

4010
Building a High Availability Solution
Victoria Symonds Hewlett Packard
19447 Pruneridge Avenue, Cupertino, Ca. 95014
408-447-6042

A Bridge Over Troubled Waters

High Availability is like a “Bridge Over Troubled Waters.” The troubled waters of disk failures, system failures, downtime associated with backup, or any inaccessibility to data presents an impediment to your goal of acquiring the information that you need. Building the appropriate High Availability Bridge, on the other hand, will ensure continuous data availability in any turbulence. Since your business increasingly operates within a multi-system, often heterogeneous environment an appropriate high availability solution needs to also ensure the continuous accessibility of data across these multitude of platforms. You can trust in the HP 3000. We have the solutions you need to lay you down on the other side of these troubled waters.

Up until 5 years ago, recovery of a full system from a backup tape provided acceptable recovery times. All business data resided on a single small system and a full system recovery only took 1 to 2 hours. Things have changed tremendously over the last five years yet we still do not consider recovery in light of these changes. As long as we have a means of recovery we think we are safe. We have not questioned whether restoring from backup alone can meet our business critical recovery requirements. We have not investigated methods other than backup that can be used to minimize the business losses incurred when recovery is needed. We have not fully assessed the time needed to get our business critical applications back up and operational. We have not considered solutions which consider recovering data residing on multiple systems often in heterogeneous environments. It is time to think differently about our Recovery Environment.

What is YOUR Cost of Downtime?

Thinking differently means thinking in terms of your Cost of Downtime per application. This dollar figure exposes which application sets are important to your business processes and determines important investment decisions affecting high availability solutions as well as recovery and backup solutions. This figure includes the cost of unplanned vs. planned downtime, the unavailability during peak usage hours vs. off-hours, and how the application affects business profits when it is unavailable. Not all downtime is equal. A system interrupt caused by an OS failure, application abort or even a planned daily backup will have different costs than a longer term outage caused by a disk failure requiring a reload or a disaster which causes a Data Center move.

In determining an applications cost of downtime consider both direct and indirect costs. Direct costs include idle employment and manufacturing costs, delayed business processes, direct profit losses and penalties. We are usually good at identifying the direct costs of downtime. However, we often do not recognize and account for the indirect costs such as negative impressions on potential investor/partner/customer visits impressions, spoilage, branch or agency communication. In a recent survey of 150 HP 3000 customers in different industries: 45% had no idea what their cost of downtime is. Of the 55% that did estimate their downtime, most had not considered all direct and indirect costs of downtime, meaning that their estimate was low.

Industry Trends

Access to data in heterogeneous environments

An important driving factor in storage management today is that customers demand continuous accessibility to huge amounts of data, very often in the terabyte range. In the past, data accessibility was a fairly simple process when a single system was the primary storage device, and the only limitations were disk size. Today, the answers to storage problems cannot be provided simply by installing bigger disks on a central server.

As customers re-engineer their business, many are choosing to migrate off the mainframe via “mainframe downsizing.” Mission-critical applications are moving to cluster environments where applications on multiple systems share the same data or to client/server computing environments consolidated across LANs and WANs. Huge amounts of company sensitive data, which used to be under central control and located in the data center are now distributed and available within a multi-system (and often heterogeneous) environment.

In addition to all the challenges raised by managing storage on distributed systems, IT managers must deal with the reality that the amount of data is outstripping the network’s capacity to handle it efficiently. For example, a company might need to back up 100 Gigabytes (GB) of data in an hour. As the storage staff looks for solutions, they see processor performance improving faster than disk performance (I/O); and both disk and processors are outstripping the installed network infrastructure (bandwidth). The amount of data being moved from system to system, across a network, or pumped to backup devices is increasing. New ways of transferring and storing large amounts of data without downtime have to be developed.

Storage Capacity Requirements are exploding

For today’s IT organizations, server based storage requirements are exploding at 50% or even 100% per year. In the last few years, there have been many new trends and technologies within the commercial/business computing industry which have influenced the dramatic change in data storage capacity requirements and the need to manage it better. Systems capabilities have grown at dramatic rates in size and storage capacity. Cost reduction directions such as system consolidations and distributed computing environments (in which server systems centralize data storage needs for many clients) have caused the measurement unit of data storage capacity to change from megabytes to Gigabytes. Processing Power has increased dramatically allowing applications to grow in complexity and size. Incorporation of new storage intensive data types such as imaging, voice and video within many commercial/business applications over the next few years will only cause data storage capacities to increase ever more. Disk storage is more affordable and provides greater storage capacity per device. The growth in storage capacity requirements results in pressure to find recovery solutions for large volumes of data.

Business critical applications.

Customers are demanding that solutions minimize downtime and inaccessibility to critical data. These customers are looking for high availability features in nearly every solution from networks to recovery solutions. Storage management software solutions need to provide much greater data availability and reliability amidst a much more complex environment. Today, there is a much greater need to recover more data more quickly than allowed by restoring all data on a system. Often Store and Restore of data alone cannot meet business critical application/data recovery requirements. We need to look for other forms of recovery to meet today’s business critical recovery requirements.

Building your High Availability Bridge

Step One: Plan, Plan, Plan: Your Blueprint

Recovery Planning

Building your High Availability Bridge requires a detailed blueprint or a recovery plan. A Recovery Plan is probably the most important deliverable that you can produce which will minimize the cost associated with an interruption in service. This is also one deliverable entirely within your control and completely customized to your environment.

One of the most important needs in enterprise-wide storage is recovery. IT managers must establish a recovery policy that provides the appropriate level of data integrity and availability. The recovery policy must ensure that critical data can be completely and quickly recovered even in the event of a disaster. Your Recovery Plan document should include specific detailed recovery strategies and procedures for all scenarios your staff will need to recover from, including but not limited to: system aborts, application aborts, disk faults, power outage, network interruption, system component faults, user and operator errors, disasters. First priority should be to get business critical applications available with minimal business loss. An applications Cost of Downtime drives this priority as well as calculating the amount of downtime you can tolerate. Ask yourself the question; do I have the appropriate recovery strategy for each application? If not, look at ways of reducing the risk and minimizing recovery time (e.g. mirror disk/iX, Fast/Wide arrays, Shareplex), reducing downtime due to backup (7x24 True-online), using faster recovery/backup devices (e.g. DDS-3 or DLT), optimizing backup/recovery configurations (e.g. user volume application sets, massive parallel store/restores, interleaving).

A recovery plan should be tested with the operations staff which will be implementing the recovery in a real disruption. Perform dry runs of the recovery plan with different failure scenarios. Test, review and update your plan regularly. A good Recovery Plan is only good as long as nothing changes.

Backup Schedules and Plans

Backup is the #1 cause of planned application downtime today accounting for 83% of total planned downtime. As with recovery planning ascertain the length of your backup window. How much planned downtime can your business afford for backup? If you cannot tolerate any planned downtime and have chosen an on-line backup solution, when and how long are the windows where there are low application/system usage? Your backup window will help drive other decisions such as the speed and number of your backup devices and configuration policies such as the number of parallel stores and backup schedules. If application availability allows, the standard full/partial backup schedule is a good schedule. An alternative may be to rotate full backups and partials of major applications. Each major applications full backup is on a different day, combined with partials of the other applications.

Policies should be in place to minimize the amount of data in each backup stream. Your current backup probably includes data which is reference data or archival data. By removing the reference data and archive data from the active data backup, backup times can be reduced significantly. Put your backup on a diet. Don't backup data that doesn't need to be recovered or will already be recovered. For example, system data is recovered with SLT and FOS tape and STDLIST spoolfiles are unnecessary. You can often recover non-production utilities from other systems. Continually monitor and remove files that are no longer needed and can be placed in archive on less expensive media. Significant amounts of current on-line data storage capacity can be freed by use of automated data management activities such as file compression, trimming, or purging.

HP 3000

Example Backup Schedule of Large System

Monday

- Full backup of volume sets 1, 2, and 5 (about 20GB)
- Partial backup of other volume sets

Tuesday

- Full backup of volume set 3 (about 20GB)
- Partial backup of other volume sets

Wednesday

- Full backup of volume set 4 (about 20GB)
- Partial backup of other volume sets

Thursday

- Full backup of volume set 7 (about 25GB)
- Partial backup of other volume sets

Friday

- Full backup of volume set 6 (about 40GB)
- Partial backup of other volume sets

Saturday

- Partial backup of all volume sets

Sunday

- Partial backup of all volume sets

Special Case

- Full backup of system volume set only when changes are made. For example, new applications installed on system volume, configuration changes, etc.

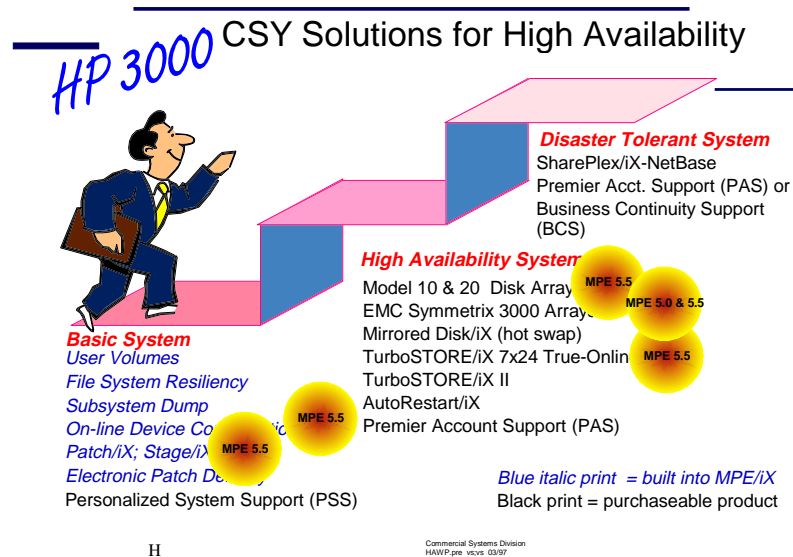
H

Commercial Systems Division
HWP/ptm vs/vs 0307

Step Two: Selecting your Recovery Environment

Determining what products are needed to reinforce and strengthen your High Availability bridge can be an intimidating task. Ultimately this decision should be based on your data availability requirements or your individual cost of downtime.

HP has identified three structures or systems for data availability. Each structure is designed to address a low, moderate, or high level of data availability. The three structures are 1) Basic Availability System, 2) High Availability System, and 3) SharePlex/iX-Netbase System.



Basic System Availability

The basic system of availability refers to the inherent robustness of HP Hardware and MPE Operating System. You do not need to buy additional hardware or software to attain this level and yet you can be assured that the system has the ability to withstand exception conditions without failing the system. Because the basic system is the foundation of your computer system (i.e. everything you do runs on top of the system) it must be very stable and robust. Driving a truck over a shaky bridge is not a safe prospect.

If recovery is required then it largely depends on recovery from a backup tape. Two types of recovery are available to you if you are running your operation within the Basic System Level; these are Full System Recovery and User Volume Set Recovery.

Full system recovery

A customer must restore the whole system when all data (system and application) is kept on the system volume set. With a single volume set, a disk fault requires recovery of all data, system & application (a full system reload). Any disk or system component fault causes the entire system and all applications to be unavailable. The recovery time is the longest of all recovery methods; approximately 4 to 8

hours. If your environment can tolerate this level of downtime this method of recovery is the least costly to implement.

HP 3000 System With A Single Volume Set


All data (system and application) is kept on the system volume set.

Availability compromised:

one disk fails, entire system reloaded

reload applies to every disk

reboot involves remounting entire volume set



H

Commercial Systems Division
HAWP.ppt v0.05 12/97


User Volume Set Recovery

You can significantly decrease recovery time just by using user volume sets. 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.

Users accessing other volume sets are not effected. Recovery of the entire system is only required if a disk failure were to occur on the System Volume Set. The advantage of this recovery method is it allows you to focus your recovery on your most critical applications. The User Volume Set recovery method reduces recovery time significantly while also increasing the fault tolerance of your critical applications. When making this segmentation be sure to separate reference data in order to avoid conflict with critical application recovery. In addition, reference data only needs to be backed up when there are changes. This recovery method is relatively inexpensive to implement but reduces recovery time significantly to generally within 1 or 2 hours.

HP 3000 System With User Volume Sets

Once system is rebooted, data on volume sets that do not have the disk fault can be accessed. Only need to recover data on faulting volume set from data backup.



H

Commercial Systems Division
HAWP.pre v6 vs 03/97

Tips for the System Volume Set

To ensure maximum fault tolerance as well as reducing recovery and backup times make the system volume set as small as possible. It is very important not to allow permanent user data on the system volume set. With all user data restricted from the system volume set full backups of the system volume set need only be done when new software, configuration changes, etc., are added to the system. If the System Volume Set is small, make a combined SLT storeset of system volume set using SYSGEN. For a larger System Volume Set, backup the volume set with multiple storesets for faster recovery.

Tips for the User Volume Set

Minimize the number of applications per volume set. If practical aim for one application per volume set. Create a general user volume set for all non-application specific user files and don't put them on the system volume set. For your business critical applications use mirrored disks within a High Availability Storage System (HASS) enclosure. Do not split the mirrored volumes during backup, instead use, 24x7 True-Online. Rotate full backups of the user volume sets and use the DIRECTORY option when doing volume set backups. When identifying application sets use the following guidelines.

- System code which is recovered with a system load tape SLT created at install/update time.
- Other system utilities and 3rd party tools which are not modified.
- Major applications (which are the reason the system is needed). Some may be critical to your business processes, some not.
- Extraneous data which does not require recovery.
- Development Data (source, tests, etc.)
- Client system data (client data storage, client backups)

High Availability System

With the basic system, we have made the foundation strong and solid. We recognize, however, that downtime still occurs and this must be addressed in order to achieve a higher level of system availability.

Businesses that require higher data availability than that of the basic system, typically associate a higher cost to either planned or unplanned downtime. For these businesses, the cost associated with purchasing products to achieve greater availability is far less than the cost of needing the data protection or fast recoverability and not having it.

This level of data availability is achieved by eliminating your vulnerability to the I/O subsystems five points of failure and by reducing or eliminating downtime associated with backup.

The I/O subsystems five points of failure are the I/O card, cable, disk controller, disks, and power supplies. Eliminating your vulnerability to these failures requires redundancy along this whole path and is achieved either through hardware solutions such as disk arrays or through software solutions such as MirrorDisk/iX. Each solution has unique benefits that meet particular needs.

Hardware Solutions: Disk arrays

The HP 3000 offers two hardware solutions. The Model 10 & 20 is positioned for customers with small capacity needs in the Workgroup or Business unit. The EMC Symmetrix 3000 is positioned for the Data center with high growth potential of >100GB or within a heterogeneous environment.

Disk arrays tolerate an outright failure of any single disk mechanism within the device without losing data or interrupting the host system. This is the most vulnerable point of the I/O subsystem. Redundancy of the I/O channel and cable is not provided with a High Availability disk array as it is when using Mirrored Disk/iX, but disk arrays do reduce the risk that a recovery from a backup will be required. We recommend that the system volume set on business critical systems be protected with disk arrays.

Taking advantage of the multiple paths that are provided by disk array solutions does require the host to recognize a failed path and to take an alternative route if necessary. This functionality will be available to HP 3000 customers in early 1998. At that time, Disk Arrays will provide complete protection for all of the I/O Subsystems 5 points of failure.


Software Solution: MirroredDisk/iX

MirroredDisk/iX provides full redundancy of the I/O card in system, the data cable, the disk drive, and the power supply. A failure along this I/O path does not result in application unavailability. With Mirrored Disk/iX when a disk fault or other I/O subsystem fault is detected, the mirror of the volume set takes over as though no error occurred. The re-activation of a failed disk can often occur without taking the system down. The cost to implement this recovery environment includes the purchase of Mirrored Disk/iX and additional disks for mirroring. However, the recovery time is minimal requiring only approximately 40 seconds to activate the mirrored volume. MirrorDisk/iX does not protect the system disks. For higher availability of the system disk then a disk array solution is recommended.

HP 3000

System With Mirrored Disk Volume Sets

Disk Fault does not fail entire system with Mirrored Disk/iX.
Once a disk fault is detected, the mirror of the volume set takes over as though no error occurred.



H

Commercial Systems Division
HAWP:pre-vol. 03/97

Multi-system considerations

The Symmetrix 3000 arrays support a broad range of platforms, guaranteeing users long-term investment protection on their storage purchase. Symmetrix arrays' multi-host connectivity allows customers to consolidate their storage requirements in both homogenous or heterogeneous environments. Simultaneous connectivity enables users to leverage the performance, availability and manageability attributes of Symmetrix disk subsystems while leveraging a single technology investment. EMC Symmetrix 3000 allows customers to support 16 heterogeneous hosts simultaneously, including UNIX, IBM OS/400, Novell NetWare and Microsoft Windows NT operating environments. A customer can reallocate the capacity among the different hosts when application needs change. Supporting a broad range of platforms, Symmetrix 3000 guarantees users long-term investment protection on their storage purchase.

Shareplex/iX-Netbase system

The High Availability system does not protect against system component failures or a disaster which effects the entire data center. The SharePlex/iX-NetBase system, on the other hand, addresses the highest level of data availability. This system provides high system and application/data availability, disaster tolerance, shared system resources, and expansion options through horizontal growth.

Disaster Tolerance

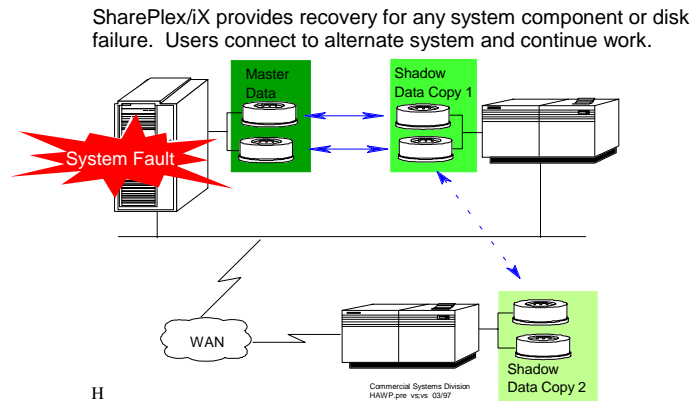
Disaster Tolerance is achieved via the Shareplex/iX-NetBase Shadowing product. Applications and data are replicated in real time on separate servers. In the event of a node failure, another system will take over applications that were running on the failed node. Recovery is available for any system component or disk failure. By Providing full redundancy for protected applications you get the best protection for very business critical applications. Cost to implement this solution includes purchase of

Building a High Availability Solution

4010-9

Shareplex/iX shadowing and access to an alternate system. The approximate recovery time for protected applications is 5 to 10 minutes.

HP 3000 Systems Using SharePlex/iX To Shadow Application Data



Shareplex/iX for an MPE/iX cluster environment

Data availability means more than resiliency from an unforeseen event or from a maintenance procedure. Data availability also means access to data that physically resides on another machine. Implementing a cluster environment with Shareplex/iX gives you access to data distributed across the cluster.

In multi-system environments you want to take advantage of all the system resources available to you. Access to resources and data is critical to a high data availability system. SharePlex/iX-NetBase can make resources on any system on the network available to all users. Files, databases, printers, and programs can be transparently shared among users on the network, regardless of geographic location.

With a cluster approach, the SharePlex-Netbase bundle gives Mid-range to high-end customers virtually unlimited growth potential. Many existing customers have run into various limits with operating systems, table limits, locking and concurrent end-user limits. Two products within the Shareplex-Netbase bundle give customers the ability to transcend these limits by implementing a clustering strategy. These products are AutoRPM and Network File Access.

The SharePlex/NetBase Network File Access (NFA) allows network administrators to spread applications across multiple machines. NFA allows you to logically extend your local file system to include remote HP 3000 file systems. NetBase maintains a centralized directory of all files and databases that are available to network users. The directory will automatically direct the application to the appropriate location of the file. Spreading applications across multiple machines isolates application usage, increases throughput, and offers maximum flexibility in upgrade paths.

AutoRPM allows users transparent access to programs located on a remote machine. Users gain instant access to virtually any software on the network. AutoRPM automatically transports the user to the appropriate server within the cluster.

Shareplex/iX for a Multi-system heterogeneous environment

Sometimes, and increasingly so, the cluster may include different Operating Systems. Extending the cluster to these divergent environment presents its own challenges. Shareplex/Netbase provides an extensive suite of networking services that allows a network of HP computers to operate cooperatively as a single computer system to maximize availability and improve overall performance. It creates a fault-tolerant heterogeneous network by integrating MPE,UNIX and desktop systems.

SharePlex Bridge provides real-time capture of transaction data on MPE. It then makes the data available through an open format and standard ASCII representations to applications running on other corporate platforms, such as UNIX and Windows (3.x, 95, and NT). Shareplex Bridge can be utilized to continuously feed TurboIMAGE, KSAM or MPE transactions to UNIX, NT or RDBMS such as Oracle, Sybase and Informix. It enable the HP 3000 server to coexist in a heterogeneous environment. New applications residing on other servers such as UNIX or NT can transparently access a true representation of data that originates in a legacy application residing on the HP 3000.

NetBase Client consists of complete Application Program Interfaces (APIs) that give UNIX, DOS, Windows, Windows NT/95/3.x and DOS client/server applications read and write access to data residing in MPE files, KSAM files and TurboIMAGE databases on HP 3000 servers. NetBase Client provides a critical link between HP 3000 legacy environments and the newer platforms, such as UNIX and Windows NT.

Step Three: Backup for Higher Availability

As you move your business from the Basic System to the High Availability System to the Shareplex/iX-Netbase system you become less dependent on a backup tape for recovery. Backup in higher level systems is extra insurance in case your primary recovery system fails. All the same, it is insurance that you do not want to forego.

Strategies for backing up data range from small shops able to back up data at night to enterprise wide backups of heterogeneous clients and servers. In a small shop there is little data to backup, a job can be scheduled at night after the close of a business day. Very little goes wrong and the jobs complete easily by the start of the next business day. This company's availability requirements are easily met with the "Basic System Availability." As companies grow with more data to back up it becomes more and more difficult to complete backups within this allotted time. This leads many companies to adopting night shifts to change tapes or to look for other ways to complete backups within an allotted window. One typical way is to break up the store process by running massive parallel stores on separate user volume sets. Additionally, faster larger capacity devices can be used such as DDS-3 or DLT tape devices. Some companies have had great success using DAT autochangers. The use of autochangers can help to eliminate tape changing delays and the need for attended backups. To completely eliminate downtime due to backup, some customers have gone to on-line backup. This allows users and jobs to continue modifying databases and files while the backup is occurring. TurboSTORE/iX 7x24 True-online allow for this functionality.

Backup solutions require selecting the appropriate Backup software as well as the appropriate hardware devices. In addition, the recovery environment should be considered. Backup in a Shareplex./iX environment will look quite different then that found in a an environment using MirrorDisk/iX or those environments using User Volume Set Recovery. Finally, special consideration should be given to backup in a heterogeneous environment.

Select Appropriate Backup Software solutions

There are a number of software products for the HP 3000. These include products from Hewlett Packard such as Store/iX, TurboSTORE/iX, and TurboSTORE/iX 7x24 True-Online. In addition there exist third party software solutions also available on the HP 3000.

HP 3000 TurboSTORE/iX on MPE/iX 5.5

Feature	STORE/iX FOS 5.5	TurboSTORE/iX B5151AA	TurboSTORE/iX 7x24 True-Online B5152AA
7x24 True-Online	NO	NO	YES
Online Backup	NO	NO	YES
Store to Disk	NO	NO	YES
Parallel Restore	NO	YES	YES
Data Compression	NO	YES	YES
Label Tapes	YES	YES	YES
File Interleaving	YES	YES	YES
Multiple Store Devices	YES	YES	YES

H

Commercial Systems Division
HAW P-pro vs vs 03/97

Store/iX

STORE is an excellent choice for small shops that do not require online backup or environments with little room for delays. STORE is included with the OS and offers basic functionality. STORE is limited in its file selection capabilities. In most cases customers use other tools to generate more precise file lists at the cost of backup time. In addition STORE is unable to specify more global parameters (e.g. full DB but partial files). STORE nor TurboSTORE offer optional tape Librarian utilities that track which file ended up on which tape. Libraries make it easy to identify tapes that have files needed for recovery.

TurboSTORE/iX Tips

TurboSTORE/iX products for the HP 3000 provide high performance backup solutions designed to meet today's backup requirements. TurboSTORE/iX products offer powerful parallel backup and recovery, data interleaving, data compression, and online backup capabilities. The on-line backup feature does require quiescing the system prior to backing up data.

To take full advantage of TurboSTORE/iX products use the following tips.

- Use multiple parallel storesets to gain throughput and performance of backup. You also gain better performance during recovery.
- Use MAXTAPEBUF for larger I/O blocks, improves performance when using fast backup devices.
- Use device hardware compression when available to minimize CPU overhead.
- Use s/w compression if hardware compression is not available for the device.
- If reading from more than 3 disks, use the INTERleave option.
- Use DIRECTORY option on all major backups.
- Use an integrated DB backup so DB and other non-DB application set data can be backed up together on-line.
- Store backup media in a safe place (a data vault or fireproof safe). Do not leave your business critical recovery data laying around the office!
- Verify your backup. This can easily be done on another system (with MPE 5.0 or later) to reduce overhead on your main system.
- When performing online backup chose a time with the lowest activity on the system to minimize CPU overhead during busy system use.

TurboSTORE/iX 7x24 True online

As businesses move closer toward continuous operations, IT managers find the growing need for solutions that can meet the demands of a 7 days per week, 24 hours per day environment. TurboSTORE/iX 7x24 True-Online Backup was specifically designed for 7x24 environments by providing backup of selected data without requiring application downtime or user log off. In addition True-Online provides the same powerful backup capabilities of previous versions of TurboSTORE/iX.

Select Appropriate Backup Hardware Configuration

Hardware Configuration for Maximum Throughput

Selecting the appropriate hardware configuration is more than just selecting the appropriate backup device. With MPE/iX 5.0 you can have up to 32 backup devices on your system. By using multiple devices in parallel you can increase your data throughput. To get the maximum throughput use the following procedures.


1. Disk devices and backup devices need to be spread across multiple device adapter cards for maximum backup performance especially if total backup speeds of over 12GB/hour are needed. For maximum backup performance the system configuration should contain no more than 4 disks per card.
2. SE SCSI cards cannot sustain more than 10GB/hour data transfer across the bus. Backup devices can share the SCSI bus, up to a total of 10GB hour. therefore, for maximum performance - no more than 3 DDS, 2 DLT, 4 7980S, 2 7980SX, should share a SCSI bus.
3. Where TurboSTORE/iX is used to compress data the INTERleave option should be used.
4. The impact of reel switch can be avoided by using TurboSTORE/iX II Sequential device or Sequential device pool functionality. Rewind times can be significant.
5. Use the INTERleave feature when storing data from more than 3 disks.

Explosion of tape technologies

There has been an explosion of new tape technologies into the market place. These technologies can be grouped into three categories: low-end, midrange, and high-end.

- At the low-end are DDS and 8mm devices. These were developed as spin-offs of consumer/entertainment applications and have had great success penetrating the computer data markets. However, larger customers are becoming impatient with inherent reliability problems and throughput and capacity of these technologies.
- The midrange includes DLT-4000 and Mammoth. Mammoth is a future technology from Exabyte based on 8mm technology. Development schedules have continued to slip significantly and it appears that the market has moved to DLT and as a result the opportunity for Mammoth has sharply diminished. DLT-4000 is one of a portfolio of products. DLT was developed specifically for computer data storage. It offers high reliability and solid throughput and capacity. HP has seen huge interest in DLT among its customers.
- DLT-4000 also penetrates high-end tape technologies. Also in this space are the StorageTek Silo technologies such as 3480/3490/3490E, TimberLine, and Redwood. All these devices attach to STK Silos.

HP 3000 Transfer Rates and Media Capacity of Industry Tape devices

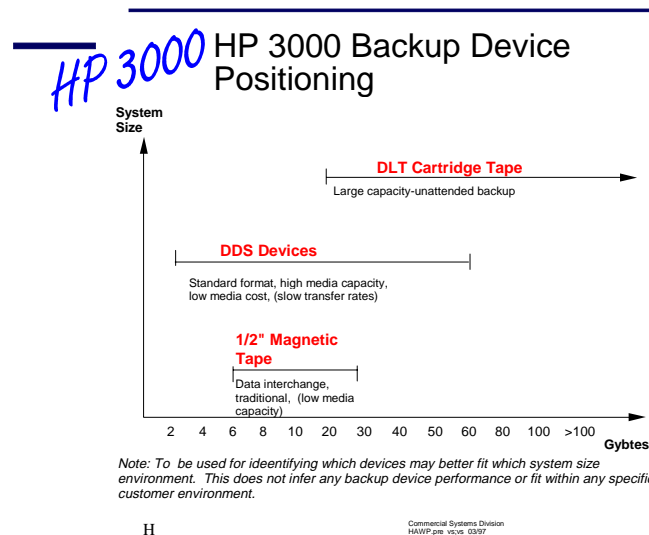


H

Commercial Systems Division
HAWP pre vs vs 03/97

Selecting a Hardware device

HP solutions for the HP 3000 are based on customer requirements for amount of data to be stored and time in which to store it. For customers with small datasets and longer backup windows, DDS may be a very appropriate solution. At the mid range and high end DLT-4000 mechanisms provide backup for customers with large amounts of data and limited backup windows. Based on throughput and capacity, multiple DLT-4000 mechanisms can meet the needs of customers with large volumes of data and aggressive backup windows.



HP DAT Products

For backup on the HP 3000 HP offers the latest DDS-3 tape drives. The new DDS-3 format has a native mode capacity of 12GB. With data compression, customers can typically store 24GB on a single tape. DDS DAT is an industry standard high capacity device with high reliability. Its compact media is easily stored in a fire proof safe.

Digital Linear Tape

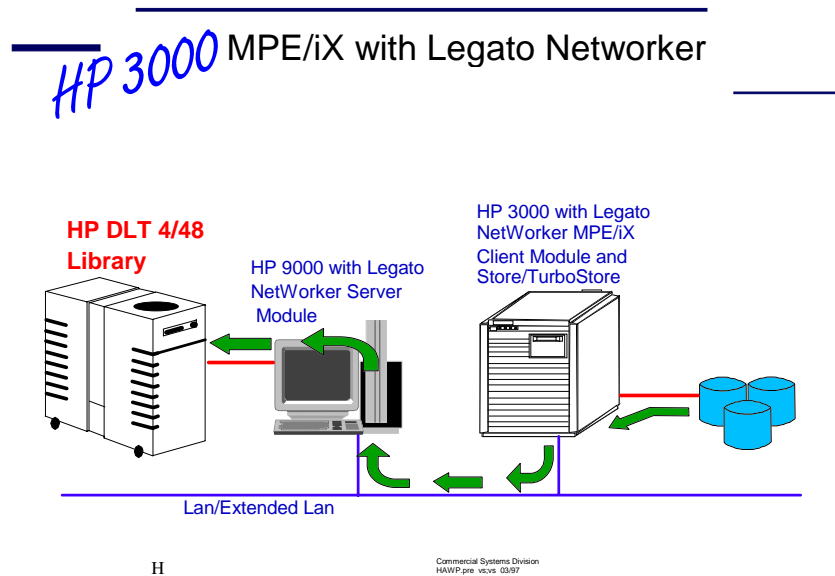
The HP 3000 servers support automated digital linear tape (DLT) mechanisms. DLT/4000 provides greater native cartridge capacity (20GB) enabling fast, unattended backup of large quantities of data within the brief windows available for backup in today's high-end, mission critical environments. The DLT native transfer rate (5.4GB/hr) is 60% faster than DDS-3 (3.6GB/hr). Besides large capacity DLT boasts superior drive head longevity.

Backup in a heterogeneous environment

NetWorker from Legato is now available on the HP 3000. NetWorker is a client/server application that provides advanced storage management capabilities to a wide variety of servers and desktop

Building a High Availability Solution

computers. It is an excellent example of a product well-established in the UNIX world that will provide HP 3000 customers with better integration in a multi-environment. HP 3000 customers can now join with other clients to backup their data to a UNIX server. An HP 9000 running NetWorker for HP-UX with its backup drives or autochangers (any tape robotics device such as DLT) functions as a backup server. Other computers on the network function as NetWorker clients. Besides the HP 3000 running MPE/iX, the clients supported can include HP-UX and other types of UNIX, Windows, Macintosh and Windows NT desktops, NT and NetWare servers.



The main barrier to large backups over a network is the available network bandwidth. With 100VGI and later with Fibre Channel, large backups become more feasible. We are seeing requests today from customers that want to perform networked backups rather than local backups. Even with today's network bandwidth policies and procedures can be put in place to make a network backup (and more importantly) a network recovery more feasible. These policies include a divide and conquer strategy by implementing user volume sets and by subdividing the data into multiple store processes.

Backup using User Volume Sets

Using User Volume Sets decreases downtime for backups. 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.

Backup in a Mirroring Environment

Due to the mission critical nature of these applications, these environments require rapid backup and restore functionality. We do not recommend splitting mirrors during a backup since the customer would be exposed during this time. Instead, with the use of 24x7 customers can continue to keep their mirrors and perform backups without any application downtime.

Backup in a Shareplex Environment

Most likely you will never need to recover your system from tape media, however, the use of Shareplex does not eliminate the need for backup. In a Shareplex/iX environment it is recommended that a 24x7 backup be performed on the less critical shadow system. In this way the overhead generally associated with on-line backup is eliminated.

If using multiple recovery methods (like Mirrored Disk/iX and SharePlex/iX), potential recovery using backup is less likely. In this case bi-weekly/monthly backup of full application and only DB logs and changed non-DB files daily. This maintains a recovery path if the higher recovery methods fails.

Future Requirements for Data Availability

As HP designs and implements stronger and more robust High Availability bridges, we can't lose sight of some important trends in the storage environment of our customers. Faster network capabilities will make centralized network more feasible. As HP 3000 systems co-exist with heterogeneous systems, customers will want to recover and backup multiple heterogeneous clients and servers from a single centralized server. FibreChannel will become the dominant storage interface over the next 3-5 years. This will provide fast interconnect as well as the availability of faster storage devices. Data capacity requirements will continue to grow requiring solutions to recover and backup Terabytes of data. Cost pressures continue to emphasize solutions that support operatorless environments. More and more customers will want storage management functionality that allows on-line backups and simple, sophisticated media management. Customers will continue to demand solutions that minimize if not eliminate inaccessibility to critical data. The HP 3000 High Availability Bridge will continue to provide more robust and higher levels of data availability in ever increasing complex and more troublesome environments. So ease your mind of those troubled waters the HP 3000 Bridge of High Availability will lay you down..