

Storage

HITACHI
DATA SYSTEMS

High-Availability Solutions for Open Systems

by Hubert Yoshida

For computing as critical as your business

Contents	
Introduction	1
Types of High-availability Middleware	2
Alternate Pathing	2
Important hardware considerations for alternate pathing	2
Hardware advantages of the Hitachi Freedom 5700E	3
Hardware advantages of the Hitachi Freedom 7700E	3
Alternate pathing for Sun Solaris	4
Hitachi SafePath	4
VERITAS Volume Manager	5
Alternate pathing for HP-UX	5
Alternate pathing for IBM AIX™	5
Alternate pathing for Windows NT	6
Host Failover	6
The increasing need for host-failover middleware	6
Important hardware considerations for host failover	6
Host failover for Sun Solaris	6
Host failover for AIX	7
Host failover for HP-UX 10.X	7
Host failover for Sun Ultra Enterprise Cluster HA 1.3	8
Windows NT Cluster Server (Wolfpack)	8
Parallel Database Clustering	9
MC/LockManager for Oracle Parallel Server	9
Sun Ultra Enterprise Cluster PDB (Parallel Database) Server	9
IBM AIX support of parallel databases	9
Summary	10

High-Availability Solutions for Open Systems

by Hubert Yoshida

Introduction

The term “open systems” generally refers to UNIX® and PC systems designed for workstation or single-user environments. These basic operating systems originally did not support high-availability features such as alternate pathing and host failover. But as these systems matured, their roles as network and applications servers grew, and so did the need for ensuring nonstop operations through high availability.

Vendors of “middleware”—software which is not part of the base operating system—including CLAM Associates, VERITAS, and Conley, stepped forward to fill the need for open systems high-availability features. Early UNIX vendors like Sun were very dependent on these third-party middleware vendors. Later UNIX providers, like HP and IBM, integrated high-availability features into their product offerings primarily through OEM agreements with many of the same vendors. Today, most major server vendors support middleware products for high availability.

High-availability middleware helps reduce downtime by automatically detecting faults and recovering data services on a redundant set of hardware. Without high-availability middleware, time is lost while a fault goes undetected. Once the fault is detected, a diagnose/repair/replace action must take place before data-service recovery can begin. High-availability middleware can begin an automated recovery process immediately on the redundant hardware. The recovery process without high-availability middleware involves time-consuming and error-prone manual operations, which may include resetting the SCSI bus, restarting drivers, reassigning IP addresses, recovering and restarting applications and transactions, and even rebooting.

Although the term “open systems” implies that such systems share standard interfaces for portability of applications, high-availability middleware is not portable across systems. The specific type of middleware used is highly dependent on the architecture of the systems platform and disk subsystem. This white paper provides a tutorial on high-availability middleware and reviews the different solutions available for use with different platforms. It also shows how the Hitachi Freedom Storage™ sub-systems provide maximum RAID protection and multipath connections to complete a total high-availability solution.

Types of High-availability Middleware

Three basic types of high-availability middleware reduce downtime in the event of a data path or host failure:

- **Alternate pathing (or I/O path switching).** This type of middleware automatically switches the I/O load on a failed primary path to an alternate path on the same host system.
- **Host Failover.** This type of middleware supports a cluster of host processors where one of the hosts automatically takes over the workload of any failed host in the cluster. This “take-over” includes the reassignment of networks and peripherals, as well as the restarting of applications. (Host clustering can also be used to balance workload and scale processor capability while sharing network and disk resources.)
- **Parallel Database Clustering.** This type of middleware is a special version of host failover middleware which supports major parallel database servers like Oracle Parallel Server (OPS), Informix XPS, and Sybase MPP. Clustering middleware supports distributed lock management, a feature which enables parallel database software running on separate cluster nodes to share access to the same database. If one host fails, the other hosts can take over its work. Database clusters allow a customer to grow a database incrementally simply by adding additional nodes. (With non-parallel database servers, the server has to be replaced or an additional server with another database instance has to be purchased and installed when the capacity of the original system is exceeded.)

Note: Oracle Version 8 provides its own distributed lock manager and requires a new interface layer between the distributed lock manager and clustering middle-ware. At the time of this writing, the availability of this interface layer was unknown. Hitachi is working with Oracle to address this area.

Alternate Pathing

Important hardware considerations for alternate pathing

Alternate pathing provides automatic I/O path switch-over from a failed I/O path to a backup I/O path. (An I/O path comprises every component from the host system to the I/O device, including the bus adapter, external SCSI cable, the disk controller’s port adapter, and the controller itself.) Connecting a single-controller disk subsystem to two separate SCSI connections on a single host protects against the failure of a single-host bus adapter or SCSI cable. However, it does not protect against the failure of the controller. Therefore, any I/O subsystem configured for alternate pathing should provide separate controllers for each path.

The Hitachi Freedom 5700E provides the option of a second controller for highest availability. This second controller can be invoked manually or automatically through the use of alternate pathing middleware. The alternate paths on the 5700E can operate in Hot Standby or Dual Active mode. In Hot Standby mode, one controller owns all the disks and the other controller takes control only when the primary path fails. In Dual Active mode, the disks are allocated between the two controllers so that both controllers are active with their own set of disks. If one controller path fails, the other controller takes ownership of the failed controller’s disks.

Successful alternate pathing requires close cooperation between host middleware and the I/O subsystem's microcode. Host middleware maintains a map of SCSI addresses to physical disk; it also monitors the health of the primary path. Alternate pathing middleware vendors provide this support in a variety of ways. In some cases, "signatures" are written on the shared disks to identify the pathing configuration. In other cases, SCSI code pages are used to identify valid paths. If a failure is detected, the middleware must attempt to gain control of the device through the alternate path.

The method by which the middleware regains control of the device from the original path depends on the host platform and architecture of the RAID subsystem. On some platforms, special SCSI commands are required to regain control of the device. In the absence of such commands, a reset of the entire SCSI bus might be required.

Hardware advantages of the Hitachi Freedom 5700E

The Freedom 5700E's RAID architecture comprises two controllers sitting at either end of a string of SCSI disks. Each controller has its own cache and is assigned ownership of a set of Logical Units (LUNs). A special bus duplicates the write data across the controller caches and links the two controllers. If one controller path fails, the other controller can take over ownership of the other controller's LUNs without any loss of data, by issuing a read or write to those LUNs. This causes the transfer of ownership of the LUNs. This transfer takes only seconds.




Figure 1: The Hitachi Freedom 5700E configured for high availability through alternate pathing

If host clustering is also involved—as shown in Figure 1—care must be taken to avoid the “ping-pong” effect. The “ping-pong” effect can occur when SCSI A on Host 1 fails and causes CNTL 2 to take over all the LUNs that had previously belonged to CNTL 1. In the meantime, Host 2 still has access to CNTL 1. Accessing CNTL 1 causes the LUNs to be transferred back from CNTL 2. The next time Host 1 accesses CNTL 2, the disks on CNTL 1 are transferred back to CNTL 2, causing a “ping-pong” effect. The effect can occur on any dual-controller disk subsystem. In order to avoid this condition in a clustered host configuration, each controller should be allocated with its own set of disks to separate hosts.

Hardware advantages of the Hitachi Freedom 7700E

The Freedom 7700E has a “share-everything” architecture which uses a common global cache for all communications between channel host interface processors and disks. Since they all use a common cache, the disks can be shared between multiple host interfaces. The 7700E is capable of supporting all modes of alternate pathing with the appropriate host software.

Because the 7700E has a “share-everything” architecture, it does not suffer the “ping-pong” effect of dual-controller subsystems. Path identifications, or “signatures,” are written on the disks to ensure that the correct path is used. This path assignment is done through the Service Processor by a service representative or by the customer through the optional LUN Manager software. The 7700E can be ordered with two, four, six, or eight Client-Host Interface Processors, which can provide 8 or 32 concurrently active SCSI-2 Fast/Wide Differential-Ended 20MB/sec paths to any disk.

Figure 2 shows a multiplatform 7700E with eight SCSI paths shared between three SCSI hosts, and eight ESCON* channels attached to a number of S/390 hosts. The first SCSI host has four paths, the other two SCSI hosts have two paths each. The SCSI paths are split between CHS pairs to ensure redundancy at the CHS level.

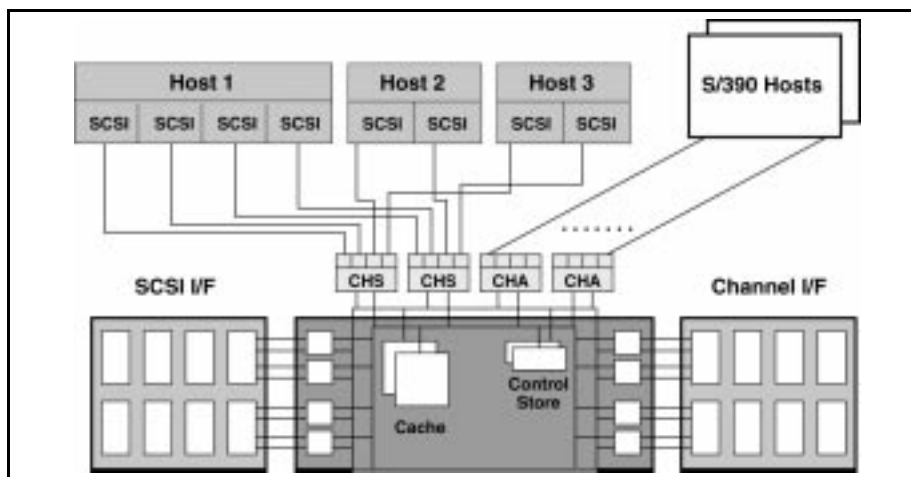


Figure 2: The Hitachi Freedom 7700 configured for high availability through alternate pathing

Alternate pathing for Sun Solaris

Hitachi SafePath

Sun Solaris 2.5 does not include proprietary alternate pathing. Solstice Disk Suite provides two paths to data through software disk mirroring, but does not provide multipath access to an external disk array. The Freedom 5700E/7700E provide alternate pathing for Solaris through the use of Hitachi SafePath middleware.

SafePath allows system administrators to add redundant SCSI connections between Solaris servers and storage subsystems, ensuring reliability and improved performance.

The initial controller path must be assigned manually. The unit of allocation is a SCSI path requiring two SCSI host adapters and two SCSI cables. SafePath provides an ACTIVE/STANDBY mode of alternate pathing. Only one path is active at a time.

When a data path fails, SafePath automatically routes all I/O to the alternate SCSI path without any loss of data access. Once the failure has been repaired, a SafePath command must be issued to return to the initial state.

SafePath is not an “off-the-shelf” product. Hitachi has licensed and optimized this middleware for Hitachi Freedom Storage subsystems as a competitive differentiator. As shown in Table 1, Hitachi SafePath is comparable to EMC PowerPath, and superior to DG Clairion ATF and Symbios RDAC.

Feature/ Function	Hitachi SafePath	EMC PowerPath	Clairion ATF	Symbios RDAC
Device driver modification delays migration to new OS levels	No	No	Yes	No
Device and path failure detection	Yes	Yes	Device	Device
Real-time performance monitoring	Yes	Yes	No	No
Alternate path with cluster support	Yes	Yes	Yes	No*

*Symbios RDAC does not provide alternate path support in a cluster environment as it provides only one path per host in a dual-host cluster. It can only provide alternate pathing in a single host configuration.

Table 1: Comparative analysis of Hitachi SafePath and other alternate pathing middleware solutions

There are two separate price models for Hitachi SafePath, one for use with Freedom 5700Es in distributed environments, and one for use with Freedom 6700/7700Es in centralized, heterogeneous platform configurations. Another version of SafePath will be available for NT in the future.

VERITAS volume manager

VERITAS, which provides logical volume management middleware for Sun Solaris, has recently added a Dynamic Management Path (DMP) capability to its Volume Manager 2.5.1. DMP provides alternate pathing capability for Solaris 2.6. The Freedom 5700E supports DMP. However, the Freedom 6700/7700E does not support DMP at this time.

Alternate pathing for HP-UX

HP9000 HP-UX (10.01 or later) provides PVLlink as a standard facility with its Logical Volume Manager (LVM). PVLlink detects a path failure and switches data over to an alternate path automatically. Unlike Hitachi SafePath, the unit of allocation is a SCSI device (LUN). PVLlink will work with one daisy-chained SCSI connection; however, for best protection against a path failure, the devices should be on separate SCSI adapters and cables.

Once the failure has been repaired, PVLlink automatically returns the data path to its initial state. Tests indicate that PVLlink polls the primary path first on every I/O, even when it is not available. This reduces I/O performance until the primary path is returned to operation, though the data is still accessible. The assignment of a primary path to a controller port requires a manual operation. PVLlink provides ACTIVE/STANDBY support of alternate pathing. The Freedom 5700E/7700/ 7700E support PVLlink to provide alternate pathing for HP-UX.

Alternate pathing for IBM AIX™

Alternate pathing for IBM AIX is provided as a standard facility with the AIX 4.2.X Logical Volume Manager. When the system is initialized, the 5700E and 7700E identify themselves as IBM 7135 devices; therefore, when these subsystems are connected to an RS/6000 AIX system in emulation mode, they are recognized and displayed by the RS/6000 as devices on an IBM 7135 disk array. (This shouldn't be a concern to AIX users, who work with logical representations called "hdisks" rather than device types.)

This support automatically assigns a primary path. It also cuts off an error path and accesses the LUN from the other path automatically.

Alternate pathing for Windows NT

HDS plans to support alternate pathing for Windows NT with a version of Hitachi SafePath. Implementation will be similar to the SafePath support for Sun Solaris. SafePath for Windows NT will also require at least two SCSI host adapters and two SCSI cables.

Host Failover

The increasing need for host-failover middleware

While most customers implement multiple path connections to their I/O sub-systems and use some form of alternate pathing middleware, the use of host-failover middleware has not been as widely accepted. The implementation of host-failover middleware requires detailed knowledge of the systems and application-recovery scenarios which are automated through scripts. Additionally, the total investment in host-failover middleware can run into the hundreds of thousands of dollars.

Until recently, host-failover middleware was most critical to customers with a large number of OLTP users. The limited number of Very Large Database (VLDB) users usually found manual takeover more than adequate. However, this profile is changing rapidly. Now, due to the widespread use of Web-enabled access, VLDBs can support large numbers of on-line users, significantly increasing the need for host-failover middleware.

Important hardware considerations for host failover

Host-failover middleware depends on some form of "coupling" between all of the hosts in a cluster; it is through this coupling that the hosts can monitor each others' "heartbeats." When a host's heartbeat is not heard, one of the other hosts in the cluster executes a series of scripts to regain control of the peripherals and networks attached to the failed host. It also attempts to recover and restart the failed host's applications.

Successful host failover requires that all hosts be physically connected to all devices and networks in the configuration. The term "loosely coupled" is used to differentiate this type of configuration from "tightly coupled" processor configurations—like symmetric multiprocessor (SMP) or non-uniform memory access (NUMA)—which share the same memory and operating system image.

Most host-failover clusters will support up to eight hosts. However, until recently, clusters have been smaller due to host-interface limitations on RAID subsystems.

In fact, before the Freedom 7700E, most RAID subsystems included only two controllers with two connections per controller. Therefore, RAID subsystems could not support more than two hosts with alternate paths to each host. With its “share everything” architecture, the 7700E provides up to 32 paths. As many as 16 hosts can be connected to the subsystem, each with a primary and an alternate path. This greatly reduces the cost of host clustering—instead of requiring a one-to-one host cluster, the 7700E allows a single host to back up as many as 15 other hosts.

Host failover for Sun Solaris

In Sun Solaris environments, VERITAS FirstWatch monitors hardware and software to provide a comprehensive high-availability solution. When a fault occurs, recovery is fully automated, and application services are restored within a matter of seconds or minutes. System configurations can support up to five nodes or host systems (five active and one standby). Each node sends a heartbeat pulse over two exclusive LANs every 0.5 seconds. When both heartbeat links fail, the standby host confirms the failure through the service network before executing takeover scripts. Typical FirstWatch configurations also incorporate Hitachi SafePath for alternate pathing. FirstWatch 2.2.3 is currently the primary high-availability software for Sun systems. It is not available for AIX or HP-UX systems. The Freedom 5700E supports FirstWatch to provide host failover for Sun Solaris; the Freedom 7700E will support FirstWatch in 2Q98.

Host failover for AIX

IBM’s High Availability Cluster Multi-Processing (HACMP) for AIX enables clustered RISC System/6000 servers to recover from server, disk, network, and network-interface failure. HACMP supports up to eight host nodes in a number of configuration modes. These include “idle standby,” where one idle host supports up to seven active hosts; “rotating standby,” where the idle processor role is rotated among up to eight processors; “mutual takeover,” where all hosts share the applications and the workload is taken over by the remaining hosts; and “concurrent access,” where all the hosts not only share the application but also share the data. The heartbeat can be monitored in a number of ways, including via LAN, RS232, or SCSI target-mode through the shared disk attachments.

Developed by CLAM Associates, HACMP is now built into the standard AIX 4.2 offering. CLAM has tested and verified Hitachi Freedom Storage’s support of HACMP. Alternate pathing can be implemented in conjunction with HACMP and is supported by Hitachi Freedom Storage through AIX LVM.

A typical scenario for HACMP recovery might be as follows: Two nodes provide database services to multiple clients. If one of these nodes fails, the HACMP facilities allow the surviving node to take over for the failed node. The failover node assumes control of all shared resources and starts a second database server identical to the one that was running on the failed node. At this point, client applications check for cluster status and reconnect to the new database. HACMP configurations with more than two nodes are best supported by the Freedom 7700E.

Most large HACMP configurations today include IBM 7133 SSA disk subsystems. But because the IBM 7133 SSA does not support caching or hardware RAID in multi-node (multi-initiator) clusters, ensuring high availability becomes extremely complex. The 7133 SSA relies on software mirroring, invoked through AIX LVM, to provide some

level of high availability.[†] Data is “mirrored”—simultaneously copied—to a paired disk. If the primary disk fails, the paired disk can provide continuous data access. However, the data is no longer protected until the mirror is re-established. Re-establishing the mirror (or “resilvering”), is a difficult, time-consuming process in any environment. In a large database environment which has suffered a rolling outage, recovery is considerably more lengthy and complex. Compared to this type of software-based RAID-1 solution, Hitachi Freedom Storage, with RAID-5 hardware and global spare, offers a clear advantage because the data remains RAID-protected if a disk fails. Recovery in the event of disk failure is much more simple, fast, and reliable.

[†] Disk mirroring is also employed by Sun to provide high availability on its Solstice Disk Suite and SpareArray.

Host failover for HP-UX 10.X

For use with HP-UX 10.X, MC/ServiceGuard (Multi-Computer ServiceGuard) protects mission-critical applications from a wide variety of hardware and software failures. MC/ServiceGuard organizes systems with between two and eight nodes into an enterprise cluster that delivers highly-available application services to LAN-attached clients. MC/ServiceGuard also monitors the health of each node and quickly responds to failures, eliminating or minimizing application downtime. MC/ServiceGuard is able to detect and respond to failures in the system processing unit (SPU), system memory, LAN media, LAN adapters, system processes, and application processes.

MC/ServiceGuard ensures that any application package will run on only one node at the same time. The cluster automatically reconfigures itself when a node fails. MC/ServiceGuard offers a safeguard in the form of a quorum, which prevents multiple clusters from forming in the event that some of the nodes cannot talk to one another (multiple nodes could form multiple sub-clusters or groupings). A cluster cannot form unless more than 50 percent of the configured nodes are members of the cluster. This is controlled through a cluster disk lock feature—a semaphore—which is used to break ties. Ordinarily, when the cluster re-forms as a result of a failure, any systems not included in the reconfigured cluster are immediately crashed to prevent the possibility of multiple clusters forming. The cluster lock disk must be accessible from all nodes in the cluster.

The 5700E supports MC/ServiceGuard to provide host failover for HP-UX (this support includes use of the 5700E as a cluster lock disk). Hitachi is testing 6700/7700E support. MC/ServiceGuard is used in conjunction with PVLINK, available as a standard feature of HP-UX 10.01 and later.

Host failover for Sun Ultra Enterprise Cluster HA 1.3

Sun has introduced Ultra Enterprise Cluster HA 1.3 for use with the Ultra Enterprise Server family. This feature provides clustering for any two of the Ultra Enterprise 2, 3000, 4000, 5000, or 6000 servers. Software requirements include Solaris 2.5.1, Solstice DiskSuite 4.0, and Solstice High Availability 1.2. Solstice HA is built on Solstice DiskSuite, which provides software mirroring, concatenation, stripes, hot spares, and UFS logging. Every Ultra Enterprise Cluster is sold packaged with installation, training, and consulting services from SunService. Solstice Disk Suite depends on software mirroring for disk redundancy. VERITAS FirstWatch, used with Hitachi SafePath, provides better flexibility and scalability with simplified management and lower costs. (See the competitive discussion on disk mirroring above, in the HACMP section.)

Solstice DiskSuite is co-packaged with Solaris operating system, application, and enterprise servers solutions. It is also sold separately for the stand-alone desktop environment. Each Solstice HA server acts as an I/O master for its respective diskset, and runs data services that export data on that diskset. In a symmetric configuration, each server is also a backup for the sibling server's data services. Solstice HA provides programs used by each server to monitor the status of data services, as well as the data services running on the sibling machine in the configuration.

Solstice HA automates the decision to take over when the sibling server has a software or hardware failure. Takeover processing includes assuming I/O mastery of the failed server's diskset, and redirecting the failed server's clients to itself. Takeover also includes actions specific to the data service. Additionally, Solstice HA provides for administrative initiated switchover—the graceful switch of a disk-set from one functional server to the sibling to reconfigure or bring a server back on-line.

Windows NT Cluster Server (Wolfpack)

Microsoft supports a multihost configuration for Windows NT with the Wolfpack product. Currently it supports up to two hosts. Wolfpack uses the SCSI Reserve and SCSI Bus Device Reset commands to manage the transfer of disks between systems. Hitachi SafePath can provide alternate path support for Windows NT. The Freedom 5700E has passed certification for Wolfpack in a single controller environment. The Freedom 6700/7700E are in process of being certified for 2Q98.

Parallel Database Clustering

Parallel database clustering middleware offers special failover support for parallel database servers. Parallel database servers provide much larger databases—such as data warehouses—and offer advantages in management, scalability, and performance for read-intensive applications. Parallel database clustering support utilizes a distributed lock manager which allows all cluster hosts to access the same data. The lock manager provides data integrity by coordinating locks associated with logical blocks of data among the cluster systems. All systems have read/write access to the data, but can write a given data block only if it holds the lock for that block. For a system to write a block of data when it does not have the lock, the system must request the lock from the other system over a network connection.

MC/LockManager for Oracle Parallel Server

HP provides a special clustering middleware for use with the Oracle Parallel Server (OPS). Unlike MC/ServiceGuard, MC/LockManager allows sharing of data. An Oracle Parallel Server instance is active on each system within the cluster. The database (and the physical disks on which it resides), may be concurrently accessed in read/write mode by all systems in the cluster that are running OPS. Failover time can be greatly reduced since OPS is already running in all the hosts and database recovery only needs to be done on the portion of the database accessed by the failed host. MC/Lock Manager recommends the use of PVLink to provide redundant attachments between the OPS hosts and the disk subsystem. Hitachi Freedom 5700E/6700/7700E testing of the MC/LockManager and Oracle Parallel Server Version 7 is under consideration. Oracle Parallel Server Version 8 support will be dependent on future MC/Lock Manager support.

Sun Ultra Enterprise Cluster PDB (Parallel Database) Server

Sun offers Ultra Enterprise PDB for customers who employ Oracle Parallel Server (OPS) Version 7. PDB provides the lock management which supports the sharing of data across multiple systems. PDB is currently only supported by the Sun SparcArray storage subsystems. Since Oracle Version 8 uses its own Distributed Lock Manager (DLM), Sun's PDB must be modified to support Version 8.

IBM AIX support of parallel databases

Oracle OPS can be supported in any of three ways by IBM. In Symmetric Multiprocessing (SMP) OPS is supported through shared memory; in HACMP through shared disks in clusters of up to four processors; and in the SPx family through the high-performance switch and the Virtual Shared Disk software that was written to support Oracle OPS.

HACMP provides a Concurrent Resource Manager which enables up to eight-way concurrent access to shared disks in a highly available cluster. This works well with Oracle Parallel Server, which has a share-everything approach to data access (each node in the cluster has access to all the data in the database). HACMP's partitioned data access works well in conjunction with parallel database products like IBM's DB2 Parallel Edition, Informix XPS, and Sybase MPP, which use a share-nothing approach to data access (each node owns a part of the database and the work or query is sent to the node that owns the data).

The Freedom 5700/6700/7700 have all been tested with HACMP. Tests with Oracle OPS are expected to be completed in 2Q98.

Summary

Alternate pathing and host-failover support are becoming key decision-making criteria in the acquisition of high-end storage subsystems. Recognizing the importance of high availability to its customers, Hitachi Data Systems is integrating support for this middleware as it develops its own open systems storage products. The combinations and permutations of base platforms, I/O architectures, alternate-pathing middleware, and host-failover middleware creates a challenging development and test effort.

Table 2 lists all the different platforms and middleware supported on the Hitachi Freedom Storage subsystems.

Platform	Alternate Pathing	Host Cluster
IBM AIX	AIX LVM	HACMP
HP-UX	HP-UX LVM	MC/ServiceGuard
Sun Solaris	Hitachi SafePath	VERITAS FirstWatch
Windows NT	Hitachi SafePath	MSCS (Wolfpack)
Digital UNIX	N/A	TruCluster
Sequent DYNIX/ptx	Hitachi Driver	ATAP

Table 2: High-availability middleware supported by Hitachi Freedom Storage

Hitachi Data Systems

www.hds.com

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia-Pacific Headquarters

11-17 Khartoum Road
North Ryde NSW 2113
Australia
02-9325 3300
info@hds.com.au

Canada Headquarters

380 Saint-Antoine Street West
Suite 7000
Montreal, Quebec H2Y 3X7
Canada
(514) 982-0707
info@hdsCanada.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
01753-61-8000
info@hds.co.uk

Latin America Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
lad@hds.com

U.S. Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1066
ussales@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd. Freedom Storage is a trademark of Hitachi Data Systems Corporation.

*ESCON is a trademark of International Business Machines Corporation. AIX is a trademark of IBM Corporation and is being used under license.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

All other brand or product names are or may be trademarks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by HDS. This document describes some capabilities that are conditioned on a maintenance contract with HDS being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local HDS sales office for information on feature and product availability.