

SharePlex/iX-NetBase

Description

SharePlex/iX-NetBase is a multi-system local or wide-area clustering approach for MPE/iX HP e3000 systems. This product is based on a partnership between Hewlett-Packard and QUEST Software, Inc. of Newport Beach, California.

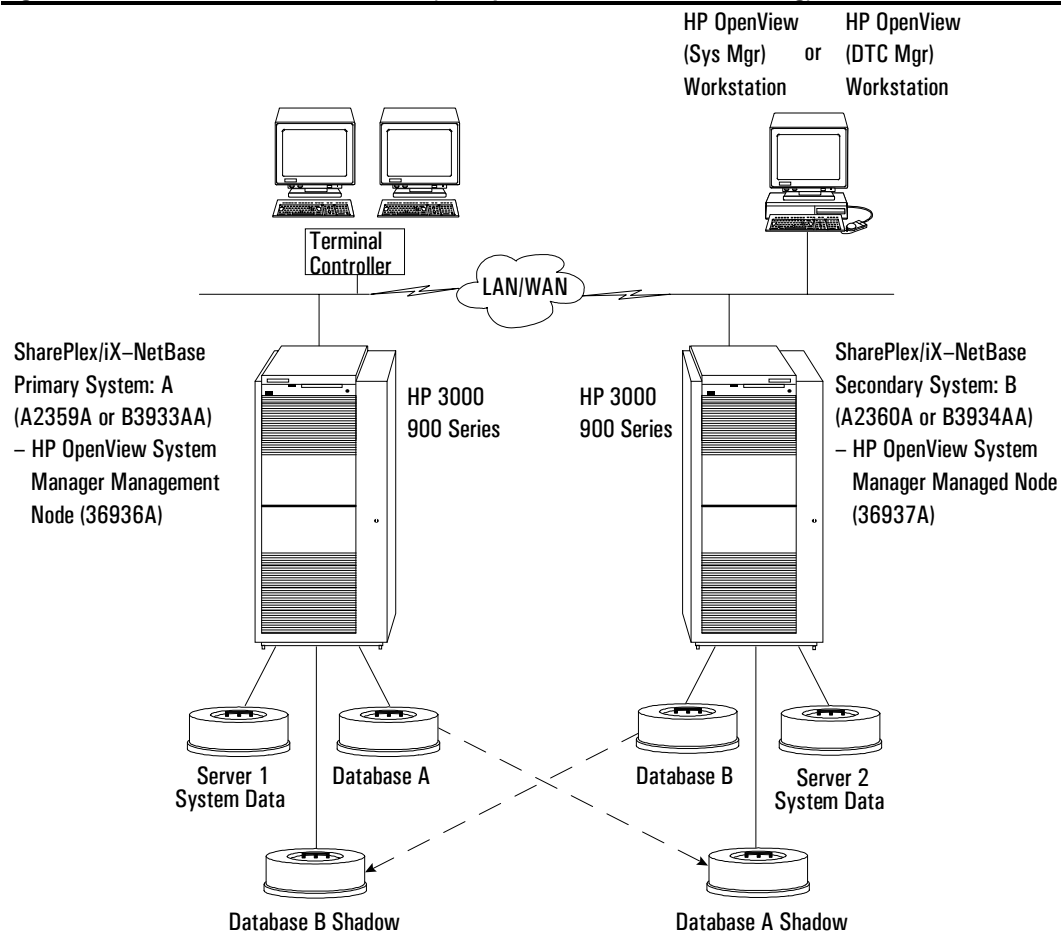
SharePlex/iX-NetBase has the following features:

- Data shadowing (ALLBASE/SQL and IMAGE/SQL databases, flat files, KSAM, etc.) for horizontal application growth, and remote site disaster recovery.
- Network scalability to 124 nodes.
- Shared/clustered applications across the network featuring a central file system for distributed systems.
- Shared remote resources for peripherals such as print queues and batch submissions.

HP OpenView products are available to supplement SharePlex/iX-NetBase to provide the systems and network management necessary in a multi-system distributed or centralized environment. The HP OpenView products are not required or bundled with the SharePlex/iX-NetBase software.

In **Figure 6.1**, you can see an example of the SharePlex/iX-NetBase solution using both the HP OpenView System Manager and DTC Manager to control the cluster. The HP OpenView System Manager can control up to 130 HP e3000s or Managed Nodes on the network.

Figure 6.1 SharePlex/iX-NetBase Solution (example of bi-directional shadowing)



Product Structure

SharePlex/iX-NetBase is system software or “middleware” that resides between the application/user and the MPE operating system. Each HP e3000 within the “cluster” or group of networked systems must have the SharePlex/iX-NetBase software installed.

SharePlex/iX-NetBase utilizes a right-to-copy product structure. All clusters will have at least one primary license configured, and one or more secondary licenses. The two SharePlex/iX-NetBase product groupings are shown in **Table 6.1**).

Table 6.1 SharePlex/iX-NetBase Product Groupings

Product Name	Primary Product Number	Secondary ¹ Product Number	Functionality
SharePlex/iX-NetBase Bundle	A2359A	A2360A	<ul style="list-style-type: none"> • IMAGE/SQL, flat file, KSAM shadowing • ALLBASE/SQL Transport ² • NFA • NBSpool • AutoRPM • Statistics
SharePlex/iX-NetBase Cluster	B6234AA	B6235AA	<ul style="list-style-type: none"> • NFA • AutoRPM • Shadowing
SharePlex/iX-NetBase Shadowing	B3933AA	B3934AA	<ul style="list-style-type: none"> • IMAGE/SQL, flat file, KSAM shadowing

¹“Secondary product” is the right to copy version. Each logical grouping or cluster of HP e3000s must have one primary system (configured on the largest system in the group) and one or more secondary systems.

²ALLBASE/SQL Transport is included in the bundle. This Transport must be paired with ALLBASE/Replicate (B2494A) in order to achieve ALLBASE Shadowing. ALLBASE/Replicate must be purchased separately. ALLBASE Shadowing can also be purchased from QUEST Software, Inc.

Key Components in SharePlex/iX-NetBase Solution

- Shadowing, which allows local or wide area network replication of IMAGE/SQL, ALLBASE/SQL, KSAM and flat files.
- Network File Access (NFA), a central file system to provide file and database access transparently to users and applications.
- NBSpool Plus—a master print manager and router to give users and applications seamless access to any printer in the cluster. Also serves as software to support HP network printers.
- AutoRPM, which gives end users transparent access and execution capabilities to any program or application in the cluster.
- Statistics, a performance tool to collect and report on file access, access response time, network overhead, etc.

Configuration Requirements

Networking

SharePlex/iX-NetBase is compatible with all HP e3000 supported networking. ThinLAN Link is required (bundled with ARPA on MPE/iX 5.0 and later) on all clustered systems. NS/3000 Services are only required when utilizing the AutoRPM feature within the SharePlex/iX-NetBase bundle (see **Chapter 8** for networking).

Licensing

A “primary” license must be purchased for the largest system in the cluster. This “primary” license does not mean it can only house master copies of data. Any system in the cluster can house master or secondary copies of shadowed data. Any data structure can also be shadowed to multiple systems.

Disk Drives

Unlike Disk Mirroring, SharePlex/iX-NetBase does not have any disk drive dependencies. Shadowing can be performed at the file, group, or account level. The master as well as the secondary data can reside on any HP-supported disk drive.

Operating System Requirements

- Release 9.7 requires MPE/iX 5.0 or later
- Release 9.7 only supports NetBase Global Installs. Account installs are no longer supported with the 9.7 Release.
- MPE/V systems are not supported with the 9.7 release. The last supported release of NetBase for MPE/V is 9.6.

Performance

SharePlex/iX-NetBase utilizes caching algorithms and pre-fetching capabilities to greatly improve network access performance. Measurements have shown an approximate 8–10 time performance improvement over NS/3000 or other network file access techniques.

Overhead and Sizing

Shadowing does add some system overhead. Typically, the master system shows a 5%–10% CPU overhead, and the secondary overhead is directly related to the number of file or database updates.

Networking overhead also varies. Experience shows that local 802.3 networking is rarely a bottleneck. FDDI is recommended, however, for applications that have very high transaction rates.

Wide-area networking is a greater challenge. The proper sized “networking pipe” is vital for quick data access and shadowing performance. Unfortunately, there are not any standard formulas to calculate this capacity. Contact an HP networking specialist, or QUEST Software, Inc., (714) 720-1434, to help determine networking performance issues.

Implementation Assistance

It is always strongly recommended to purchase 3–10 days worth of SharePlex/iX-NetBase implementation assistance, depending on the complexity of the system environment. This service can be purchased from a limited number of people in the HP PSO, or through QUEST Software, Inc. directly. HP and QUEST have had very good success at partnering on the implementations.

Quarterly customer training courses are offered by QUEST in Newport Beach, California. Call QUEST Software, Inc. directly or go through your HP sales representative to book the one-week class.

High Availability Cluster/iX

Overview

High Availability Cluster/iX is a high availability solution that protects customers against loss of server access to data. The Cluster consists of a minimum of two HP e3000s connected to a multi-ported array. In this topology, all HP e3000s connected to the array have access to disk drives on the other port. However, only one HP e3000 has direct access to the MPE/iX sharable volume set. All computers in the topology agree on a single owner and all access to the shared data is through the computer that owns the volume set.

Sharable disks are called Cluster volume sets. Ownership of a cluster volume set is determined by an “ownership label” written to the volume set at its creation. Ownership is assigned only to the master volume of the volume set and not the members of the volume set. The ownership mechanism prevents multiple systems from trying to access the same volume set simultaneously. High Availability Cluster/iX is not limited to 2-way topologies. It can manage N-way systems. The only limitation is the ability to configure and access multiple independent pathways from each of the computers to the cluster volume.

Usage

The intended usage of High Availability Cluster/iX is to reduce planned or unplanned downtime due to system failures, routine maintenance, power outages or other unforeseeable events. High Availability Cluster/iX provides increased up time by moving the mission critical data from the unavailable system to a secondary system that is up and running. It in effect becomes the other system.

Cluster Control

Cluster change over between servers is accomplished via a command issued manually. Commanding an HP e3000 to take control of another computers volume set causes cluster utilities to run and take control of disks that have been designated as “Cluster volume sets” over from another computer. Once another HP e3000 has taken control of the cluster volume, it becomes the primary system and does not need to be switched back. The previous primary then becomes the secondary system.

Failure Detection

There are no internal automated controls to detect and exchange an unavailable server with one that is available. Operators still need to monitor system availability through OpenView ITO manager or other system management tools. Automation is planned for the next release.

Supported Disk

The disk technologies supported in High Availability Cluster/iX topologies are any HP e3000 supported disk arrays:

- HP Surestore Disk Array XP48, XP256, XP512
- HP Surestore Disk Array 12H
- HP Surestore Disk Array Model 20

Operating System

- MPE/iX 6.5 Express 2 and later

Ordering Information

High Availability Cluster/iX is ordered by using part number B9480AA to order the primary system software and part number B9842AA to order the right to copy license for the secondary or HOT Standby system. Customers ordering a primary license must have at least one secondary license on the same order or the order will not be processed. Customers may separately order any number of secondary licenses as long as they have at least one primary license. The right to use High Availability Cluster/iX on an HP e3000 computer system is subject to the terms and conditions of the Hewlett-Packard Software Purchase Agreement.

Product Structure

Product Number	Option Number	Description
B9480AA		High Availability Cluster/iX Primary
B9480AA	OCD	Upgrade credit for option
B9480AA	OCF	Upgrade credit for option
B9480AA	310	System license for HP 3000 Tier 1
B9480AA	330	System license for HP 3000 Tier 4
B9480AA	340	System license for HP 3000 Tier 6
B9842AA		High Availability Cluster/iX Secondary
B9842AA	OCD	Upgrade credit for option
B9842AA	OCF	Upgrade credit for option
B9842AA	310	System license for HP 3000 Tier 1
B9842AA	330	System license for HP 3000 Tier 4
B9842AA	340	System license for HP 3000 Tier 6

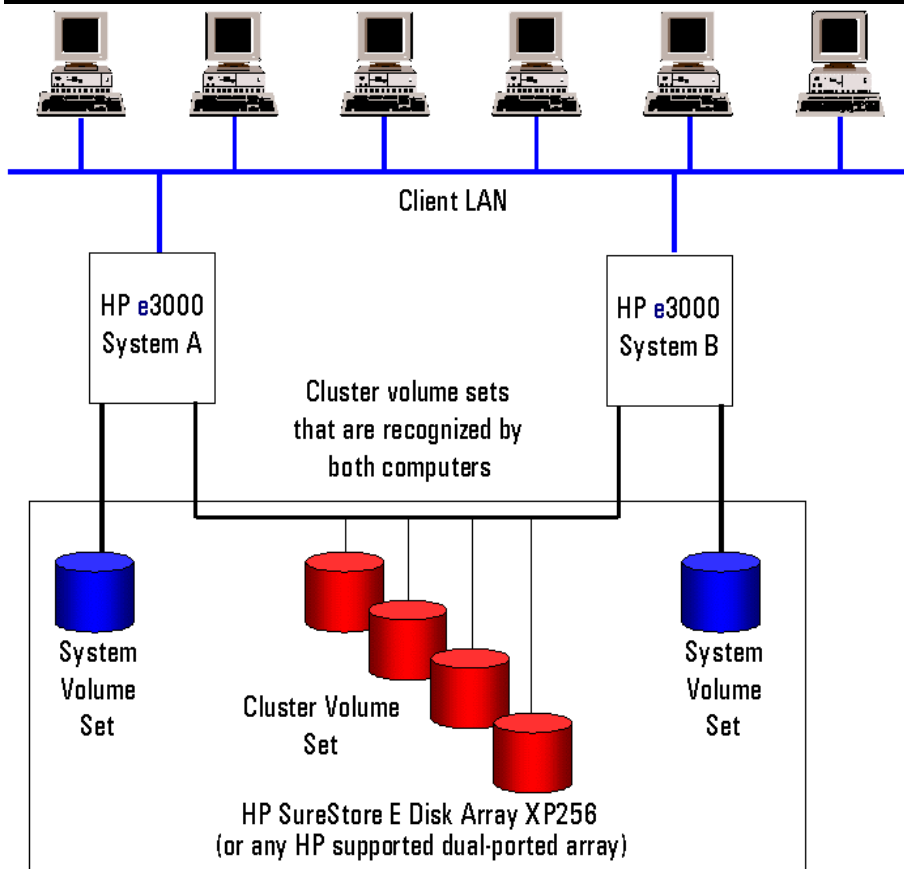
Documentation

The operation and functionality of High Availability Cluster/iX is documented in the High Availability Cluster/iX User and Configuration manual, which is available from Direct Marketing Division part numbers (B9480-90001).

HP Worldwide Response Center Support

HP Response Center engineers have been fully trained on High Availability Cluster/iX operations.

Figure 6.2 High Availability Cluster/iX



High Availability FailOver/iX

Overview

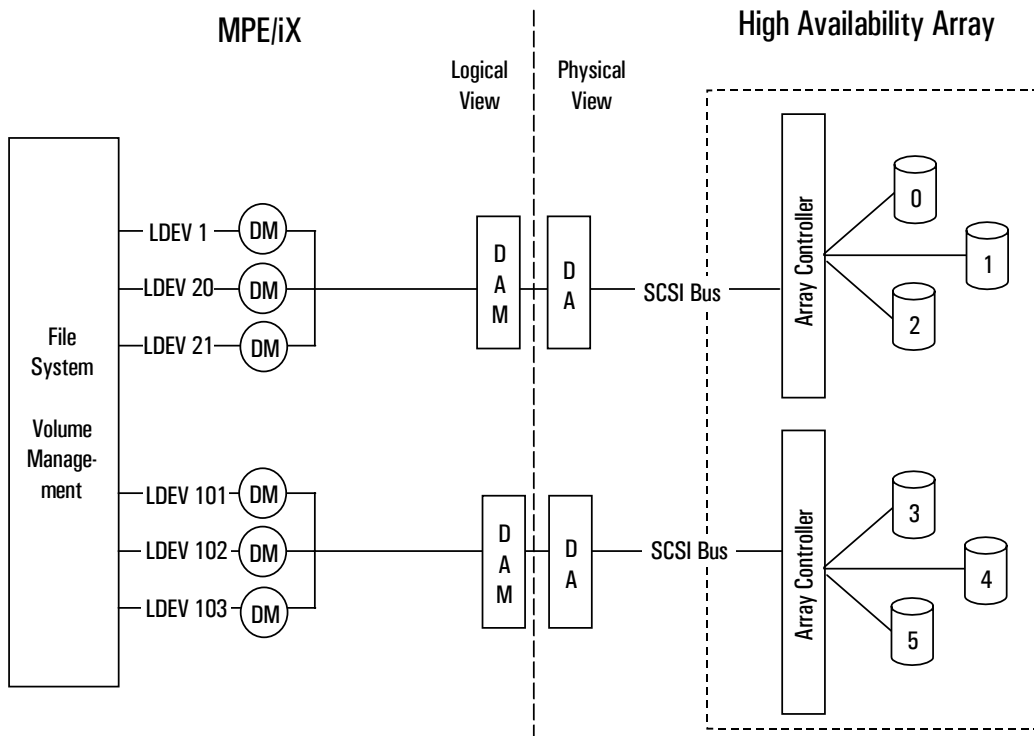
High Availability FailOver/iX is a dual active path failover solution for dual multi-ported arrays that are supported on the HP e3000. At first release, HAFO is only supported on the HP SureStore Disk Array XP256. It will be supported at a later time on the whole HP XP family of disk arrays. HAFO allows MPE/iX to switch from a failed I/O path to another active I/O path that is connected to the same array. HAFO is not a separate product, it is included free with MPE/iX.

Once installed and configured, the High Availability FailOver/iX (HAFO) continually monitor SCSI reply messages for failed SCSI data path components. This adds minimal overhead to the I/O subsystem operation. HAFO event information and data structures are memory resident. This eliminates the need for disk file access to perform high availability failover. This is an advantage especially if a SCSI path failure to LDEV#1 should occur. No failover action is taken until you configure the ldevs for HAFO using the new "ha" section of SYSGEN. Device data path and alternate data path information is entered and saved in SYSGEN's HAFOCONF configuration file. HAFOCONF configurations are read and validated during each system boot.

Specific configuration information is provided in the Chapter, "Configuration." In the HAFO configuration guide

Figure 6.3 illustrates a sample configuration. This figure can be compared against **Figure 6.4**, which illustrates a failover of the same system.

Figure 6.3 Normal System Layout



Triggering a Failover

HAFO only acts on specific error types that indicate a data path failure. Any of these three occurrences will generate a failover:

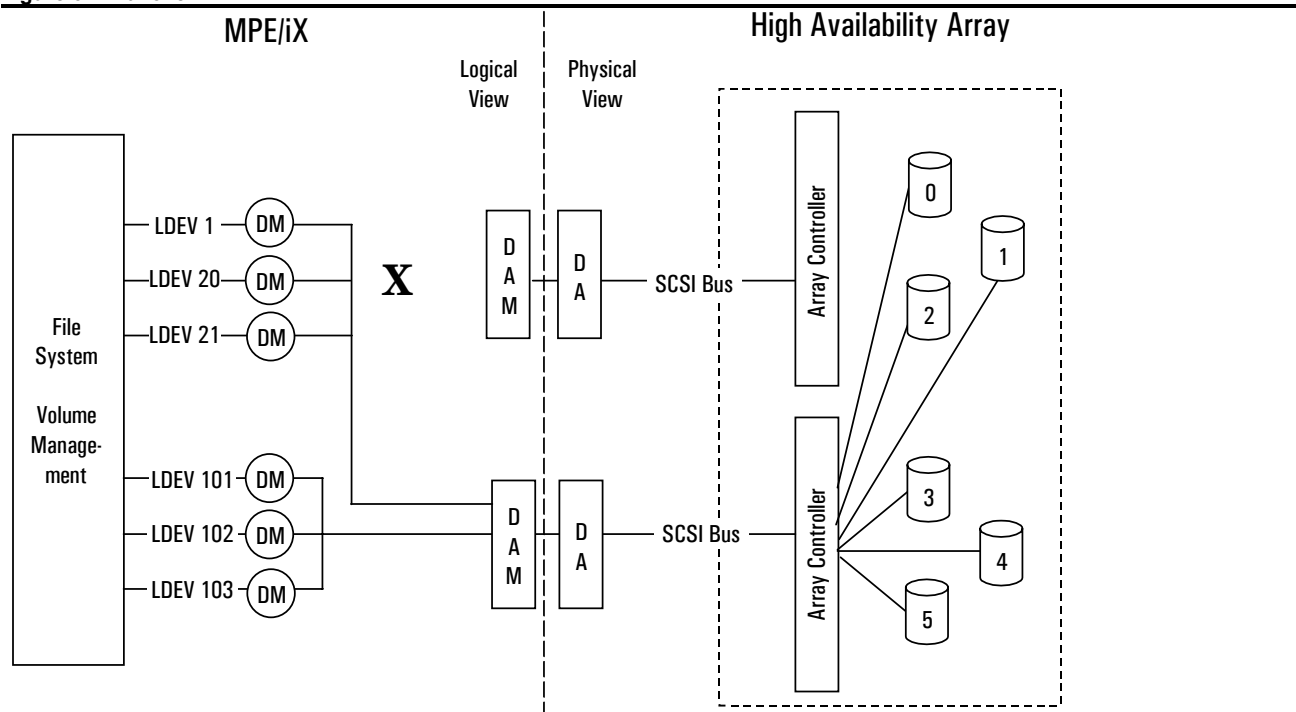
- hung I/O
- failed high availability array controller (communicated by a SCSI reply status)
- failed host device adapter card

If any other error type occurs (such as a data transmission or device error), the I/O subsystem will manage the error and perform corrective action. HAFO will remain idle and not participate. In addition, HAFO remains idle when error types are received from non-high availability devices.

Executing a Failover

When a trigger status is received, HAFO will immediately begin the failover sequence. This sequence activates the alternate data path and reroutes I/O to it. Failover occurs on a per ldev basis. Each device manager (the piece of the I/O subsystem that manages a specific ldev) learns of the data path failure during a subsequent I/O to its ldev. For example, if ten ldevs on a fast-wide SCSI bus experience an array controller failure, the associated ten device managers will perform a HAFO failover event independently.

Figure 6.4 Failover



No application or higher level MPE/iX Operation System (above I/O subsystem—for example, file subsystem, database management, or memory management subsystem) will experience an abnormal event. All I/Os complete as normal using the alternate data path and alternate array controller.

User Notification of Failover

The I/O and system logs will document the failover event. In addition, the following console error message will be displayed upon failover and every five minutes thereafter. The repeating message can be turned off with a [CTRL] [A] reply.

```
HIGH AVAILABILITY FAILOVER IS STARTED FOR LDEV# IN DISK ARRAY. NO DATA LOSS OR
CORRUPTION. SYSTEM OPERATION WILL CONTINUE. PLEASE PLACE SERVICE CALL SOON.
```

```
ACKNOWLEDGE HAFO FAILOVER IN DISK ARRAY FOR LDEV# (Y/N)?
```

Once a Failover Has Occurred

After an HAFO event, system I/O activity will resume via the alternate data path. There are no limits to the kinds of normal I/O that can be processed on the alternate data path. Throughput may be affected since I/O is shared with bus activities from other ldevs configured for that bus.

For additional information, see the Chapter, "Recovering from a Failover." In the HAFO configuration manual.

Components

The HAFO product has multiple components:

- System Boot Failover Initialization Utility
- SYSGEN HAFO commands
- HASTAT HAFO status report
- HAFOCONF configuration file

Each of these components is briefly described below.

System Boot Failover Initialization Utility

During system boot, the HAFO configuration utility reads the HAFOCONF configuration file and arms HAFO. The alternate path information becomes frozen in memory and is available to any configured devices when needed to complete an HAFO event.

High Availability Failover Commands

The following commands are used to configure and maintain the HAFOCONF configuration file.

- addconf (ad)—adds high availability device configuration
- delconf (de)—deletes high availability device configuration
- listconf (li)—lists high availability device configurations
- doha—validates high availability device configurations, and activates the HAFO configuration
- goback (go)—instructs device manage to use the primary data path
- help—get help on HAFO commands
- exit—exit the "ha" section and return to the "io" section of SYSGEN.

These commands, located in SYSGEN, are documented in the Chapter, "Configuration." Goback is documented in more detail in the Chapter, "Recovering from a Failover."

High Availability Failover Status

With the installation of HAFO, HASTAT displays current HAFO data path status. The report lists each configured LDEV, its primary and alternate paths, and their status. This feature is documented in detail in the Chapter, "Monitoring Status."

HAFOCONF Configuration file

Each configured HAFO ldev will have a record entry in the HAFOCONF configuration file. This file will reside in SYSGEN's CONFIG groups. The HAFO information configured or modified using the HAFO commands resides in this file. HASTAT displays the High Availability Failover Status information available from this file.

System Requirements

HAFO is supported *only* with the HP XP256 high availability array.

Mirrored Disk/iX

Product Overview

Mirrored Disk/iX (30349A) virtually eliminates application downtime from disk failure and dramatically increases application uptime by providing full data redundancy through disk mirroring.

Hardware Configuration

Mirrored Disk/iX can be configured on all HP e3000s that run MPE/iX and requires disk drives equipped with either HP-FL or SCSI interfaces. To mirror disks, you must mirror them with drives that have the same model number as the drive being mirrored. This ensures that the capacity and the performance characteristic of the two disks are equal.

SCSI disk mirroring is supported with MPE/iX Release 4.0 and later, and utilizes User Volumes that are built into the MPE/iX Operating System. Because Mirrored Disk/iX uses volume management, users do change how they manage their system, including backups.

A string of SCSI drives must be mirrored by a string of SCSI drives (again the rule that you have to mirror with identical model numbers).

Fast/Wide SCSI mirroring is supported with MPE/iX Release 5.0 and later. You can mix SCSI and HP-FL drives on one system by configuring a string of SCSI drives to mirror a string of SCSI drives, and a string of HP-FL drives to mirror a string of HP-FL. In other words, you can have both SCSI and HP-FL on one system, but the SCSI drives cannot mirror the HP-FL drives and vice versa.

Each string of drives should be allocated its own dedicated SCSI or HP-FL interface card in order to maximize data protection. **Figure 6.5** shows an HP e3000 system configured with Mirrored Disk/iX. Mirrored Disk/iX can provide greater data availability than standard High Availability disk arrays for several reasons. Mirrored Disk/iX, when configured properly, protects against 1) I/O card failure, 2) cable failure, 3) controller failure, 4) disk mechanism failure, and 5) disk power failure. See **Figure 6.6**. High Availability disk arrays only protect against disk mechanism failures and disk power failures (labeled No. 4 in **Figure 6.7**). In 1H98, MPE/iX will support automatic controller failover. At that time, all five points of failure will be protected with disk arrays.

In addition, since the disks in a mirrored partner are identical, read requests go to the disk in the pair that can service the read in the shortest time, thus optimizing performance. Intense I/O per second applications can experience performance gains using Mirrored Disk/iX over disk arrays.

Sites with very large disk configurations and a limited number of slots may not be able to use Mirrored Disk/iX because of the lack of slots. In these cases, customers should definitely choose the High Availability disk arrays that give the customer protection from disk mechanism failures—soon all points of I/O failure.

High Availability disk arrays should be used on the system domain since it cannot be mirrored.

Mirroring the internal drive on a Series 9x8/xx system is not supported because the SCSI interface supporting that drive is also supporting the DDS device.

Mirrored Disk/iX is NOT supported with Disks connected through the SCSI-FC Distancing Router (A5814A). Mirrored Disk/iX is also NOT supported for usage with High Availability Disk Array products such as the AutoRAID 12H, VA family or XP family of Disk Array products.

Figure 6.5

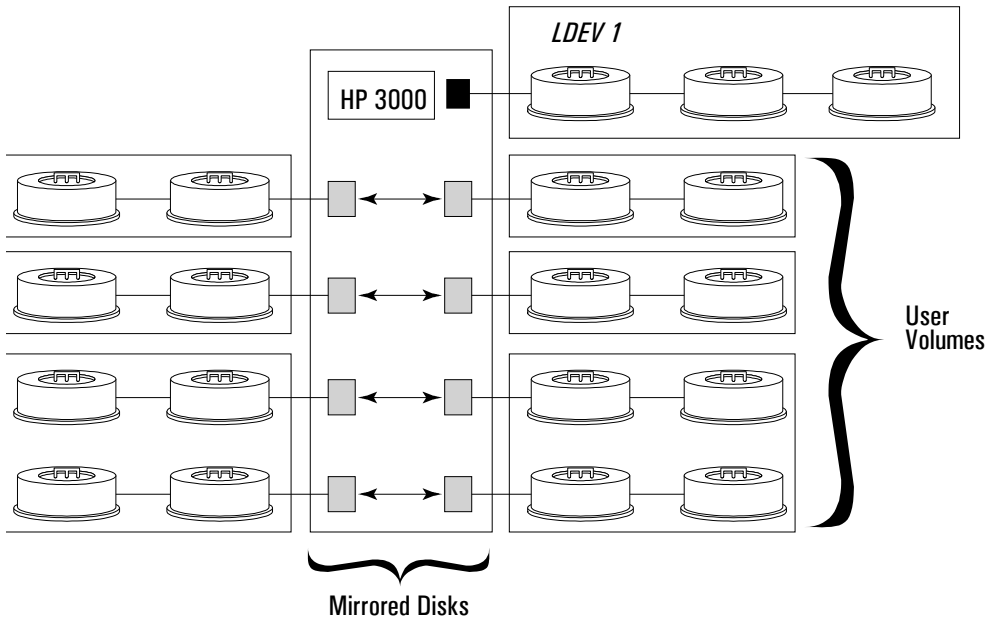
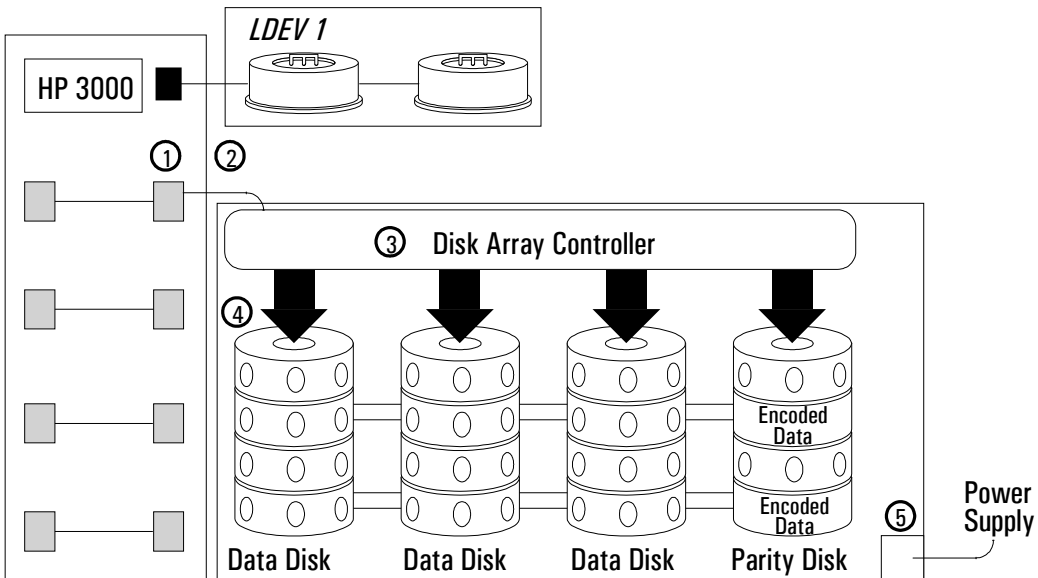


Figure 6.6



AutoRestart/iX

Product Overview

AutoRestart/iX is a software product that performs automatic system restart in the event of a system software failure. System recovery time is minimized and system state information is automatically saved for later analysis. The “dump to disk” feature of AutoRestart/iX dumps key software information to a preallocated disk file, which is substantially faster than the traditional dump-to-tape method. The disk file needs to be large enough to store the largest possible dump that the system can generate. AutoRestart/iX does perform data compression as it performs the automatic save of error information, thus minimizing the amount of disk space required. When a soft failure occurs, AutoRestart/iX automatically saves the system state information to disk (or tape) prior to restarting the system.

Hardware Configuration

AutoRestart/iX is supported on all HP e3000s that run the MPE/iX operating system (Release 2.1 or later).

The question of how much added disk to configure on the system to accommodate this potential need sometimes arises. The following is a guideline for this purpose:

1. Assuming that each user on the system will generate 3 processes, estimate the typical or maximum number of users and jobs and multiply by three. This gives you a number of active processes.
2. If the number of active processes is between 3-200, multiply that number by 1.5-MB disk. If the number is 200 or above, multiply it by 1.0-MB disk.
3. Now add in the amount of main memory.
4. Now add in 96 MB of “fudge factor.”
5. Multiply this total number by .4

This resulting number is the amount of disk space you should allocate to receive the dump containing the system state. This is a general guideline. AutoRestart/iX provides a utility to assist in predicting the disk file size needed.

The mini-dump feature of AutoRestart/iX enables MPE/iX to write a summary of system failure information to a disk file for later analysis, allowing the system to reboot immediately and reduce recovery time.

TurboSTORE/iX II

Table 6.3 STORE and TurboSTORE/iX II Products

Product	Part Number	Description
STORE/iX		Single device store/restore, fastsearch for DDS as of 2.2 and forward
TurboSTORE/iX	30319A	Up to 8 parallel tape drives, up to 8 sequential tape drives, support dissimilar devices (7980XC and 7978), interleaving, fastsearch as of 2.2 and forward—Discontinued
TurboSTORE/iX II	36387A	Software data compression, up to 8 parallel or serial devices, DDS and fastsearch on DDS, parallel restore, IBM/ANSI label tapes, dissimilar tape devices, interleaving—Discontinued
TurboSTORE/iX II with Online Backup	36388A	All functionality in P/N 36387A plus capabilities for online backup—Discontinued
TurboSTORE/iX II with support for Rewritable Optical Disk	36397A	All functionality in P/N 36387A plus support for Rewritable Optical Disk—Discontinued
TurboSTORE/iX II with Online Backup for Rewritable Optical Disk	36398A	All functionality in P/N 36397A plus capabilities for online backup—Discontinued
TurboSTORE/iX II	B5151AA	For MPE/iX 5.5. Same functionality as 36397A
TurboSTORE/iX 7×24 True Online	B5152AA	All the functionality of B5151AA, plus support for 7×24 True Online backup and store to disk

Table 6.4 STORE and TurboSTORE Features
MPE/iX Release 5.0

Feature	STORE/iX	TurboSTORE/iX II			
		36387A	36388A	36397A	36398A
Multiple backup devices		X	X	X	X
File interleaving		X	X	X	X
Data compression		X	X	X	X
Parallel restore		X	X	X	X
Online backup			X		X
Store to disk					
Optical device				X	X
Labeled tapes	X	X	X	X	X

MPE/iX Release 5.5

Feature	STORE/iX	TurboSTORE/iX II	
		B5151AA	TurboSTORE/iX 7×24 True Online B5152AA
Multiple backup devices	X	X	X
File interleaving	X	X	X
Data compression		X	X
Parallel restore		X	X
Online backup			X
Store to disk			X
Optical device		X	X
Labeled tapes	X	X	X

Using TurboSTORE/iX II More Effectively

Use Parallel Devices

TurboSTORE/iX can store data to multiple backup devices running in parallel during backup. This will overcome the performance bottlenecks experienced by backup solutions using only a single backup device. This can significantly reduce backup times.

To enhance performance with TurboSTORE/iX or TurboSTORE/iX II, modified configurations are suggested. For Series 9x8 systems with greater than 12 GB storage, spread disks across 3 SCSI busses. Back-up performance is enhanced with TurboSTORE/iX II, SCSI disks, and DDS by having the disks on a separate SCSI bus from the DDS tapes.

If HP-IB backup devices will only be used sequentially, they may share an HP-IB DA. If devices are used in parallel or parallel in device pools, it is suggested that each backup device have its own HP-IB Device Adapter (DA). Additional performance increases can occur with:

- HP-IB disks spread over at least four HP-IB DAs
- A maximum of two HP-FL DAs per CIO
- Have HP-IB and HP-FL DAs on separate CIO
- A maximum of 3 HP-IB DAs per CIO

Use Sequential Pools

TurboSTORE/iX supports up to eight drives for sequential store. After filling one tape, backup continues automatically by switching to a second tape drive while the previous tape rewinds and unloads. Optimally, you can utilize both sequential and parallel backup features in device pools.

When using large capacity devices, such as DDS DAT tapes or Optical disks, using all available devices in parallel will generally provide greater benefits. Using Sequential devices or sequential device pools is only beneficial if tape rewind and unloading is necessary. Generally, pools of 2 sequential devices are optimal.

Use Data Compression

For products without hardware data compression, such as the 7980S, Optical Disk Libraries, and 1/2-inch cartridge tape without IDRC (data compression), TurboSTORE/iX II boosts backup performance and capacity significantly with software data compression. For maximum flexibility two different data compression algorithms are provided. The low algorithm provides on average 2.1 to 1 data compression and more than doubles backup device performance and capacity while requiring the least CPU resources.

Use the high compress option to reduce the amount of actual data being transferred to the backup media. However, compression does require significant additional CPU time and memory. Therefore, when using TurboSTORE/iX II data compression for unattended online data backup the less CPU intensive low data compression algorithm should be used.

Although HP e3000 systems with typical loads will experience satisfactory on-line backup performance using TurboSTORE/iX II data compression, heavily loaded systems may achieve higher on-line backup performance by utilizing the hardware data compression versions of DDS-DATs and Digital Linear Tape drives.

The 1/2-inch tape compressed on the 7980XC or 7980SX cannot be read by TurboSTORE/iX II if restored from a drive other than the 7980XC or 7980SX.

Online Backup Considerations

Online backup as supported by TurboSTORE/iX II with online backup (36388A) and TurboSTORE/iX with online backup and optical disk support (36398A) are both available on pre-5.5 systems. Instead of the traditional application shut down of 2–6 hours, system unavailability is limited to a small quiescence time to ready the system for backup. Downtime varies between 5 minutes and 20 minutes depending on the number of files, number of processes, and number of open files residing on the system.

Additionally, you should also include application restart time in your system backup and recovery strategy. Application restart time is application dependent and in most cases elapsed time will be short. However, startup time can be significant in some environments such as a distributed computing environment with a large number of networks. In this environment, each network needs to be restarted.

In addition, TurboSTORE/iX 7×24 True Online (B5152AA) is available on MPE/iX Release 5.5 and supports online backup and optical disk. No application modification is necessary to take advantage of either 7×24 True Online or Online backup offered by B5152AA.

The system load during online backup is assumed to be low. On-line backup needs adequate CPU resources to compress data and keep devices streaming.

Use Interleaving

Using the interleave option of TurboSTORE/iX II increases backup performance by optimizing the retrieval of data from disks. The interleave option allows the store process to perform multiple reads on multiple files. This helps to eliminate idle time waiting for data to be retrieved and keeps the backup devices streaming.

TurboSTORE/iX 7×24 True Online backup allows users to remain active within their applications while the backup occurs. Databases and files remain available for access and modification seven days per week, 24 hours per day.

TurboSTORE/iX 7×24 True Online automatically quiesces and synchronizes all TurboIMAGE and ALLBASE/SQL databases to achieve logical database integrity. Applications need not be closed, nor is a separate backup required.

No File Cleanups

To decrease system disk usage and system overhead during an online backup, avoid file cleanup activities (file purge, rename, EOF cutback) that can be done before or after the online backup is complete. If file cleanup is performed during online backup additional disk space is required to keep images of files as they were before the backup started.

Use MAXTAPEBUF Option

The MAXTAPEBUF option can be used to increase the output buffer size from 16 KB to 32 KB. This will improve backup performance.

The 1/2-inch cartridge tape drive from StorageTek can provide your customer with media interchange in mainframe environments. This standard (3480 compatible) cartridge is usable in IBM's 3480 1/2-inch cartridge tape drives. Data stored to this cartridge with IBM label tape format can be easily transferred to an IBM system, and vice versa. For maximum backup performance with the 1/2-inch cartridge tape drive from StorageTek, STORE's MAXTAPEBUF option should be specified.

Avoid TRANSPORT Option

TurboSTORE/iX II can create compatibility mode tapes that can be recovered on MPE V systems and MPE/iX systems using the "TRANSPORT" parameter. You cannot create tapes with TurboSTORE/iX or TurboSTORE/XL II in Native Mode and recover in Compatibility Mode. Compatibility Mode tapes can be created, but no TurboSTORE/XL or TurboSTORE/iX II options (data compression, online, etc.) can be used.

Use the TRANSPORT option of TurboSTORE/iX II only when an alternative format (pre MPE/iX 4.5 or MPE V) is required for transferring data between systems. Since system backups are system specific and will most likely be used only on the source system, the use of the TRANSPORT option is generally not required. Using this option will increase overall backup time.

High Availability Backup/Recovery

Today more businesses require high availability features. Following are specific backup recommendations to meet high availability requirements for these businesses.

User Volume Sets

Improves recovery time. Partitioning the disk subsystem into User Volume Sets makes it easier and faster to recover in the event of a disk media failure. Under this strategy, the operator stores all accounts by volume sets. If a drive fails within the volume set, the operator recovers only the files on the affected user volume set, not the entire system.

Decreases downtime. Not all users are restricted from system access during backup. Only the users of the volume set being backed up are impacted and they are impacted for a shorter time. Users accessing data on other volume sets can still be accessing the system during the backup period.

Figure 6.7 shows a typical high-end system configuration. The system is divided into a system volume set and two-user volume sets. One volume set contains FL arrays with 13.4 GB of data while the second volume set contains three 2.0-GB SCSI disks.

With the configuration shown in **Figure 6.7**, **Table 6.5** compares how long each device would need to backup or recover each volume set or the entire system if no volume sets are used. This scenario clearly indicates the advantages of partitioning your system into volume sets.

Figure 6.7 Typical High-End System Configuration

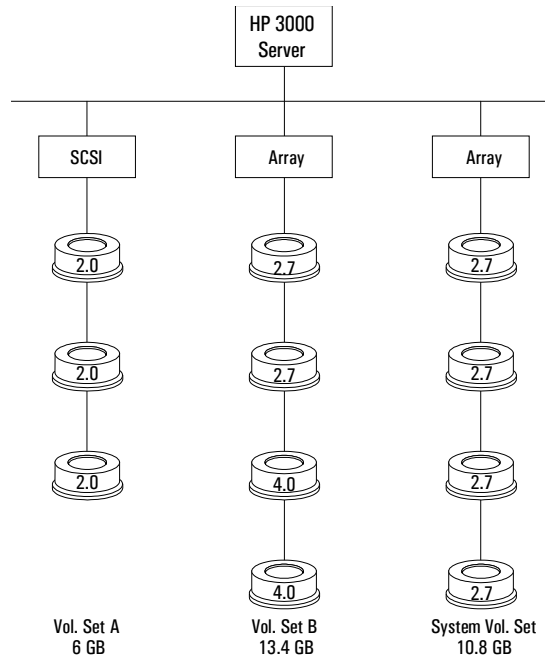


Table 6.5 Backup and Recovery Time Estimates (in hours)

Device	System Volume Set (10.8 GB)		Volume Set A (6 GB)		Volume Set B (13.4 GB)		Total*	
	Backup	Recovery	Backup	Recovery	Backup	Recovery	Backup	Recovery
DDS-2	1.35	1.70	.75	.95	1.7	2.1	3.8	4.7
DDS-3	.5	.65	.35	.4	.56	.70	1.36	1.75
DLT4000	.3	.4	.15	.2	.35	.45	.8	1.0

* If using User Volume Sets, it is extremely unlikely that you would ever need to recover the entire system at one time. That is the primary benefit of user volumes. If a media failure does occur, only users of that volume set experience downtime while the disk is repaired and data restored.

NOTES:

- Assumes 987 processor
- 2.1 : 1 data compression
- Does not include time for media switching or rewind time if applicable

System Backup with TurboSTORE/iX II

To minimize backup time and system downtime we recommend the use of TurboSTORE/iX II. Follow the guidelines starting on **page 6-12** for using TurboSTORE/iX II more effectively. The following features of TurboSTORE/iX II should be used to meet high availability requirements.

Online Backup Capability.

With TurboSTORE/iX II with online and True Online backup capabilities, users can access their data while it is being transferred from disk to the backup media. In addition, choosing a low activity period and a fast backup device or using multiple backup devices will minimize the impact to users during on-line stores. Some overhead is required to perform on-line backup and users may experience performance degradation.

Use of Parallel Devices.

TurboSTORE allows you to backup to more than one device simultaneously. This will obviously increase the throughput of your backup.

Data Compression.

Whenever available use the compression capability of the hardware because device compression does not utilize additional CPU resources. DDS-DC DATs, 7980SX products, and the IDRC version of the 1/2-inch cartridge tapes use hardware data compression. However, if your backup device does not have data compression capability in the hardware, TurboSTORE/iX II can provide software data compression. This will reduce the amount of data that needs to be stored to the backup device, but will increase CPU utilization.

The compression ratio will vary depending on the content of your data, but a typical compression rate for hardware data compression is 3.5:1. When using TurboSTORE/iX II to store data on-line we recommend utilizing the low data compression rate with typical compression ratios of 2:1 to minimize performance degradation.

Flexible Backup Scheduling

You may wish to consider more flexible backup scheduling. For example, backing up different user volumes at different times. Or selectively choosing what data is backed up when.

Speed and Capacity of the Device

The speed and capacity of the backup device is always a consideration. You should determine the time you are willing to dedicate to backup and choose the appropriate device with required speed and capacity. See **Chapter 4** for a discussion of backup devices.