**HP WORLD Õ96**

# Backup and Recovery Solutions for the HP 3000

Presentation Number 1022

**Vicki Symonds**

Hewlett-Packard Co.
c/o Liz Ynegas

*Phone:* 408/447-3213
*Fax:* 408/447-4916

# Why Backup?

Typically we think in terms of backup only. We develop backup plans not recovery plans, we monitor backup times and resources used, we use metrics to reduce backup costs and resources used. We often do not even consider the reason that we are always backing up systems—RECOVERY. In our attempts to be realistic, we focus and gain expertise on backing up systems efficiently or reducing the cost of backup. We focus on making efficient use of backup resources and inadvertently compromise the ability to recover. 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.

Up until 5 years ago, recovery of a full system from a backup tape provided acceptable recovery times. Systems were still quite small and tape backup and recovery only took 1 to 2 hours for full system recovery. Things have changed tremendously over the last five years. We now need to think differently about recovery and backup.

## 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 then 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.

# Data Storage Trends

## 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 backup/recovery solutions for large volumes of data. Increasingly, HP is being asked for backup/recovery solutions for Terabytes of data to high end peripheral devices.

## 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. You can see this demand occurring by the increased use of the Internet and on-line services such as CompuServe and America On-line and the emergence of new applications such as imaging and multimedia. In the past, data accessibility was a fairly simple process when mainframes were the primary storage devices, 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 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 in the network. Today, the average amount of distributed data has surpassed the average amount of data in the data center.

Companies must begin viewing storage management as integrated to their network solutions. 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.

### Business critical applications.

Customers are demanding that solutions minimize downtime or inaccessibility to critical data. These customers are looking for high availability features in nearly every solution from networks to backup/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.

### Cost reductions

IT managers are looking to reduce costs of operations. Data storage management strategies must be tuned to require minimal operator intervention and make use of less expensive storage devices. Toward this end unattended backup solutions and automated tape libraries are attractive to customers. Many IT managers are looking to lower costs by eliminating backup altogether at branch offices and/or remote sites.

## HP 3000 Customer Recovery Environments

### Backup/Recovery Planning

Recovery Planning

One of the most important needs in enterprise-wide storage is backup/recovery. Very early in a solution deployment, IT managers must establish a backup/recovery policy that provides the appropriate level of data integrity and availability. The backup/recovery policy must ensure that critical data can be completely and quickly recovered from a backup 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-2 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 application downtime today accounting for 83% of total current 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.

## Selecting a Recovery Method

Select recovery methods or combinations of methods that will decrease downtime revenue losses by more than the cost to implement the recovery method. Several recovery methods are available to you providing differing levels of high availability; Full system recovery, Application set recovery from backup, Application volume set recovery with disk arrays and mirrored volume sets (mirrored disk/iX), application recovery using shadowing (Shareplex/iX).

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.

Application 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.

***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 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 Disk Arrays

High Availability disk arrays tolerate an outright failure of any single disk mechanism within the device without losing data or interrupting the host system. Redundancy of the I/O channel, cable, and power 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.

Mission Critical Environment using mirrored disks

Disk failure of a mirrored disk does not make system or applications unavailable. With Mirrored Disk/iX when a disk fault is detected, the mirror of the volume set takes over as though no error occurred. The re-activation of failed disk can often occur without taking the system down. Mirrored disks provides full redundancy of I/O card in system, data cable, disk drive, and power into disk drive however it does not protect against system component failure nor can you mirror the system volume set. The cost to implement this recovery environment includes the purchase of Mirrored Disk/iX (approximately $1500 to $26,300 US dollars) and additional disks for mirroring. However, the recovery time is minimal requiring only approximately 40 seconds to activate the mirrored volume.

OLTP - Mission critical Environment using Shareplex/iX

In this high availability segment, 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 Shareplex/iX (approximately $14,000 to $125,000 US dollars, depending on system size and bundle) and access to alternate system. The approximate recovery time for protected applications is 5 to 10 minutes.

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 still maintaining a recovery path if higher recovery methods fail.

# Backup for Recovery and System Availability

## Backup Environments

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. 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-2, DLT or 8mm (from IEM). Some companies have had great success using DAT autochangers (also from IEM). 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. Both TurboSTORE/iX 7x24 True-online and a utility from Orbit Software allow this functionality.

To improve your data availability the following environments should be considered..

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.

## Network Backup

Many HP 3000 users have initiated projects to develop an enterprise-wide backup/recovery strategy. The single greatest limiting factor to a totally centralized backup architecture is network bandwidth. With large volumes of distributed data totally centralized backup architectures are unfeasible. Additionally, full volume restores resulting (e.g. disk head crash) can require significant time tying up the network during business hours. Most customers end up with a multitiered architecture with local tape and limited centralized backup. These circumstances favor backup/recovery vendors supporting heterogeneous systems such as Legato.

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. With a higher speed backup data from the clients on the network are moved to a single server which then stores the data to local tape drives.

## Select Appropriate 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 such as Orbit's Backup/3000 and Unison's RoadRunner.

### 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 a snap to identify tapes that have files needed for recoveries. More sophisticated media managers are also not available.

### TurboSTORE/iX Tips

TurboSTORE/iX products for the HP 3000 provide high performance backup solutions designed to meet todays backup requirements. TurboSTORE/iX products offer powerful parallel backup and recovery, data interleaving, data compression, and online backup capabilities. 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 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. 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 4 DC-DDS, 2 DLT, 4 7980S, 2 7980SX, should share a SCSI bus.
3. Where TurboSTORE 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 for 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.

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-2 tape drives in addition to DDS and DDS DC drives. The new DDS-2 format combined with 120-meter tapes, has a native mode capacity of 4GB. With data compression, customers can typically store 8GB 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) than DDS (4GB)

enabling fast, unattended backup of large quantities of data within the brief windows available for backup in today's high-end, mission critical environments. DLT native transfer rates (5.4GB/hr) is three times faster than DDS-2 DAT(1.8GB/hr). Besides large capacity DLT boasts superior drive head longevity.

# Future Requirements for massive backup/restore

As HP designs and implements solutions for the various environments discussed, 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 back-up multiple heterogeneous clients and servers to a single centralized server. FibreChannel will become the dominant system interconnect and peripheral interface over the next 3-5 years. This will provide fast interconnect as well as the availability of faster disks, arrays, and tapes. Data capacity requirements will continue to grow requiring solutions to backup and restore 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. Overall of these customer requirements is the umbrella of high availability solutions. Customers will continue to demand solutions that minimize if not eliminate inaccessibility to critical data. Storage management solutions need to provide much greater availability and reliability amidst a much more complex environment.