	SK4071	05:30	A71	Stockholm	
Go to gate	SK4072	05:45	A72	Stockholm	
Go to gate	SK4073	06:00	A73	Stockholm	
Go to gate	SK4074	06:15	A74	Stockholm	
Go to gate	SK4075	06:30	A75	Stockholm	
Go to gate	SK4076	06:45	A76	Stockholm	
Go to gate	SK4077	07:00	A77	Stockholm	
Go to gate	SK4078	07:15	A78	Stockholm	
Go to gate	SK4079	07:30	A79	Stockholm	
Go to gate	SK4080	07:45	A80	Stockholm	
Go to gate	SK4081	08:00	A81	Stockholm	
Go to gate	SK4082	08:15	A82	Stockholm	
Go to gate	SK4083	08:30	A83	Stockholm	
Go to gate	SK4084	08:45	A84	Stockholm	
Go to gate	SK4085	09:00	A85	Stockholm	
Go to gate	SK4086	09:15	A86	Stockholm	
Go to gate	SK4087	09:30	A87	Stockholm	
Go to gate	SK4088	09:45	A88	Stockholm	
Go to gate	SK4089	10:00	A89	Stockholm	
Go to gate	SK4090	10:15	A90	Stockholm	
Go to gate	SK4091	10:30	A91	Stockholm	
Go to gate	SK4092	10:45	A92	Stockholm	
Go to gate	SK4093	11:00	A93	Stockholm	
Go to gate	SK4094	11:15	A94	Stockholm	
Go to gate	SK4095	11:30	A95	Stockholm	
Go to gate	SK4096	11:45	A96	Stockholm	
Go to gate	SK4097	12:00	A97	Stockholm	
Go to gate	SK4098	12:15	A98	Stockholm	
Go to gate	SK4099	12:30	A99	Stockholm	
Go to gate	SK4100	12:45	A100	Stockholm	

EXECUTIVE SUMMARY

Information has become our most valuable asset. The challenge for global businesses lies in making information both available and secure. Hitachi Data Systems solutions help organizations transform raw data into valuable information by delivering on a vision that IT must be virtualized, automated, cloud-ready and sustainable. Data drives our world — and information is the new currency.

Solutions from Hitachi Data Systems help organizations transform their data centers into information centers that can quickly deliver capacity and compute resources, increase IT staff productivity while improving data access and operating expense, and enable pay-per-use, multitenancy storage services. At the same time, these solutions maintain enterprise-class security and privacy and achieve energy efficiency without compromising the reliability and availability of the infrastructure. At the heart is a virtualized platform for all data with the ability to manage multivendor environments.

With the introduction of Hitachi Virtual Storage Platform, Hitachi Data Systems is delivering the only 3D scaling storage platform designed for all data types. 3D scaling describes the fact that this storage platform can scale up to support increasing workloads, scale out to deliver the highest performance and capacity growth, and scale deep to extend the value of legacy and lower cost storage. With support for file, content and block storage, all data types can be integrated on a single platform. It is the only storage architecture that flexibly adapts data for performance and capacity and also extends to support storage from multiple vendors. With its 3D scaling, storage virtualization and high performance, Virtual Storage Platform, combined with the unique management capabilities of Hitachi Command Suite, transforms the data center.

Virtual Storage Platform is designed around a newly designed intelligent controller chassis with purpose built processors, global cache and a new, fifth-generation switch architecture. Together, these capabilities deliver the scalability, performance, high availability and advanced storage functionality to satisfy all data center file, object and block storage needs.

This storage platform requires less floor space and cooling due to innovative new technologies and packaging. It offers tighter integration with server virtualization and improved storage administrator productivity through Hitachi Command Suite.

Hitachi Command Suite is an open, SMI-S standards-based suite of tools providing comprehensive management, control and reporting for all file, object and block services. This software suite enables “single pane of glass” management of all virtualized internal and external storage attached to Virtual Storage Platform.

The unique controller-based virtualization of Virtual Storage Platform aggregates all storage, including internal and externally attached Hitachi branded and third-party storage, to create a common pool of capacity. Reusable storage services, including thin provisioning, host port virtualization and nondisruptive heterogeneous data migration, can then access all storage in the virtual pool.

Unlike alternative storage systems, Virtual Storage Platform optimizes return on storage assets by combining unmatched performance, capacity and utilization of multivendor assets. In combination with unique Hitachi Dynamic Tiering capabilities, it makes the storage infrastructure more fluid and IT more agile.

CONTENTS

Executive Summary	3
Contents	4
Chapter 1: Data Drives Our World	6
The Four Building Blocks of Our Product and Solutions Portfolio	6
Virtualization	
Automation	
Cloud-readiness	
Sustainability	
Our Approach: Provide a Single Virtualized Platform for All Data with Hitachi Virtual Storage Platform	8
3D Scaling	
Fifth-generation Hi-Star-E Network	
One Storage Solution for all Data Types: File, Block and Object	9
Hitachi Command Suite	
Business Continuity Solutions	
Hitachi Universal Replicator	
Hitachi NAS Platform, Powered by BlueArc®	
Hitachi Content Platform	
Hitachi Data Discovery Suite	
What Is Storage Virtualization?	17
A Brief History of Controller-based Storage Virtualization	
Intelligent Controller-based Virtualization	
How Controller-based Virtualization Works	
One Platform for All Data	19
Operational Simplicity	
Dynamic Provisioning	
Tiered Storage	
Hitachi Dynamic Tiering	
Quality of Service	
End-to-end Secure Virtualization for VMware	
Host Storage Domains	
Mainframe Storage Virtualization	
Business Continuity	
Surviving Link Failure	

Three Data Center Business Continuity Solution	
Four Data Center Business Continuity Solution	
High Availability Manager	
Security and Compliance	
The Green Data Center	
Hitachi Data Systems Global Solution Services and Best Practices	27
Chapter 2: Virtual Storage Platform Architecture	28
Hi-Star-E Network — the Industry's Highest Performing Storage Architecture	29
Data Cache and Control Memory	31
Data Cache	
Control Data	
Advanced Data Cache Algorithms Provide Read-ahead for High Performance Sequential Reads	
Hitachi Cache Residency Manager Allows Datasets to Be Permanently Placed in Data Cache	
Front-end Design	33
Connecting to Mainframe Hosts	
High Performance, SAS Back-end Design	35
RAID Implementation Advantages	
RAID-1+0	
RAID-5 Distributed Parity	
Concatenated RAID-5 Array Groups	
RAID-6 Redundant Distributed Parity	
RAID Rebuild	
Connecting External Storage	
Packaging and Configuration	40
The Control Chassis	
Power Subsystem	
The System Monitoring Network	
High Availability	42
Appendix A: Dynamic Tiering	44
Tiered Storage for Files, Volumes and Pages	45
Hitachi Dynamic Tiering	
Appendix B: The Evolution of Hitachi Storage Virtualization	46

DATA DRIVES OUR WORLD

Information has become the new currency in this world driven by data.

Send it. > Receive it. > Store it. > Manage it. > Process it. > Protect it. > Use it.

Information has become our most valuable asset. The challenge for global businesses lies in making this information available both when and where needed, while keeping it secure.

Solutions from Hitachi Data Systems assist organizations in transforming raw data into valuable information by delivering on a vision that IT infrastructure must be:

VIRTUALIZED



CLOUD-READY



AUTOMATED



SUSTAINABLE



Our customers, whether they are in financial services, healthcare, media, entertainment, retail, manufacturing, etc., trust Hitachi Data Systems to protect their data, the raw resource that lies at the heart of their businesses. Every industry is being or has been transformed by digital data and information.

The Four Building Blocks of Our Product and Solutions Portfolio

To help summarize and simplify the broad portfolio of technologies, products and solutions that Hitachi Data Systems delivers, we use the concept of four building blocks. They are four key enabling technologies that serve as foundations for our entire product portfolio (see Figure 1). Though some competitors may offer point products that provide one or perhaps two of these enabling technologies, only Hitachi Data Systems combines all four throughout our product portfolio. Hitachi Data Systems solutions are:

Virtualization

Virtualized IT environments enable the automation of the most labor-intensive data management and data protection operations. They support fast delivery of capacity and storage processing resources, and they reduce costs while increasing asset utilization at the pace of your business.

Automation

Automation helps accelerate the introduction of new applications and business systems with less cost by automating time-consuming operations. It also facilitates continuous operations while reducing capital expenditure (CAPEX) wasted on underutilized storage capacity. Automation can increase your IT staff's productivity while maintaining or improving data access, security and operating expenses (OPEX). Automation can also enable your IT staff to more efficiently align the IT infrastructure with and react to your business objectives.

Cloud-readiness

Cloud-ready IT environments enable the adoption of or evolution to a cloud computing model or IT as a service at your own pace, with minimal risk. A cloud-ready infrastructure provides seamless access to all information and resources; it is cost-effective, offering dynamic scaling and

provisioning for your lines of business, customers or partners. It also enables consolidation and aggregation of storage, server, networking and middleware assets to optimize operations through virtualization. Cloud-readiness is based on multitenancy and service level agreement (SLA)-oriented storage services, but it must also ensure that enterprise-class security and privacy are maintained.

Sustainability

Sustainability is becoming increasingly important across all aspects of business. Understanding what the providers of IT infrastructure components are doing can help organizations plan for and move towards a more sustainable IT infrastructure. This enables organizations to achieve energy efficiency without compromising the reliability, availability or security of their IT infrastructure. Having a blueprint or plan for the availability and implementation of sustainable infrastructure elements also helps in planning for centralization and consolidation of data management and data protection; it supports reduced power consumption and a lower overall total cost of ownership (TCO).

Virtual Storage Platform with Hitachi Command Suite delivers on these four building blocks. The industry leading virtualization capabilities of this platform and the automation provided through enhanced Command Suite features and capabilities, such as Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering, are key elements of this foundation. Additionally, the reductions in power consumption and footprint along with the new cooling characteristics of Virtual Storage Platform make this an excellent player in a sustainable IT infrastructure. This combination of hardware and software from Hitachi Data Systems is completely cloud-ready, enabling you to move to the cloud now or in the future.




Figure 1. There are four building blocks or key enabling technologies in the Hitachi Data Systems solution portfolio.

Our Approach: Provide a Single Virtualized Platform for All Data with Hitachi Virtual Storage Platform

Hitachi Data Systems has a unique approach to storage management: a single virtualized platform for all data (structured and unstructured) with “single pane of glass” management for heterogeneous environments.

Hitachi Data Systems is focused on the evolving requirements of structured and unstructured data and is the only vendor who can provide you with a single, virtualized and scalable data management architecture for all open systems platforms (x86, VMware, Sun Solaris, Linux and IBM® AIX®) and mainframe operating environments, as well as all data and information types.

This single platform, Virtual Storage Platform, enables a common set of management tools across all storage assets, whether they are the products of Hitachi Data Systems or our competitors. Our common platform and management tools improve the productivity of storage administrators and enable better utilization, cost management (power and cooling), chargeback, simplicity and data protection.

The technologies at the core of Virtual Storage Platform have been developed through multiple generations and over many years with the overriding goal of providing a highly scalable high performance platform for storage services. The strategy driving this development has had four basic tenets:

- Expand the scalability of the storage platform to: enhance the ability to nondisruptively increase the internal processing capability and cache, improve the response time, increase IOPS and increase internal and external storage capacity
- Use control unit virtualization to support massive consolidation of storage services on a single platform
- Offer comprehensive solutions for storage management, business continuity, data mobility and advanced storage functionality
- Provide a common storage management suite with access to the metrics necessary to measure and automate storage services

3D Scaling

Virtual Storage Platform delivers the industry's only 3D scaling storage platform. With its unique ability to concurrently scale up, out and deep in a single storage system, the new Virtual Storage Platform easily adapts

to meet performance, capacity, connectivity and virtualization needs.

Today, with the virtualization of more and more servers on powerful platforms with multicore processors, their increasing I/O workloads are being consolidated onto a single storage system. To handle the growing I/O demands of these virtualized servers, the storage system must be able to scale up dynamically by nondisruptively integrating new processing capacity, such as processors, cache and internal switching. This ability of Virtual Storage Platform to scale up will enable even more virtual server consolidation, better utilization of resources and lower cost.

With the growing number of servers and applications accessing the storage system, it is important for the storage system to be able to scale out by adding storage resources to grow the base system for greater capacity, performance and connectivity. Thanks to the Virtual Storage Platform Hi-Star-E network architecture, host servers can access needed resources out of a common pool of processing power, global cache, ports and disk capacity.

Scaling deep extends the advanced functions of Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform pool of storage resources. Once the external storage is virtualized behind Virtual Storage Platform, the data resident on this external storage can then be migrated, tiered, replicated and managed by Virtual Storage Platform. As a result, more cost-effective modular storage can be used as an archive tier and older data storage systems can gain a longer useful life. Distance replication for business continuity can be extended to lower cost, lower function storage systems, which are virtualized behind a Virtual Storage Platform.

Fifth-generation Hi-Star-E Network

At the heart of Virtual Storage Platform is the fifth-generation Hi-Star-E network, which provides a system where all the basic components, including front-end directors, back-end directors, global cache modules and virtual storage directors, are connected through this high performing, nonblocking network. The design of Virtual Storage Platform and its implementation of this

network enable redundant pairs of directors and cache modules to be added as required without disruption to connected host servers. In addition, a second control chassis can be added and, through these network links, tightly coupled to the existing chassis. The internal

processors along with the front-end and back-end directors share a global cache, which creates a common pool of storage resources and disk capacity. This disk capacity can include external storage, which is connected through front-end director initiator ports.

One Storage Solution for all Data Types: File, Block and Object

Hitachi Data Systems storage solutions support a full range of file, object and block storage services (see Table 1). The solutions give IT the ability to mix and match services to satisfy a desired goal. All solutions use open protocols, rather than proprietary APIs, to promote interoperability and maintain scalability. When combined with the Virtual Storage Platform, this complete offering of software solutions allows one storage solution to support all data types.

Table 1. Hitachi Data Systems Solutions Providing File, Object and Block Services

	File	Object	Block
Hitachi Adaptable Modular Storage 2000 family	X	X	X
Hitachi Basic Operating System	X		X
Hitachi Basic Operating System V			X
Hitachi Business Continuity Manager			X
Hitachi Content Platform	X	X	
Hitachi Copy-on-Write Snapshot			X
Hitachi Data Discovery Suite	X	X	
Hitachi Data Protection Suite, powered by CommVault®	X	X	X
Hitachi Dynamic Link Manager Advanced			X
Hitachi Dynamic Provisioning			X
Hitachi NAS Platform, powered by BlueArc®	X		X
Hitachi Replication Manager			X
Hitachi ShadowImage® Heterogeneous Replication			X
Hitachi Storage Capacity Reporter, powered by APTARE®			X
Hitachi Storage Services Manager	X		X
Hitachi Tiered Storage Manager			X
Hitachi TrueCopy® Synchronous remote replication			X
Hitachi Tuning Manager			X
Hitachi Universal Replicator			X
Hitachi Virtual Storage Platform	X	X	X
Virtual Tape Library Solutions by Hitachi Data Systems	X	X	X

The following sections describe the complementary storage capabilities and software components in depth.

Hitachi Command Suite

Hitachi Command Suite is an integrated software suite for efficient management of virtualized storage and server infrastructures that lowers capital and operational costs. With its usability, workflow, performance, scalability and private cloud enablement, Hitachi Command Suite lets IT organizations build sustainable infrastructures with leading storage technologies. It helps IT align more flexibly with changing business requirements and it helps maximize return on IT investments.

Hitachi Command Suite delivers unique 3D management in support of Virtual Storage Platform, the only 3D scaling storage platform designed for all data types and all other Hitachi storage systems. 3D management simplifies operations and lowers costs with three distinct dimensions:

- Manage up to scale with the largest infrastructure deployments.
- Manage out with breadth to manage storage, servers and the IT infrastructure.
- Manage deep with the integration required to manage the multivendor resources of today's complex data centers.

Hitachi Command Suite is an application-centric storage management solution (see Figure 2) that incorporates the internationally recognized Information Technology Infrastructure Library (ITIL) services management best practices to simplify administration of a common pool of multivendor storage. This suite of management software gives administrators the tools to deliver dependable storage system availability and performance (see Figure 3). The software offers comprehensive management, control and discovery for file, object and block storage services and reduces complexity, costs and risk in the storage infrastructure.

Figure 2. Hitachi Command Suite provides a unified management solution for managing the storage environment.






Figure 3. Hitachi Command Suite provides deep integration through a common GUI along with usability, data repositories and task workflows.

The software eliminates the need for separate and isolated tools to manage each vendor's storage. The Command Suite management framework improves operational efficiency and resilience, and integrates solutions for storage resource management, tiered storage, performance management and business continuity.

Hitachi Command Suite includes the following products:

- Hitachi Basic Operating System
- Hitachi Basic Operating System V
- Hitachi Command Director
- Hitachi Device Manager
- Hitachi Dynamic Link Manager Advanced
- Hitachi Dynamic Provisioning
- Hitachi Dynamic Tiering
- Hitachi Replication Manager
- Hitachi Storage Capacity Reporter, powered by APTARE®
- Hitachi Storage Services Manager
- Hitachi Tiered Storage Manager
- Hitachi Tuning Manager
- Hitachi Virtual Server Reporter, powered by APTARE

Hitachi Basic Operating System (BOS) and Hitachi Basic Operating System V (BOS V)

Hitachi Basic Operating System is a set of fully integrated tools for simplifying storage management and enabling the partitioning and virtualization capabilities of Virtual Storage Platform (see Table 2). BOS includes Hitachi Storage Navigator 2, Hitachi Device Manager, Hitachi Dynamic Provisioning, a 50 server license for Hitachi Dynamic Link Manager Advanced, all software that was formerly included in the Hitachi Resource Manager™ utility package, Hitachi Virtual Partition Manager and Hitachi Server Priority Manager. BOS V adds Hitachi Universal Volume Manager and expanded cache partitioning to the BOS package (see Table 3).

Table 2. Hitachi Basic Operating System (BOS) Integrated Tools

Software Name	Open Systems	Mainframe	Function
Hitachi Device Manager	✓	✓	■ Primary control point for the Basic Operating System ■ Provides a single point of management for all Hitachi physical and virtual storage ■ Acts as the interface for integration with other systems
Hitachi Dynamic Link Manager Advanced	✓	✓	■ Provides advanced SAN multipathing with centralized management
Hitachi Dynamic Provisioning	✓	✓	■ Provides thin provisioning for simplified provisioning operations, automatic performance optimization and storage space savings
Hitachi Performance Monitor	✓	✓	■ Provides an intuitive, graphical interface to assist with performance configuration planning, workload balancing, and analyzing and optimizing storage system performance
Hitachi Storage Navigator 2	✓	✓	■ Facilitates storage system maintenance functions
Hitachi Virtual Partition Manager	✓	✓	■ Supports up to four cache partitions
Storage System-based Utilities	✓	✓	■ Includes LUN manager/LUN expansion, LUSE, virtual LVI/LUN manager (customized volume size), volume port security and volume security port options, audit log, command control interface (CCI), Microsoft® Volume Shadow Copy Service (VSS) support, and volume shredder

Table 3. Hitachi Basic Operating System V (BOS V) Integrated Tools

Software Name	Open Systems	Mainframe	Function
Expanded Cache Partitioning	✓	✓	■ Supports up to 32 cache partitions
Hitachi Universal Volume Manager	✓	✓	■ Enables virtualization of external heterogeneous storage

Hitachi Command Director

Hitachi Command Director centralizes management and reporting across Hitachi Command Suite. By consolidating reporting of storage configuration, capacity, performance and tier information, Hitachi Command Director provides a business-oriented view into the storage environment to easily align Hitachi storage assets with critical business applications and simplify application-to-storage reporting. It also ensures compliance to application-specific storage service levels.

Hitachi Dynamic Provisioning

Hitachi Dynamic Provisioning provides Virtual Storage Platform with thin provisioning services. Thin provisioning gives applications access to virtual storage capacity. Applications accessing virtual, thin provisioned volumes are automatically allocated physical disk space, by the storage system, as they write data. This means volumes use enough physical capacity to hold application data and no more.

All thin provisioned volumes share a common pool of physical disk capacity. Unused capacity in the pool is available to any application using thin provisioned volumes. This eliminates the waste of overallocated and underutilized storage.

Dynamic Provisioning also simplifies storage provisioning and automates data placement on disk for optimal performance. Administrators do not need to micro-manage application storage allocations or perform complex, manual performance tuning. In addition, physical storage resources can be added to the thin provisioning pool at any time, without application downtime.

Hitachi Dynamic Tiering

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

Hitachi Replication Manager

With Hitachi Replication Manager, administrators can configure, monitor and manage any replication systems that are based on Hitachi storage systems. Combining robust reporting, mirroring and features previously available in separate offerings, it manages in-system or distance replication for open systems and mainframe environments. Replication Manager simplifies and optimizes configuration, operations, task management, automation and monitoring of the critical applications and storage components of your replication infrastructure.

Hitachi Storage Capacity Reporter, Powered by APTARE®

Hitachi Storage Capacity Reporter software provides end-to-end storage capacity reporting from applications and hosts to heterogeneous storage systems to improve capacity planning, increase storage utilization and reduce infrastructure costs. It offers a host application view of storage usage for Oracle, Microsoft® SQL Server® and Microsoft Exchange without the use of host-based agents. It helps to identify underutilized or overutilized storage capacity for better utilization and defer additional storage capacity investments.

Hitachi Storage Services Manager

Hitachi Storage Services Manager software provides storage infrastructure management and storage resource management (SRM) across heterogeneous storage assets. It autotdiscovers hosts, switches and disk storage systems while analyzing and visualizing their interdependencies. With accurate reporting of storage assets, application availability and quality of service is ensured while use of storage resources is optimized.

Hitachi Tiered Storage Manager

Hitachi Tiered Storage Manager provides the unique ability to migrate volumes between heterogeneous tiers of storage, without affecting application access to data. The software supports up to 64 simultaneous migrations, using 64 processor threads, and it offers a single interface for all data movement between storage virtualized by Virtual Storage Platform.

Hitachi Tuning Manager

Hitachi Tuning Manager software is a path-aware storage resource management application that maps, monitors and analyzes storage network resources from the application to the storage device. It provides the end-to-end visibility required to isolate and diagnose

performance bottlenecks with a focus on business applications, such as Oracle, Microsoft SQL Server, Microsoft Exchange and IBM DB2®.

Hitachi Universal Volume Manager

Universal Volume Manager simplifies migration of applications using third-party storage to the Virtual Storage Platform. The software eliminates the need to reformat LUNs or perform complex extent mapping, providing a minimally disruptive migration process. Once virtualized behind Virtual Storage Platform, external LUNs gain access to all Hitachi storage services.

Hitachi Virtual Partition Manager

Hitachi Virtual Partition Manager logically partitions Virtual Storage Platform cache.

Hitachi Virtual Server Reporter, powered by APTARE

Hitachi Virtual Server Reporter provides comprehensive heterogeneous storage reporting for VMware virtual server infrastructures. By properly monitoring a virtual machine's use of networked storage, virtual machines can be assured to have the right storage resources available to meet their particular needs.



Business Continuity Solutions

Virtual Storage Platform supports a comprehensive range of business continuity and data protection tools, including:

- Hitachi ShadowImage® Heterogeneous Replication for full volume local replication of data
- Hitachi TrueCopy® Synchronous for remote replication
- Hitachi Universal Replicator for efficient, cost-effective replication and recovery of data hosted on multivendor storage attached to Virtual Storage Platform
- Hitachi Replication Manager and Hitachi Business Continuity Manager for common management of and access to data movement products
- Virtual Tape Library Solutions by Hitachi Data Systems for disk-based backup on Virtual Storage Platform, without the need to change established backup policies or procedures
- Hitachi Data Protection Suite, powered by CommVault®, an integrated portfolio of storage services to automate and simplify data and storage management tasks
- Hitachi Copy-on-Write Snapshot for local volume snapshots
- A global file namespace for seamless file migration across heterogeneous storage systems without the need to remount

Hitachi Universal Replicator

Universal Replicator delivers the industry's first multiplatform, multivendor, remote copy capability. Supporting asynchronous, real time replication of data from one Virtual Storage Platform to another, the software gives storage administrators a consistent replication approach, using a common set of procedures, tools and skills, for all data center storage.

Conventional replication solutions push updates from the primary storage system to a remote location. Universal Replicator software changes things. Universal Replicator running on Virtual Storage Platform at the remote site pulls changes from journal volumes maintained by Virtual Storage Platform running Universal Replicator at the primary site. This approach offloads replication overhead from the primary storage system and dramatically reduces the impact of replication on the business application. Transmitted record sets include a time stamp (mainframe only) and sequence number, enabling the replication engine to verify receipt of all records at the remote site and to sequence writes in the correct order for storage.

Hitachi NAS Platform, Powered by BlueArc®

The Hitachi NAS Platform interfaces with Virtual Storage Platform to consolidate unstructured data, enabling highly efficient content indexing and delivering intelligent file tiering. These capabilities support policy-based migration of data and content among storage and archive tiers, delivering on the promise of one platform for all of your data.

With leading performance and scalability, Hitachi NAS Platform allows administrators to consolidate third-party NAS devices or file servers to simplify the IT infrastructure. They can also reduce the complexity of storage management and lower TCO. The advanced virtualization and dynamic provisioning capabilities significantly improve efficiency, agility and utilization across file sharing environments. The native policy manager software enables administrators to determine when and where to move the files or folders, and how often. Intelligent file tiering helps organizations move data among storage and archive tiers with automated, policy-driven migration tools. It enables policy-based hierarchical storage management (HSM) within Hitachi NAS Platform.

With intelligent file tiering, the job of the storage administrator is to classify and segregate data, create storage tiers with separate file systems and then create the policies to manage and move the files to the appropriate tier based on the policy rules the administrator builds. This automation through policies significantly reduces the administrator's workload once the policies are defined. However, they do need to continue to monitor and classify the files to ensure that the policies are correct or determine whether there is a need to create additional policies to handle specific situations.

Hitachi Content Platform

The Hitachi Content Platform is a multipurpose distributed object store that stores files and associated metadata (information about the file) as an object. The system imbues these objects with intelligence that enables them to automatically take advantage of advanced storage and data management features to ensure proper placement, retention, security and protection of content. With massive scale, multiple storage tiers, Hitachi reliability, cloud capabilities, multitenancy and configurable attributes for each tenant, the platform can support a wide range of applications on a single physical cluster. By dividing the physical cluster into multiple, uniquely configured tenants, administrators create virtual content platforms that can be further subdivided into namespaces for further organization of content, policies and access.

One of the key applications of this technology is the archiving of content for compliance with legal, government or self-imposed regulations. Regulatory compliance and the threat of legal discovery now require that companies prove the authenticity of stored data, give assurances that data has not been altered and provide access to archived data on demand. To do this, an archiving platform must offer encryption for data at rest, secure role-based access and access logs for audit and forensics. Further, it must have the ability to retain records in a protective storage environment for many years across technology refreshes and ensure the destruction of data when the retention period has expired.

Hitachi Content Platform offers SAN-based digital archiving, backed by Virtual Storage Platform for scalability, availability and performance of long-term, fixed content storage. The object storage solution offers

automated policy-based authentication, protection, retention and destruction to ensure content integrity with integrated indexing, search and retrieval.

Hitachi Content Platform uses standard file and Internet protocols to ingest the data and metadata developed by independent software vendors (ISVs). The metadata from different data types, such as files, documents, email, web content and PACS (picture archiving and communication system), is stored in a common repository. Virtual Storage Platform enables Hitachi Content Platform storage to scale to petabytes of data and billions of objects. In addition, virtualization offers seamless migration of archived content across generations of storage systems.

Hitachi Data Discovery Suite

Hitachi Data Discovery Suite simplifies cross-platform search and intelligent file tiering with a single user interface for improved productivity and business competitiveness. It also supports proactive e-discovery and identification of content for compliance and governance. Hitachi Data Discovery Suite supports federated searching of file systems and archive data on behalf of Hitachi NAS Platform, Hitachi Content Platform and NetApp storage systems.

This suite helps search, identify, collect and preserve data across the storage infrastructure from one secure location. In addition, it lets end users and IT staffs conduct self-service file restoration, cross-platform search and retrieval, and file-based tiering, using existing tiered storage. Hitachi Data Discovery Suite helps in efficiently managing, retaining and discovering volumes of information and meeting legal obligations.

Hitachi Data Discovery Suite enables searching across multiple data silos from a single interface, according to the user access privileges via Microsoft Active Directory, LDAP and RADIUS integration. The data silos can include email archives and snapshots or versions from remote locations.

Hitachi Data Discovery Suite also provides advanced storage management tools for Hitachi NAS Platform and Hitachi Content Platform with intelligent file tiering. These tools provide a flexible and extensible file content and metadata model for assigning data to the most appropriate media, whether it is online, nearline or archive. This capability increases the overall efficiency of the storage environment by reducing the backup times and duplication of data as well as enabling faster retrievals of data.



What Is Storage Virtualization?

PART 1

The act of abstracting, hiding or isolating the internal function of a storage (sub) system or service from applications, compute servers or general network resources for the purpose of enabling application — and network — independent management of storage or data.

PART 2

The application of virtualization to storage services or devices for the purpose of aggregating, hiding the complexity or adding new capabilities to lower level storage resources.

While many storage virtualization solutions address part one of the SNIA definition (as shown above), the controller-based storage virtualization solution offered by Hitachi Data Systems in the Virtual Storage Platform is unique in its ability to satisfy the second part.

A Brief History of Controller-based Storage Virtualization

Virtual Storage Platform storage virtualization technology has been in continuous development since 1995, when Hitachi Data Systems debuted Hitachi Freedom 7700, with the ability to consolidate storage for mainframe, UNIX and Microsoft Windows® servers on a single platform. Hosting multiple operating platforms was achieved by separating control information and application data in the storage controller. Three years later, the controller technology was refined to enable the dynamic movement of LUNs between RAID groups; this functionality is still unique to Hitachi Data Systems almost 15 years after it was introduced.

Hitachi Lightning 9900™ Series enterprise storage systems introduced the first generation of the Hitachi Universal Star Network™ crossbar switch architecture in 2000. The revolutionary crossbar switch dramatically increased scalability and connectivity. This was enhanced two years later with the ability to support virtual Fibre Channel ports.

In 2004, Hitachi Data Systems introduced the first-generation Hitachi Universal Storage Platform®, with the ability to virtualize third-party storage through standard Fibre Channel port connections. The second generation of this family of storage systems, Universal Storage Platform V and Universal Storage Platform VM, were introduced in 2007. Powered by the fourth-generation Hitachi Universal Star Network V crossbar switch architecture, these storage platforms delivered high performance, massive scalability and total security. They also provided administrators with the ability to scale from very small to very large implementations, using the same architecture and the same unified management strategy.

Virtual Storage Platform is powered by the fifth-generation Hitachi switch architecture: the Hi-Star-E PCI Express switched grid. Virtual Storage Platform not only is built on this fifth-generation switch architecture, but it also has separated the management and control processing from the front-end and back-end processing, delivering improved performance for the unique processing requirements of different data types. This latest architecture provides an entirely new level of scalable enterprise storage and offers much higher performance, scalability, reliability and flexibility than any competitive offering available today.


Intelligent Controller-based Virtualization

Intelligent, controller-based storage virtualization provides a platform for aggregating all file, object and block storage services and all heterogeneous storage hardware. This is the only approach to storage virtualization that gives lower level and legacy storage systems access to enterprise-class performance, functionality and services, in accordance with the SNIA definition of storage virtualization. The Virtual Storage Platform Hi-Star-E network architecture allows direct access to application data in global cache. Unlike other storage architectures, Virtual Storage Platform maintains mappings that describe global cache data in the control memory of individual control processors. This control memory is in the RAM local to each control processor in the storage controller, giving application data full access to the capacity and bandwidth of the global cache. Separating control information from data in the global cache also allows mappings to change, dynamically altering the data configuration. This process provides the basis for controller-based virtualization.

How Controller-based Virtualization Works

To virtualize external storage with Virtual Storage Platform, the Fibre Channel host ports of the external storage systems are connected to the Fibre Channel front-end director ports of Virtual Storage Platform (see Figure 4). The ports on Virtual Storage Platform are then reconfigured to change their behavior from SCSI targets, which is storage as viewed from the application server, to SCSI initiators, which is the application server view of storage. Hitachi Universal Volume Manager administers connectivity between Virtual Storage Platform and external storage devices. When a storage system is connected to Virtual Storage Platform, the software queries the Fibre Channel attachment to discover existing LUNs. These can then be reconnected to application host servers, without compromising file system consistency or integrity.

Figure 4. Virtual Storage Platform supports up to 255PB of internal and externally attached storage, combined, and offers common data management and services for all internal and externally attached storage.



One Platform for All Data

Controller-based storage virtualization is a key enabling technology to providing one platform for all data. In addition, the new internal architecture of Virtual Storage Platform, with management and control processing separated from front-end and back-end processing, delivers improved performance for any unique processing required by different types of data. It provides the flexibility to mix and match front-end and back-end storage services to rapidly achieve a desired goal. Consolidated and aggregated multivendor storage resources provide a single platform of enterprise capacity for all file, object and block services functionality.

Operational Simplicity

Hitachi controller-based storage virtualization dramatically simplifies the enterprise storage infrastructure. Virtualization enables the consolidation and aggregation of all multivendor storage resources in a single pool of enterprise storage capacity.

Virtual Storage Platform is massively scalable. In addition to supporting direct attached Fibre Channel external storage, including virtual tape libraries, the storage platform supports NAS and SAN. Existing LUNs and data on externally attached storage systems are automatically discovered by the Universal Volume Manager and presented through global cache as if they were still connected to the application server.


This Hitachi storage controller provides a common storage management platform for all virtualized multi-vendor storage. Hitachi Command Suite offers a comprehensive set of management, control and discovery tools for file, object and block storage services, and it supports open systems and mainframe environments. The software integrates solutions for storage resource management, tiered storage, business continuity and thin provisioning, giving storage managers a unified management solution for all enterprise storage.

Dynamic Provisioning

Virtual Storage Platform offers a large-scale, enterprise-class virtualization layer with thin provisioning. Dynamic Provisioning uses thin provisioning techniques to give applications access to virtual storage capacity (see Figure 5). As an application writes data to a virtual volume, the storage system dynamically allocates physical disk space from a common pool. Unused capacity in the pool

is available to any application that needs it. This eliminates the waste of overallocated and underutilized application storage capacity. The automated allocation of application storage capacity also dramatically simplifies the storage provisioning process. Dynamic Provisioning reduces the need to perform complex storage system tuning to improve I/O performance. Virtual Storage Platform automatically stripes thin provisioned volumes across the entire width of the physical pool of disks. This wide striping balances the I/O load across many disk spindles, giving consistently high performance for all I/O.

Figure 5. Hitachi Dynamic Provisioning reduces the size, cost and complexity of the enterprise storage environment.



Virtual Storage Platform implementation of thin provisioning is unique. All internal and externally attached storage resources virtualized by the Hitachi storage controller can participate in the common physical pool used by thin provisioned volumes.

How Hitachi Dynamic Provisioning Works

Dynamic Provisioning uses a paging technique that lets administrators provision application storage without actually allocating physical disk capacity until it is required. Dynamically provisioned volume capacity is virtual. The space is allocated to the application, however, and the space is not physically provisioned until needed. As an application writes data to the virtual volume, disk space is assigned in 42MB page increments from a common pool of capacity. Once allocated, pages are striped across the pool of logical devices, which can span hundreds of physical disk spindles. The storage system manages wide striping dynamically, reducing the need for storage administrators to manually tune I/O for performance.

Tiered Storage

Tiered storage delivers substantial infrastructure cost savings by matching storage system capabilities and application sensitivity to performance, availability, price and functionality. For example, a tiered storage configuration enables nonessential copies of primary application data, i.e., point-in-time snapshots, replication copies and data mining, etc., to be hosted on legacy and low-cost storage systems. This frees capacity on high-end storage for critical application data.

Although the benefits of tiered storage are widely recognized, implementation has lagged due to problems performing nondisruptive data movement between heterogeneous tiers. The Virtual Storage Platform virtualization engine solves this problem by providing a single interface to all multivendor tiered storage resources through Hitachi Universal Volume Manager.

Hitachi Tiered Storage Manager is unique in its ability to migrate data between heterogeneous tiers of storage, without affecting application access to data. The software supports up to 64 simultaneous migrations, using 64 processor threads, and it offers a single interface for all data movement between storage virtualized by Virtual Storage Platform.

Tiered Storage Manager manages groups of volumes. Migration groups are typically associated with a single

business application and represent a list of volumes requiring similar storage resources. Storage tier is a logical label given to a group of storage resources that match a predefined set of capabilities, such as minimum I/O response time. During data migration, a migration group is moved between one storage tier and another. Migration copies the contents of a source LDEV to a target LDEV. When the two LDEVs are fully synchronized, Tiered Storage Manager swaps the LDEV identifiers. Hosts accessing the LDEV will not notice any change. The LUN address path, as seen by the host, remains constant. In addition, source and target LDEVs are not required to have the same RAID configuration.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering extends Hitachi Dynamic Provisioning and takes the automation of tiered storage to a new level. It enables the management of multiple storage tiers as a single entity. It presents a new kind of virtual volume with embedded smart tiering that monitors access and moves data at the fine grain page level. Dynamic Tiering breaks the volume into pages and automatically promotes only the pages that are frequently referenced, avoiding the time and storage space to move the entire data set, file or volume.

Dynamic Tiering uses a dynamic storage pool, the same one that is used by Dynamic Provisioning, which is divided into 42MB pages. Virtual volumes are made up of pages from the dynamic storage pool. Each page is mapped to a tier within the pool and can be anywhere in the pool on any tier within that pool. Virtual Storage Platform, through Dynamic Tiering, determines which pages it is best to place in higher tiers and which in lower tiers. The pages can be moved up and down between the tiers within the pool for performance optimization purposes.

In operation, a Hitachi Dynamic Tiering environment will have most of the highly referenced pages in the highest tier and the least referenced pages in the lowest tier. Dynamic Tiering will try to use highest tiers as much as possible. This can provide the effective appearance that all the data is on the fastest tier of storage; in fact, most of the data resides on lower cost, lower performance storage tiers comprising SAS and/or SATA drives. The lowest tier also contains much of the spare capacity.

A Dynamic Tiering volume is another type of volume in a virtualized storage architecture. Just like disk types and RAID configurations, multitier volumes are another way

of delivering storage with tailored cost and performance service levels. Capacity can be dynamically added to or removed from a tier at any time, which means that sizing a tier for a pool is much easier.

With Dynamic Tiering, the storage administrator does not need to classify data or define policies; once the tiers are configured, the storage system does all the work, freeing up the storage administrator to focus on other storage related responsibilities.

Quality of Service

A prerequisite for any storage consolidation strategy is the ability to host multiple applications on a single storage platform, without allowing the actions of one set of users to affect quality of service (QoS) of others. Potential problem areas for shared storage access include:

- Workloads with I/O and cache conflicts, such as online transaction processing (OLTP) and data warehousing
- Tiered storage access restrictions, such as development and production applications
- Peak processing demands of critical applications versus maintenance activities like backup or a database reorganization


Virtual Storage Platform supports logical partitioning. The Hitachi Virtual Partition Manager allows resources (internal and externally attached) such as capacity, cache and ports to be dynamically partitioned into virtual machines (see Figure 6). Each virtual machine has its own virtual serial number (for asset tracking and chargeback purposes). Up to 32 of these virtual machines (logical partitions or LPARs) can be created on Virtual Storage Platform. Each partition is separately managed and password-protected, to provide better resource allocation and enhanced protection through isolation between the various partitions.

Each Virtual Storage Machine is a fully autonomous environment, with its own cache, ports, disks, storage administrator and serial number, with support for logging to a SYSLOG server for auditing. This capability enables users to build different internal service levels and separate test from production environments. It also helps to reduce the costs for users that previously, for data security reasons, may have required separate storage systems.

The characteristics of each Virtual Storage Machine can be dynamically adjusted using policies. For example, this enables an application to receive more resources during peak processing and fewer resources during off-peak periods.

Virtual Storage Machines allow multiple users to safely share the resources of Virtual Storage Platform, including all attached and virtualized multivendor capacity, without the risk of activities in one region affecting the performance, availability or privacy of other users.

Figure 6. Resource groups are the next-generation version of Virtual Partition Manager and will allow greater administrator granularity and flexible partitioning.



End-to-end Secure Virtualization for VMware

Virtual Storage Platform along with its predecessors, Universal Storage Platform V (USP V) and Universal Storage Platform VM (USP VM), are enterprise-class storage systems that can virtualize external storage and are certified to work with VMware ESX Server. The abilities to securely partition port, cache and disk resources, and to mask the complexity of a multivendor storage infrastructure, make Virtual Storage Platform the ideal complement to a VMware ESX Server environment.

Virtual Storage Platform fully supports the vStorage API for Array Integration (VAAI) for integration with VMware. This API significantly enhances the communication between the ESX server and Virtual Storage Platform. Integration with VMware Midnight Riders via the VAAI combined with Hitachi Command Suite provides a robust platform for support of VMware environments.

With up to 1024 virtual ports for each physical Fibre Channel port, Virtual Storage Platform provides the connectivity to support large VMware server farms. Each virtual port is allocated a storage partition, called a “host storage domain.” The fully secure host storage domains ensure applications in one VMware guest cannot access data belonging to another guest machine.

Virtual Storage Platform will support vMotion using the High Availability Manager vMotion Anywhere Support that is built on Hitachi High Availability Manager and uses Hitachi Dynamic Link Manager on ESX (see Figure 7). This support is built on active-active storage controllers with the hosts using Hitachi Dynamic Link Manager owner and non-owner path management configurations. Thick, thin or virtualized volumes will fail over to provide continual access to data stored on both internal and virtualized storage without requiring a restart of the host-based application.

Host Storage Domains

Host storage domains are created from the common pool of internal and virtualized external Hitachi and third-party storage attached to Virtual Storage Platform. This gives VMware guests access to a wide range of storage with different price, performance and functionality profiles. Partitions allocated to a VMware guest, through its host storage domain, can be migrated between different tiers of storage, according to the needs of the business application. Data can also be replicated locally, or remotely, to accommodate business continuity needs. Movement and replication of host storage domain partitions is nondisruptive and totally transparent to applications running in the VMware guest machine.

Each host storage domain has unique LUN addressing, including a bootable LUN0. This enables VMware guests with different host operating systems to securely share a single physical port without risking LUN and data integrity. Virtual Storage Platform cache can also be partitioned into virtual storage machines, which support logical partitioning of ports, cache and disks. Virtual storage machines have the added benefit of their own storage partition administrator and virtual serial number, for chargeback. With up to 1024 virtual ports for each physical Fibre Channel port, Virtual Storage Platform provides the connectivity to support large VMware server farms.

Virtual Storage Platform gives applications running on VMware guest machines access to the most advanced storage system functionality available, including dynamic tiering, thin provisioning, common management of all virtualized storage, content management services and storage security services. Combining VMware ESX Server and Virtual Storage Platform overcomes physical infrastructure inflexibility and delivers true end-to-end virtualization.




Figure 7. Hitachi High Availability Manager vMotion Anywhere Support is depicted using Hitachi Dynamic Link Manager on ESX.

Mainframe Storage Virtualization

Virtual Storage Platform gives mainframe users access to the same virtualization technology available to open systems platforms. Virtual Storage Platform offers both IBM FICON® and Fibre Channel connectivity. This gives mainframe users access to virtualized solid state drives (SSD), SAS and SATA storage solutions, without requiring expensive protocol converters and network extenders.

For example, SATA storage capacity is generally unavailable to mainframe applications, because the modular systems hosting SATA disks do not support FICON. Using Hitachi Data Systems 3390-Memulation mode, Virtual Storage Platform provides mainframe users with access to low-cost SATA storage. This enables dataset migration processes, which would otherwise consume costly mainframe processing cycles performing compression, to write directly to SATA storage. With the availability of cost-effective SATA storage, mainframe migration solutions, such as Data Facility Hierarchical Storage Manager (DFHSM), can target SATA disks instead of tape. This eliminates tape handling and latency issues and provides greater responsiveness during recalls and restores. Hitachi ShadowImage Heterogeneous Replication also offers a DFHSM alternative, offloading mainframe processing cycles by archiving volumes in the storage controller.

With Hitachi Dynamic Provisioning, mainframe users will be better able to manage storage utilization. A large amount of storage can be allocated to an application in anticipation of future need but the storage will not actually be provisioned for the application until it is actually needed.

Virtual Storage Platform also gives mainframe users access to controller-based, Fibre Channel data replication technology. This can substantially reduce mainframe business continuity costs by eliminating the need for expensive FICON extenders between local and remote data centers. Similarly, mainframe-based virtual tape solutions that currently rely on FICON for connectivity can direct writes to any virtualized local or remote Fibre Channel-based storage system.

Business Continuity

Virtualization brings the availability benefits of data replication technology within reach of more enterprise applications. Virtual Storage Platform supports a single suite of replication tools for all virtualized multivendor storage. This enables data stored on a high-end storage system at a primary site to be cost-effectively replicated to a modular SATA storage system virtualized behind Virtual Storage Platform at a remote site. Similarly, performing in-system mirroring of data on Tier 1 storage to a legacy device can speed local recovery. By lowering the cost of business continuity, Virtual Storage Platform makes these solutions available to more enterprise applications.

The Hitachi business continuity framework combines local high availability and remote disaster recovery capabilities in a centrally managed infrastructure (see Figure 8). Hitachi Universal Replicator performs asynchronous replication of any storage system virtualized by Virtual Storage Platform, across any distance. Hitachi TrueCopy Synchronous supports synchronous remote replication between Hitachi storage systems; ShadowImage Heterogeneous Replication delivers fast, nondisruptive full volume copy to mirror mission critical data locally; and Tiered Storage Manager provides policy-based data lifecycle management across all tiers of storage.




Figure 8. Business continuity offerings from Hitachi Data Systems offload CPU processing and optimize bandwidth utilization through disk journals with unique pull technology. These offerings leverage virtualization that uses a common engine for internal and external replication on open and mainframe systems.

Surviving Link Failure


Universal Replicator uses disk-based journaling and pull-style replication, enabling it to survive temporary communication interruptions for long periods, without the risk of data loss. If the replication link between the primary and remote sites fails, changes are still logged in the local journal. When the replication link is restored, changes are transmitted, without interrupting the protection process or the application. Virtual Storage Platform preserves data consistency and supports more current recovery point objectives, even during outage situations. These processes also help when balancing replication bandwidth against data currency objectives.

Three Data Center Business Continuity Solution

Virtual Storage Platform makes three data center business continuity an affordable reality for a broad range of enterprises and applications. Using a combination of TrueCopy Synchronous and Universal Replicator, cascade, pass-through and multitarget replication configurations are possible (see Figure 9).


Cascade configurations provide complete copies of application data at an intermediate data center, to support server clusters and application failover. Pass-through configurations store only journal volumes at the intermediate site, reducing business continuity cost. Multitarget configurations offer the best protection against local and regional disasters, albeit at higher cost.

Figure 9. Three data center replication configurations ensure data is always available.



Four Data Center Business Continuity Solution

Virtual Storage Platform also makes four data center business continuity an affordable reality for a broad range of enterprises and applications. Using a combination of TrueCopy Synchronous and Universal Replicator with Virtual Storage Platform, administrators can create a combined cascade and multitarget replication configuration (see Figure 10).



Hitachi High Availability Manager

High Availability Manager (HAM) enables open systems server-based applications to fail over between two Virtual Storage Platform storage systems in the event of an I/O failure to the primary storage system (see Figure 11). This failover may be accomplished without restarting the application. This capability provides continuous accessibility and availability to data located on internal or virtualized (i.e., external) storage. HAM is designed to ensure system availability for organizations running 24/7/365 mission critical applications.

Both storage systems in the HAM configuration must be of the same type, (i.e., Virtual Storage Platform to Virtual Storage Platform, USP V to USP V, USP VM to USP VM); however, they do not have to be the same exact configuration. The storage systems involved in the high availability configuration form a HAM cluster. The two Virtual Storage Platform storage systems of the HAM cluster may be fully active in supporting separate applications and they may be configured to fail over to the other system.




Figure 11. High Availability Manager ensures system availability.

Security and Compliance

Storage Navigator 2 has improved the security available for restricting and controlling access for managing Virtual Storage Platform. It provides role-based access control (RBAC) to provide for granular separation of duties between storage administrators, security administrators and maintenance personnel. It uses external authentication of privileged users, such as storage administrators, and allows for the storage system to integrate into the organization's existing security and compliance infrastructure (e.g. existing authentication data) using directory services environments like Active Directory. This enables administrators to provision and deprovision storage management access using their existing authentication infrastructure. It can use the RADIUS protocol to provide the flexibility of authenticating against a wide variety of authentication sources. Active Directory is the most widely used external authentication service. However, with support for the RADIUS protocol, it allows use of EMC's multifactor authentication service (SecurID).

Virtual Storage Platform supports the Fibre Channel Security Protocol (FC-SP) for end-to-end authentication of application data. Hitachi virtual storage machines further enhance security by offering rigid segregation of individual user and application cache, port and disk resources. These features are being evaluated for the Common Criteria EAL2 certification. Common Criteria certification gives organizations the assurance that these solutions conform to the highest internationally recognized standards for security.

Every Virtual Storage Platform is encryption capable and only needs a software license to enable the encryption capability. The encryption uses the recently approved XTS-AES 256 Bit Encryption for increased data security in environments like healthcare, financial services, government and military environments. Each Virtual Storage Platform is capable of supporting up to 32 encryption keys per platform. This allows for encryption to be used as an access control mechanism within the storage system and for different classifications of data to be stored on the same storage system, with encryption providing a data leakage prevention mechanism. There is no performance impact when using encryption. Because the encryption is done in the controller there is no major reconfiguration to the fabric or hosts required. It is completely transparent to the upstream host and application environments, and it can be used by open systems and mainframe environments. Encryption can be applied to any media type within the storage system with no measurable increases in heat output or energy consumption.

The Green Data Center

Virtual Storage Platform offers a highly effective foundation for a multifaceted, environmentally friendly data storage strategy. This platform underscores the alliance of cost-effective data management and sustainable IT in one efficient footprint. Hitachi has implemented the latest in sustainability characteristics in Virtual Storage Platform:

- 2.5 in. small form factor disk drives
- Next-generation SAS hard disk drive technology
- Larger, flash-protected cache
- Faster processors and high speed memory
- Interconnected blade-based architecture
- Fan speed control
- Front-to-back airflow
- Simplified and more efficient storage management
- Dynamic tiering to automate data placement

These characteristics have resulted in a design that consumes less power per floor space and per capacity utilized. With cache and control storage backed up to solid state disk, the battery size has been reduced to only provide power to cache and the SSD drive until the data from cache is written to the backup disk.

With an abundance of environmental benefits, Virtual Storage Platform is tailored to minimize costs and foster eco-friendly, sustainable storage practices. In particular, its high-density design allows for more efficient use of power, cooling and floor space. These reductions have a direct correlation to lower operating expenses, as well.



Hitachi Data Systems Global Solution Services and Best Practices

Hitachi Data Systems Global Solution Services (GSS) has developed specific services to help IT staffs install, configure and customize the technologies and services associated with data center transformation and Virtual Storage Platform (see Table 4). With the help of GSS, IT staffs can plan, design, implement, integrate, optimize and manage new solutions to meet specific business requirements. They can upgrade existing technologies to accommodate changing workloads, solve performance problems and perform capacity planning to ensure future growth is accommodated without infrastructure disruption.

Table 4: Global Solution Services Offerings for Hitachi Virtual Storage Platform

Assess and Consult

- Data Classification Service
- Storage Economics Strategy Service
- Transformation Services

Tiered Storage and Storage Management

- Data Migration Services
- Design and Implementation Service for Hitachi Dynamic Provisioning
- Disk Physical Destruction and Recycling Service
- Implementation Service for Hitachi Command Director
- Implementation Service for Hitachi Device Manager
- Implementation Service for Hitachi Dynamic Link Manager
- Implementation Service for Hitachi Tuning Manager
- Implementation Service for Hitachi Storage Capacity Reporter, powered by APTARE®
- Implementation Service for Hitachi Storage Services Manager
- Onsite Disk Data Eradication Service
- Tiered Storage Planning and Design Services
- Upgrade Service for Hitachi Storage Services Manager

Business Continuity

- Basic Replication Design Service
- Hitachi Storage Cluster for Microsoft® Environments Services
- Hitachi Storage Cluster with Microsoft® Hyper-V™
- Implementation Service for Hitachi High Availability Manager
- Implementation Service for Hitachi Replication Manager
- Implementation Service for Hitachi ShadowImage® In-System Replication
- Implementation Service for Hitachi SplitSecond™ Solutions for Hitachi Protection Manager
- Implementation Service for Hitachi SplitSecond™ Solutions for Hitachi Protection Manager for Microsoft® Exchange
- Implementation Service for Hitachi TrueCopy® Remote Replication
- Implementation Service for Hitachi Universal Replicator
- Implementation Service for Three Data Center Business Continuity Solutions
- Onsite Disk Data Eradication Service
- Remote Copy Planning and Design Service

Managed Services: Remote Storage Management

- Remote Operational Services
- Storage Scorecard Service

Support Services

- Select Care Program



VIRTUAL STORAGE PLATFORM ARCHITECTURE

Virtual Storage Platform is the most intelligent and powerful enterprise storage system in the industry. It is based on the Hi-Star-E network architecture: a fifth-generation implementation of the switch architecture first introduced in Hitachi Lightning 9900™ Series enterprise storage systems. With this new generation of the switch architecture and the newly designed internal components, Virtual Storage Platform supports over 2.37x I/Os per second, up to 255PB of internal and external virtualized storage capacity and up to 1TB of directly addressable cache (see Table 5).

Hi-Star-E Network — the Industry's Highest Performing Storage Architecture

At the heart of Virtual Storage Platform is the next-generation, high performance, fault tolerant and truly nonblocking switched network that is used to communicate and transfer information between the various components (see Figure 12). The performance of Virtual Storage Platform will vary based on the actual components that are installed and active in the system. All cache data and control data paths are part of the switched network. This switched network is used to communicate and transfer information between the various components of Virtual Storage Platform.

Hi-Star-E network is made up of PCI Express grid switches (GSW). Each switch has 24 high speed ports that each support 1GB/sec send and 1GB/sec receive for a total bandwidth of 2GB/sec per port. Eight of the ports are used to connect to ports on the front-end director (FED) and back-end director (BED). These ports pass both metadata and user data traffic, which is intermixed. Four of the ports are used to connect to ports on the virtual storage directors (VSDs) and only pass job request traffic or system metadata, e.g., cache address for a target slot. Eight of the grid switch ports are attached to the data cache adapters (DCA). These ports pass user data reads and writes as well as copies of control data from the virtual storage director (VSD) boards. The final four ports are used to cross connect

the grid switches to the corresponding grid switches in the second chassis, if used.

Each switch has only one connection per component. However, because the switches are installed in pairs, all connections between the components and the Hi-Star-E network are redundant, which provides higher performance as well as protection in case of a link failure.

Each chassis can have two or four grid switches. The DCA boards, VSD boards, BED boards and FED boards connect to the Hi-Star-E network using PCI Express 4-lane paths, providing a 1GB/sec send path and a 1GB/sec receive path. Each GSW port has 24 full duplex ports resulting in aggregate bandwidth for each GSW of 24GB/sec send and 24GB/sec receive.

The primary internal subsystem components that make up Virtual Storage Platform include the VSDs, DCAs, BEDs, FEDs and the GSWs, as discussed above. Each contributes to achieve the overall internal bandwidth of the Hi-Star-E network by bringing incremental paths (and bandwidth) to the internal network. A subset of the components can be installed and configured to meet the specific requirements of a particular customer or application.

Table 5. Hitachi Virtual Storage Platform Specifications

Characteristics	Virtual Storage Platform (minimum configuration)	Virtual Storage Platform (maximum configuration)	
Number of Frames or Racks (basic control frame or rack plus optional array frames or racks and second control frame or rack)	1	6	
Number of Grid Switches in Hi-Star-E Network	2	8	
Raw Cache Bandwidth (GB/sec)	8	128	
Raw Control Memory Bandwidth (GB/sec)	8	32	
Aggregate IOPS (60:40 mix Read:Write – 8Kb)		237,000	
Fibre Channel Port Performance	2/4/8Gb	2/4/8Gb	
Supported Logical Devices (LDEVs)		65,280	
Number of Disks per System (min./max.) (3.5")		1280	
Number of Disks per System (min./max.) (2.5")		2048	
Number of Disks per System (min./max.) (SSD)		256	
Number of Spare Disks per System (min./max.)	1	256	
Max. Internal Raw Capacity (2TB disks)		2.52PB	
Max. Internal and External Capacity (PB)		255PB	
Max. Useable Internal capacity (TB) Open Systems, RAID-5 (7D+1P)		2,192.2 (Open-V)	2,080.8 (IBM® z/OS® 3390M)
Max. Useable Internal Capacity (TB) Open Systems, RAID-6 (6D+2P)		1,879 (Open-V)	1,779.7 (z/OS 3390M)
Max. Useable Internal capacity (TB) Open Systems, RAID-1+0 (2D+2D)		1,256.6 (Open-V)	1,190.2 (z/OS 3390M)
Cache Boards			
Number of Boards	2	16	
Max. Board Capacity (GB)	32 (64)	32 (64)*	
Capacity (GB)	16	512 (1,024)*	
Control Memory			
Number of Boards	2	8	
Max. Board Capacity (GB)	4	4	
Capacity (GB)	8	32	
Front-end Director Boards	2	24	
Fibre Channel Ports	0	192	
Virtual Ports per physical port	1,024	1,024	
IBM FICON® Ports	0	176	
Fibre Channel over Ethernet (FCoE) Ports	0	(88)*	
IBM ESCON® Ports	n/a	n/a	
Max. Hitachi Virtual Storage Machines	1	32	
Back-end Directors (SAS 6Gb)	2	8	

* Cache capacity and FCoE will be added in an upcoming release.

Data Cache and Control Memory

The complete set of eight DCAs in a single chassis system has up to 32 links connecting the data cache to the Hi-Star-E network. This allows for a total 64GB/sec of full-duplex aggregate switch cache data bandwidth for a single controller chassis. Each controller must have at least two data cache adapters, i.e., one data cache adapter feature.

The connections between DCAs and the GSWs only see user data reads and writes as well as control memory updates. It is not necessary for the data cache adapters to have their memory fully populated to achieve the maximum internal bandwidth; however, each DCA must be populated with the same amount of memory.

Data cache in Virtual Storage Platform is divided into two equal areas, and data to be written and the copies of control storage are always duplicated across logic and power boundaries. This redundancy ensures full data integrity in the event of a power-related failure.

Data Cache

The data cache memory is used to cache user data blocks that either come from disk drives via the back-end director during a read or to cache data from a front-end director as part of a data write operation. All write data is mirrored in cache. However, read data is not mirrored; there is only one copy of cached data when performing read operations from disk. Each DCA can contain up to 32GB (64GB) of data cache memory with a total of 256GB (512GB) per controller chassis.

The complete set of data cache adapters in a single chassis has up to 32 links (eight DCAs with four links

each) connecting the cache adapters to the Hi-Star-E network for a total 64GB/sec of full-duplex aggregate cache to switch bandwidth in a single controller chassis.

In the case of a power failure, component failure or other unscheduled outage, there is a new protection and recovery mechanism to back up cached user data that has not been




Figure 12. Virtual Storage Platform employs Hi-Star-E network for high performance data movement.

destaged to disk and also to protect the state of the system, i.e., control memory. As part of the data cache adapters, all write data that has not been destaged to disk is written to an onboard 32GB SSD device, including a backup of control memory (see Figure 13). The DCA has an onboard battery that

allows the cache adapter microprocessor to destage the data to the embedded SSD on system shutdown.

Data Cache Adapter (DCA)




Figure 13. Data cache adapter with SSD backup device allows for de-staging of data on system shutdown.

Control Data

The operational control memory of Virtual Storage Platform is physically located on the virtual storage directors (VSDs). This centralization of control memory with the main processors on the VSDs enables the main processors to access the control data memory at internal memory speeds, resulting in the elimination of processing and transfer latency. This means that there is no internal bandwidth usage for access to control data, which results in faster processing.

The VSDs with their control memory and main processors store and manage all of the internal operational metadata and state information about the system. The metadata related to array groups, internal and external LDEVs, runtime tables, mappings, data for various software products and the overall state of the system are stored, referenced and executed from this memory. This metadata is distributed as needed to the appropriate I/O offload processors.

There can be a total of 4GB of control data memory on a single VSD and a total of 16GB of control data memory in a single controller chassis. This control data is not mirrored, so the entire 16GB of memory capacity is available for operations use in a single chassis. However, this control data is backed up to cache on the first data cache adapter pair. The control data in the data cache adapter cache is not a mirror but rather a backup, which is stored in an optimized recoverable format. In addition, since any data written to data cache is mirrored, the backup of the control data is also mirrored. As a result, there are three copies of control data: the operational copy in control memory on the virtual storage directors and two (mirrored) backup copies on the data cache adapters.

Advanced Data Cache Algorithms Provide Read-ahead for High Performance Sequential Reads

Virtual Storage Platform has a variety of advanced cache algorithms and software that combine to deliver exceptional performance. Read clustering uses built-in heuristics to determine whether data is being accessed sequentially. If sequential reads are detected, Virtual Storage Platform allocates a set of 256KB cache slots for the sequential data to be written to or read from the LDEV. Sequential pre-fetch will read several 256KB

chunks into cache slots without being instructed to do so by the controlling FED so that subsequent read requests will be satisfied immediately from memory, speeding I/O.

Hitachi Cache Residency Manager Allows Datasets to Be Permanently Placed in Data Cache

The Hitachi cache residency manager feature allows administrators to dynamically lock and unlock data in cache in real time. The software makes it possible to allocate a portion of cache to specific data. When activated, read and write functions are performed from cache, eliminating performance delays from disk latency. Administrators can add, delete or change which datasets cache residency manager feature manages at any time.

In IBM mainframe environments defined by the logical volume image (LVI), cache data can be as small as a single track or as large as a full disk volume. Using cache residency manager in conjunction with Hitachi Cross-OS File Exchange increases data transfer speed, improving the performance of mainframe-to-open-systems data exchange.

Front-end Design


Virtual Storage Platform supports Fibre Channel for open systems and SAN connectivity and FICON for IBM z/OS platform connectivity. Virtual Storage Platform with two chassis supports up to 128 Fibre Channel or FICON ports from up to 8 FEDs per controller chassis. Additional ports, up to 64, can be added by replacing one to four BEDs per controller chassis with FEDs. With the additional FEDs, Virtual Storage Platform can support up to 192 ports. The Fibre Channel and FICON ports support industry standard 2Gb/sec, 4Gb/sec and 8Gb/sec transfer rates.

Unique port virtualization technology dramatically expands connectivity from open systems platforms to a Virtual Storage Platform (see Figure 14). Each physical Fibre Channel port supports 1,024 virtual ports. Each virtual port has access to a private storage partition called a host storage domain, which includes a bootable LUN0.

Connectivity from all major open systems hosts is supported, including :

- HP: HP-UX, Tru64 UNIX, OpenVMS
- IBM: AIX
- Microsoft: Windows Server® (2000, 2003, 2008)
- Novell: NetWare, SUSE Linux Enterprise Server
- Red Hat: Enterprise Linux
- Oracle: Sun Solaris
- VMware: ESX Server

Figure 14. Port virtualization technology supports up to 1,024 virtual ports for each Fibre Channel port.



Connecting to Mainframe Hosts

Virtual Storage Platform supports both FICON and Fibre Channel connectivity to mainframe hosts. Each FICON FED supports eight FICON connections and can autonegotiate to one of three speeds: 2Gb/sec, 4Gb/sec and 8Gb/sec. Each FICON feature consists of two FICON FEDs for a total of 16 ports per feature and four features per chassis or 128 FICON ports in a dual chassis configuration with a total bandwidth of 204.8GB/sec (128 ports @ 8Gb/sec full duplex).

The FICON FED modules of Virtual Storage Platform feature negotiating logic that automatically sets the port speed to match the host's CHPID (channel) speed.

Open systems and mainframe hosts can connect simultaneously to Virtual Storage Platform (see Figure 16). Supported mainframe host operating environments include: the IBM OS/390®, z/OS, z/OS.e, VSE/ESA, VM/ESA®, z/VM®, z/VSE®, Linux for S/390® and z/Series®.




Figure 15. Virtual Storage Platform provides FICON and Fibre Channel attachment to mainframes and open systems hosts.

High Performance, SAS Back-end Design

Virtual Storage Platform is available with a choice of disk drives ranging from high performance SSDs to SAS drives to high capacity SATA drives (see Table 6). Both large form factor (LFF) 3.5 in. drives and small form factor (SFF) drives are supported. All three drives types use a 6Gb/sec interface to the BED. The SATA drives actually have a single SATA port; however, the SATA disk canister for Virtual Storage Platform has a logic board that creates two SAS ports from the single SATA port. This logic board also does a read-after-write check to guarantee that the data was accurately written since SATA drives don't validate proper functioning of the write heads as SAS drives do.

Virtual Storage Platform supports command tag queuing to offload much of the seek optimization processing from the storage system to the disk drives. This allows more simultaneous back-end I/O operations and offers greatly improved back-end performance.

BEDs control the transfer of data between disk drives and cache. Each BED has four SAS ports. Each port accepts a cable that has two independent 6Gb/sec SAS links within it (i.e., 2W). A BED feature is composed of a pair of BEDs to provide redundancy and high performance. Each

Table 6. Hitachi Virtual Storage Platform Drive Types

Drive Type	Size	Capacity	Speed (RPM)
SSD	2.5"	200GB	n/a
	3.5"	400GB	n/a
SAS	2.5"	300GB	10K
	2.5"	600GB	10K
	2.5"	146GB	15K
SATA	3.5"	2TB	7.2K

Virtual Storage Platform chassis supports one or two BED features and a total of four BEDs features in a fully configured, two chassis system. These eight BEDs provide 64 6Gb/sec SAS connections with concurrent data transfers to and from the disk drives. Each four-port BED has two data accelerators and two SPC SAS I/O controller processors and two banks of local RAM. There is a lot of processing taking place on each back-end directors.

RAID Implementation Advantages

Hitachi BEDs are the most advanced in the enterprise storage industry and employ advanced algorithms to manage performance. The intelligent controllers provide disk-interface and RAID management circuitry, offloading these tasks to dedicated embedded processors. Each Virtual Storage Platform BED supports RAID-1+0, RAID-5 (parity RAID) and RAID-6 (see Table 7). All user data disks in the system are defined as part of a RAID array group of one type or another.

Table 7. Virtual Storage Platform Back-end Director Support of RAID Configurations⁶

RAID Type	Supported Disk Configuration
RAID-1+0 Mirroring	2D+2D 4D+4D
RAID-5 Distributed Parity	3D+1P 7D+1P
RAID-6 Redundant Distributed Parity	6D+2P

⁶ When using Open V-volumes for RAID-5, RAID-6 or RAID-1+0 configurations, Virtual Storage Platform uses two adjacent 256KB chunks from the same disk before moving to the next disk so the effective chunk size is 512KB. For example, in a RAID-5 7D + 1P configuration, the stripes are actually 3,584KB (7 x 512KB).

RAID-1+0

RAID-1+0 is available in two data plus two data (2D+2D) and four data plus four data (4D+4D) disk configurations. The configurations include a rotating copy, where the primary and secondary stripes are toggled back and forth across the physical disk drives for performance. RAID-1+0 is best suited to applications with low cache-hit ratios, such as random I/O activity and with high write to read ratios.

RAID-5 Distributed Parity

RAID-5 disk arrangements consist of four disks (3D+1P) or eight disks (7D+1P). Data is striped across disks similar to RAID-1+0. However, RAID-5 keeps parity information on each stripe of data for fault resilience. If a failure occurs, the contents of the failed block can be recreated by reading back other blocks in the stripe and the parity. Parity information is distributed throughout the array to minimize bottlenecks when rebuilding data from a failed disk. The overhead of RAID-5 is equivalent to one disk drive, regardless of the size of the array. RAID-5 is best suited to applications using mostly sequential reads.



Concatenated RAID-5 Array Groups

Concatenated array groups enable users to stripe LDEVs across more disks than are in a single RAID-5 array group for better performance. The feature supports the concatenation of two or four RAID-5 7D+1P array groups. An LDEV created on a standard RAID-5 array group becomes a LUN when it is mapped to a host port. For concatenated array groups, the LDEV, striped across multiple elements in the joined array groups, forms a virtual device (VDEV). Mapping the VDEV to a host port creates the LUN.

RAID-6 Redundant Distributed Parity


Like RAID-5, RAID-6 stripes blocks of data and parity across an array of drives. However, RAID-6 maintains redundant parity information for each stripe of data. This redundancy enables RAID-6 to recover from the failure of up to two drives in an array, i.e., a double fault. Other RAID configurations can only tolerate a single fault. As with RAID-5, performance is adjusted by varying stripe sizes. RAID-6 is good for applications using the largest disks and performing many sequential reads.

Figure 16. Virtual Storage Platform is available with RAID-1+0, RAID-5 and RAID-6.




RAID-5: (7D + 1P)

Data is striped with parity over RAID members



RAID-6: (6D + 2P)

Data is striped with parity over RAID members



RAID Rebuild

If a disk in a RAID group fails, the information on the failed drive is rapidly and automatically rebuilt on a hot spare drive. Virtual Storage Platform also allows logical volumes to be expanded online, and for administrators to set rebuild priority, based on the daily processing cycle and RAID configuration installed.

Connecting External Storage

One of the most significant features offered by Virtual Storage Platform is the ability to connect external, multivendor storage (see Table 8). This creates a single virtualized pool of capacity for administrators to centrally manage using Hitachi storage management software. Virtual Storage Platform supports a wide range of third-party storage systems.

Table 5. Externally Attached Storage Systems

Vendor	Product
Hitachi	Adaptable Modular Storage
	Adaptable Modular Storage 2000 family
	Workgroup Modular Storage
	Simple Modular Storage 100
	Thunder 9500™ V Series
	Thunder 9200
	Lightning 7700 Series Lightning 9900™ Series Lightning 9900 V Series Universal Storage Platform® Universal Storage Platform V Universal Storage Platform VM Network Storage Controller Virtual Storage Platform
3PAR	InServ S400
ADTX	Array MasStor L
EMC	Symmetrix 3830
	Symmetrix 3930
	Symmetrix 8530
	Symmetrix 8730
	Symmetrix 8830
	Symmetrix DMX-3
	Symmetrix DMX-4
	Symmetrix DMX800
	Symmetrix DMX1000
	Symmetrix DMX2000
	Symmetrix DMX3000
	Symmetrix VMAX
	CLARiiON CX3 CLARiiON CX4 CLARiiON CX200 CLARiiON CX300 CLARiiON CX400 CLARiiON CX500 CLARiiON CX600 CLARiiON CX700
Fujitsu	Eternus3000
	Eternus4000
	Eternus8000
	FibreCAT CX3-80


Vendor	Product	
Gateway	840FC	
HP	StorageWorks EVA3000 StorageWorks EVA4000/4100 /4400 StorageWorks EVA5000 StorageWorks EVA6000/6100 StorageWorks EVA8000/8100 StorageWorks XP10000 StorageWorks XP1024	StorageWorks XP12000 StorageWorks XP128 StorageWorks XP20000 StorageWorks XP24000 StorageWorks XP48 StorageWorks XP512 Storage Works XP256 SVS200
IBM	TotalStorage ESS 2105 System Storage® DS4700 System Storage DS4500 (FAStT900) System Storage DS4400 (FAStT700) System Storage DS4300 (FAStT600) System Storage DS3400	TotalStorage DS4100 (FAStT100) System Storage DS4800 System Storage DS6000 System Storage DS8000 SAN Volume Controller (SVC) XIV® Storage System
NetApp	FAS3000	FAS6000
Nexsan Technologies	ATABeast SASBeast	SATABeast SATABeast2
Pillar Data systems	Axiom 300 Axiom 500	Axiom 600
Sun Microsystems	StorEdge T3B (T3BES and T3BWG) StorEdge 6120/6320 StorEdge 3510 StorEdge 3511 StorEdge 6130 StorEdge 6140 StorEdge 6540 StorEdge 6580//6780 StorageTek BC82 StorageTek D178 StorageTek D240	StorageTek D280 StorageTek FlexLine 240 StorageTek FlexLine 280 StorageTek FlexLine 380 StorageTek SVA StorageTek 9960/9910 StorageTek 9980/9970 StorageTek 9990/9985 StorageTek 9985V StorageTek 9990V StorageTek 2540
Xiotech	Emprise 5000	

For the latest information see the Externally Attached Storage Systems on www.hds.com.

Packaging and Configuration

Virtual Storage Platform is available in configurations from one to six cabinets (see Figure 17), depending on capacity, performance and connectivity needs. The base configuration is easily and nondisruptively upgraded by adding features, such as virtual storage directors, data cache, FEDs, BEDs, grid switches and disk drives to the existing chassis, adding a second chassis and adding cabinets with disk modules. A single cabinet chassis with two VSDs, two DCAs, two FEDs, two BEDs and two grid switches serves as the base configuration and manages up to two array frames connected via 6GB/sec links (see Figure 18). A second chassis can be added and connected to the first chassis via the Hi-Star-E coupler, which is a 32GB/sec connection that tightly couples the chassis directly at the Hi-Star-E network switches.

Figure 18. The disk chassis can hold 3.5 in. or 2.5 in. drives.



Fans

Drives 2.5" or 3.5"
form factor
(Front and Back)



Figure 17. Virtual Storage Platform is available in configurations from one to six cabinets.

The Control Chassis

The control chassis contains the virtual storage directors, FEDs, BEDs, data cache, grid switch boards and up to two disk chassis, as well as power supplies (see Figures 19, 20 and 21).

The disk chassis contains the hard disk drives, power supplies, battery modules, cooling fans and communication interfaces.

Figure 19. Control Chassis (front view)




Figure 20: Control Chassis (back view)





Figure 21. Virtual Storage Platform Racks Showing Control and Disk Chassis



Power Subsystem

The power subsystem used in Virtual Storage Platform has redundant power supplies in the control frames and array frames. The power supplies are hot pluggable, hot swappable and N+1 redundant. This means power supplies can be removed or installed without affecting system operation.

Each controller chassis in Virtual Storage Platform consists of two redundant controller halves. Each half contains all physical and logical elements needed to sustain processing. Both storage halves support alternate connection path schemes, so if one fails, the other can continue processing for the entire storage system.

Redundant power supplies are provided to power each half of the controller. The power supplies in each half are capable of providing power for the entire system, in the event of a power supply failure.

The System Monitoring Network

The Virtual Storage Platform system monitoring network connects to all cabinets and reports on a variety of system conditions, such as component failure, fan speed, voltages and temperature. The network enables efficient predictive maintenance, so that components can be replaced before they fail. It also expedites system repair after a failure, through the failure alert system.

The system-monitoring network passes information to the service processor or remotely across a private LAN on a UNIX or Microsoft Windows console. Failure information is also available via Simple Network Management Protocol (SNMP). Copies of all system status and alerts are sent to Hitachi Data Systems Customer Support Centers via the Hi-Track® Remote Monitoring system

High Availability

In addition to offering the most scalable and highest performing platform for enterprise storage, Virtual Storage Platform places maximum emphasis on high availability computing to satisfy the needs of the most demanding enterprise applications. Hitachi Data Systems offers a combination of industry unique solutions and best-in-class third-party software integration to deliver unbeatable resilience and availability.

- Heterogeneous asynchronous remote replication over any distance
- Disaster recovery
- Copy-on-write snapshots
- Mainframe replication for IBM Geographically Dispersed Parallel Sysplex™ (GDPS™)
- Nondisruptive microcode upgrades
- Nondisruptive hardware upgrades
- Automatic failover architecture with redundant, hot-swappable components
- Dual data paths and dual control paths connecting every component
- Active-active dual-ported disk drives
- Mirrored cache for all write data
- Nonvolatile backup of cache using a combination of battery and flash disk drives
- Hi-Track Remote Monitoring, which immediately sends notification of any incident that may affect availability to the Hitachi Data Systems Customer Support Center

Hitachi Data Systems demonstrates its confidence in the high availability of Virtual Storage Platform by offering a 100% data availability warranty.



APPENDIX A: DYNAMIC TIERING


With the availability of various classes of storage with different cost and performance characteristics, such as drive type, speed and RAID configuration, it is possible to create tiers of storage. Each tier can be based on the devices' different characteristics. The use of storage tiers enables companies to use their storage more efficiently and reduce costs. Most tiered storage solutions today require the storage administrator to manually classify the data and schedule the movement of the data. The administrator uses tools provided to receive the benefits of reduced cost and more efficient use of storage.

Tiered Storage for Files, Volumes and Pages

File tiering or file migration between tiers refers to the movement of files or content between storage tiers; the tiers may reside within the same storage system or separate storage systems connected through a network. Volume and page-based tiering or migration applies to block-based environments, which in turn are used by higher-level file and content interfaces (see Figure 22).

With volume-based tiering, a complete storage volume is moved or migrated between tiers. The tiers may include high performance, expensive storage such as solid state drives (SSDs), and good performance storage such as Fibre Channel or SAS, low cost, high capacity storage such as SATA or even tape drives. Page-based tiering offers a finer granularity of tiering within a designated volume. Any page for a volume can dynamically reside on any tier. Page-based tiering always occurs independently and transparently from the host application server(s), whether it is presented as block or file level storage.

Figure 22. Volume, Page and File System Tiering



Hitachi Dynamic Tiering

Hitachi Dynamic Tiering (HDT) takes the automation of tiered storage to a new level by enabling the management of multiple storage tiers as a single entity. It presents a new kind of virtual volume with embedded smart tiering that monitors access and moves data at the fine grain page level. It breaks the volume into pages and automatically promotes the pages that are being referenced frequently, avoiding the time and storage space to move the entire data set or file. It provides the facility to have the right data in the right place at the right time by self-optimizing for high performance and space efficiency.

Hitachi Dynamic Tiering uses a dynamic storage pool, which is divided into 42MB pages. The pages are mapped to tiers within the pool. Virtual volumes are made up of pages from this pool. The pages can be anywhere in the pool on any tier within that pool. Virtual Storage Platform, through HDT, determines which pages it is best to place in higher tiers and which should be moved to lower tiers. The pages are moved up and down between the tiers within the pool for performance optimization purposes. In operation, a HDT environment will have most of the highly referenced pages in the highest tier and the least referenced pages in the bottom tier. HDT tries to use as much of the higher tiers as possible. This can provide the effective appearance that all the data is on the fastest tier of storage. In fact, most of the data resides on lower cost, lower performance storage tiers composed of SAS and/or SATA drives. The lowest tier also contains much of the spare capacity. By default, HDT automatically initiates page migration on a periodic basis. However, the storage administrator can change this to manually initiate the migration between tiers.

Hitachi Dynamic Tiering volumes are just another kind of volume in a tiered storage architecture. Just like disk types and RAID configurations, HDT multitier volumes are another way of delivering tailored storage cost and performance service levels. These dynamically tiered volumes do, however, deliver superior service levels at less cost, maximizing service levels and minimizing storage total cost of ownership. Capacity can be dynamically added to or removed from any tier at any time, which means that sizing a tier for a pool is much easier.

Hitachi Dynamic Tiering and SSD make an excellent match. SSD storage delivers excellent performance but is relatively expensive when compared to SATA storage or even SAS storage. Organizations typically choose to purchase less SSD capacity than they do SAS or SATA. As a result, they have to decide how to make the best use of the SSD storage they do purchase. With Hitachi Dynamic Tiering, this is no longer a problem. By defining the SSD storage as the highest tier, the system will use this storage for those pages that are the most active and move less active storage to lower tiers with the lowest activity storage being moved to the lowest tier. Pages will migrate to the appropriate tier based on their usage. Because Hitachi Dynamic Tiering only moves the highly referenced pages, a smaller amount of high performance disk, such as SSDs, can significantly improve the storage system performance.

With Dynamic Tiering, the storage administrator does not need to classify data or define policies; once the tiers are configured, the storage system does all the work, freeing up the storage administrator to focus on other storage related responsibilities.

Hitachi Dynamic Tiering is very good for database applications. It can provide improved performance for databases stored in a file system by keeping the highly referenced data, such as control tables and indexes on highest storage tier, while moving less frequently reference pages to a lower tier of storage. HDT is also excellent for rich media applications. The self-adjusting system can optimize performance based on demand enabling to content service providers to manage their storage assets more effectively and allow the use of less expensive, higher capacity disk drives such as SATA.

APPENDIX B: THE EVOLUTION OF HITACHI STORAGE VIRTUALIZATION

The storage virtualization engine used by the Hitachi Universal Storage Platform family is a result of evolutionary development through successive product iterations and technology refreshes. This unique solution offers unparalleled scalability, performance, manageability, security and functionality, and makes it possible to deliver storage functionality to business users as a service.

In 1995, Hitachi Data Systems introduced the Hitachi Freedom 7700. This storage system uniquely separated control information and data by placing them in different caches accessed across different busses. This enabled support of multiple server platforms, including, mainframes, UNIX and Microsoft Windows.

In 1998 Hitachi Data Systems introduced Dynamic Optimizer, which enabled the Hitachi CruiseControl™ feature to dynamically change the configuration mapping in cache to move LUNs between RAID groups (LDEVs) in the storage system. This functionality, now known as Hitachi Volume Migration, is still unique to Hitachi Data Systems, and it laid the foundation for virtualization in the storage control unit.

In 2000, Hitachi Data Systems introduced the Hitachi Lightning 9900™ Series enterprise storage systems and the revolutionary Universal Star Network™ crossbar switch architecture. The crossbar switch increased connectivity between storage ports, data cache modules and array control processors (back-end directors), increasing scalability and connectivity. This enabled a Lightning 9900 Series storage system to replace between six and eight competing storage systems. In the eight years since the revolutionary crossbar switch debuted, major storage vendors have yet to match this innovative technology.

In 2002, the Hitachi Lightning 9900 V Series added support for up to 128 virtual ports for each Fibre Channel port, increasing SAN connectivity. Each virtual port supported a dedicated address space. This guaranteed no data leaks between applications sharing the same physical port. Virtual ports also simplified LUN management and enabled multiple applications to securely share a single physical storage device through safe multitenancy. Virtual ports also enabled shared access to a single storage system from applications running on different operating environments. UNIX, Linux and Microsoft® Windows® applications could share a single storage system and each maintains different levels of RAID protection, different performance disks and different port access priorities.

2004 saw the introduction of Hitachi Universal Storage Platform®. The Universal Storage Platform enabled lower cost and legacy external storage systems, from different vendors, to attach to the storage system through standard Fibre Channel ports. This let users move data across internal and external tiers of storage, migrate data from old to new storage systems and replicate copies across different cost tiers, without disrupting business applications. Port virtualization was extended to 1024 virtual ports per physical port, and the policy-based Hitachi Tuning Manager and Hitachi Tiered Storage Manager enabled nondisruptive data movement triggered by time, events and alerts.

In 2007, Hitachi introduced Universal Storage Platform V and Universal Storage Platform VM. These advanced storage systems incorporated higher performance front-end storage ports, flash memory in the port processors, the next-generation crossbar switch and higher performing control memory and data cache. Dynamic provisioning capabilities were added, extending the benefits of storage virtualization all the way to the business application by creating virtual storage capacity and automated wide striping for performance.

In 2010, Hitachi Data Systems introduced Hitachi Virtual Storage Platform in an effort to increase efficiency while expanding function. With a continued focus on virtualization and automation, Virtual Storage Platform introduced 3D scaling. This platform can scale up by nondisruptively integrating new processing capacity, i.e., processors, cache,

ports and internal switching to handle the growing I/O demands. It can scale out by adding storage resources to grow the base system for greater capacity, performance and connectivity. It can also scale deep by extending the advanced functions of Virtual Storage Platform to external multivendor storage; it dynamically virtualizes new and existing storage systems and makes them part of the Virtual Storage Platform pool of storage resources. With Virtual Storage Platform Hitachi Data Systems also introduced Hitachi Dynamic Tiering, which uses a dynamic storage pool divided into pages. These pages from the dynamic storage pool are used to make up virtual volumes. Each page is mapped to a tier within the pool and can be anywhere in the pool on any tier within that pool. Through Dynamic Tiering, Virtual Storage Platform determines which pages should be migrated to higher tiers and which to lower tiers to optimize performance. With the fifth generation of the internal network, the Hi-Star-E, Virtual Storage Platform introduced a whole new level of performance and expandability.

Hitachi Data Systems Corporation

Corporate Headquarters

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

Regional Contact Information

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

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

IBM, OS/390, S/390, z/OS, FICON, ESCON, AIX, VM/ESA, z/VM, z/VSE, GDPS and z/Series are registered trademarks and z/OS.e, Geographically Dispersed Parallel Sysplex and VSE/ESA are trademarks of International Business Machines Corporation.

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

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

© Hitachi Data Systems Corporation 2011. All Rights Reserved.