

Are backup and disaster recovery measures necessary in today's world? white paper



Executive summary.....	2
Why backups?	2
Defining the data status (consistency)	4
New disk technologies	4
Understanding disaster recovery, restart, and recovery.....	4
Disaster recovery	4
Restart	5
Recovery	5
Drawbacks with backups.....	6
Data backup techniques	7
Backup in mainframe	9
Schematic presentation.....	10
Mirroring in a local subsystem	10
What can be done with these local mirror copies?	11
Mirroring at a second location	11
For more information.....	13

Executive summary

This white paper discusses the concepts and definitions of backup, disaster recovery restores and recovery of data, as well as the drawbacks of each solution. This paper also discusses how, through the use of HP products, companies can be assisted in maintaining their Service Level Agreements (SLAs) and in optimizing their workflow, which in turn increases their value.

Why backups?

A common thought today by IT technicians is that it is not really necessary to take backup precautions because of new disk technologies with RAID processes as well as mirroring on peripheral sites. This presumption is false. Some common rebuttals often given by IT departments against implementing a disaster recovery solution are:

“Disk-failures are so rare; it’s not worth the trouble making backup copies. Besides, our current RAID-system fixes itself.”

“Disk prices are so cheap that we can mirror all of our data.”

“Intelligent disk-systems offer functions that activate data mirroring in seconds.”

“If worse comes to worse, we can always access our remote disk system.”

“We copy all of the server data over a network connection to a backup server on a weekly basis.”

“We don’t need magnetic tape. Time is not a factor.”

What can possibly happen?

- Databank crashes?
 - We have the redo log files online.
- Fire? Flooding?
 - We can move to a mirror version on a remote site.
- Bugs in our house software?
 - That is the programmer’s problem.
- Virus attacks?
 - We are protected by a firewall.

But what happens when:

- A databank failure is first noticed after several days?
- A batch program “corrects” a large section of data?
- Viruses invade the network through a user’s PC?

Then, help is needed:

- “My mirrored data is worthless; it’s as bad as the original.”
- “The remote mirror doesn’t look any better.”
- “The online logs don’t help much.”

What about the IRS auditors or the in-house auditors wanting to access and examine the programs and data from last year?

The fact is:

- Users are responsible for most errors (release changes, upgrades, software bugs, operating errors, and so on)
- Auditors want to understand the business flow.
- The legal requirements for data storage must be maintained—“and it will not be possible without backups and archiving.”

Backups are important:

- For databanks necessary for business:
 - Full backups once a week
 - Incremental backups daily
 - Independent of the size and how often updates are made
- For other dynamic data:
 - Every day, completely, for at least one week
 - Eventually, accumulated weekly and monthly data
- For financially relevant data:
 - Yearly and if necessary quarterly
- For data that must be legally archived for a certain time:
 - Backup on WORM media

For businesses, data backups, duplication, or both help eliminate the problems in the event of a failure or malfunction, and have become legal requirements. Changes in data, such as account transactions, must be recorded before and after they are applied, and such that these step-by-step modifications and changes can be monitored. Additional measures are necessary to help avoid the loss of data in the event of some physical failure or accident, such as the deliberate erasing of data, user error, or an error in the program. RAID systems cannot protect against these kinds of error events.

With RAID, techniques ensure uninterrupted access to the data in the event of a hardware failure, and increase the application availability. The logical integrity of the data cannot be assured through hardware functions.

System hardware is responsible for the data being written and how it is made available to the server and the application. The logical structure of the data is not monitored by the hardware. To minimize the intrusion of logical errors in the data system over time, a backup copy of the original data should be retained so that, in the event of a failure, it is possible to recover the data to some defined error-free point.

Defining the data status (consistency)

Data status refers to data with the same content and level available at a specific point in time (PIT copy). A PIT copy refers to specific data, a logical subunit of a disk subsystem, or a device (LUN/LDEV), and reliable data backup requires the synchronous backup of all associated data at a defined PIT (consistency groups). Copying the data onto magnetic tape is the most common technique, but this causes an interruption in the system availability. Online techniques such as concurrent copy or techniques, which allow for copying active databanks, help to increase the backup window and application availability. But these techniques may lead to a decrease in performance.

New disk technologies

Disk technologies like HP StorageWorks arrays make it possible to define consistency groups (that span a multi-array configuration) that can be split in at precisely defined timestamps (AT-split). This requires the use of HP StorageWorks Business Copy or FlashCopy and allows the data to be available for a forward recovery in the event of an error. The second copy can then be transferred to another media such as magnetic tape, because saving the copy on another disk, or disk subsystem, will not ensure that the data is available in the event of a failure. Regular copying of the data over short time intervals reduces the recovery time necessary after a loss of the logical integrity of the data. It is therefore essential that all of the dynamic data, that is, redo logs (protocol file of databank changes), be part of the copying procedure.

Understanding disaster recovery, restart, and recovery

The concept of restart and recovery is often confused with that of disaster recovery. Disaster recovery procedures are not the same as the normal restart and recovery procedures necessary to recover business capability as a result of inconsistent data or a loss in the logical integrity of the data. For this reason, disaster recovery, restart, and recovery will be discussed separately.

Disaster recovery

To adequately describe disaster recovery it is first necessary to define what is meant by a disaster. The dictionary describes disaster as “a great and sudden misfortune.” Disaster recovery measures will therefore describe and lay out plans as how to best prepare for such a “misfortune,” even though it is not possible to know which form this will take. It is therefore necessary to prepare beforehand a computer center **disaster plan** or **disaster handbook**, which describes what the various degrees of a disaster are and their implications, for example, the complete loss of the computer center building (highest degree) or the loss of the disk system through fire or an electrical malfunction (low degree). The most important part of a disaster plan is the naming of the persons with the responsibility for declaring a disaster, so that steps can be taken according to the disaster plan. This responsibility includes defining when and what measures will be taken to regain the availability of the system. The disaster plan defines which applications will be started first and which can be started later. It defines which limitations will be accepted, that is, performance or the number of users, in reestablishing the system availability. Disaster recovery therefore defines the responsibilities of the computer centers in bringing the applications with the highest priorities online promptly.

As stated earlier, application availability is assured through RAID technology in a box and mirroring assures disaster recovery. The aims of these techniques should not be confused, for example, using remote copy without RAID technology. Mirroring has increased the application availability, but as the mirrored copy contains the primary and secondary data, there is although no disaster security. A switchover during a disaster event (rolling disaster) results in data on the primary and secondary system that may no longer be consistent and that is unusable.

Restart

Restart, on the other hand, means the ability to bring an application and its data, from a certain PIT (checkpoint), back online, and is the fastest way to reactivate a crashed application. For example, after a power outage, which is in itself no disaster, and a restart of the system, the applications can be restarted from a predefined point without needing additional data, as the data included in the copy made at this checkpoint is complete and can be used to restart the application, ensuring data consistency. Unconfirmed databank changes are returned to the pre-error condition; this is referred to as back-out.

A manufacturer solution to this problem is Oracle® Parallel-Database Server and the IBM Mainframe side DBRC-Recovery procedure.

It is important to note that for a successful restart, not only does the data have to be physically available, but its logical consistency must be assured.

Recovery

Recovery is necessary when the data is lost or logically inconsistent. Recovery copies a former version of the data, such as a copy available on the same disk subsystem or from magnetic tape, back onto the active disk system. This is commonly referred to as a restore procedure. Restoring from a backup copy located on the same disk system is the fastest method to regain data consistency.

Databank systems can restore the system availability using the data from these copies with help from the transaction records (redo logs). It should be obvious that backups must not only secure the databank system, but all information pertaining to the databank, that is, transaction records. Backups alone of the databank system without copies of the transaction records will result, in the event of a disk failure, that although the databank can be restored, data loss will have to be accepted. For example, a backup point is made the day before at 6:00 p.m. When a failure occurs the next day at 11:00 a.m., the records of the dynamic data generated since the last backup point cannot be restored. This shows that backing up is more than just a copy of the databank or the work data; it is an organized procedure that ensures data consistency, at any time.

Available disk technologies make it possible to speed up these procedures increasing the productivity and value for the computer center and the customer. But an increased overhead and cost have increased disk capacity necessary for duplicating the data. Copying data onto remote or peripheral magnetic tape systems is not eliminated, but less frequently necessary. This means that in a computer center environment, both disaster recovery and backup/recovery must be thoroughly considered. Only through a well-planned strategy can unnecessary costs be avoided.

For the customer, optimizing the existing procedures with the currently available techniques brings not only an increased value in terms of the service availability, but also an increased value through a reduction in the administrative overhead caused by older procedures.

Drawbacks with backups

Figure 1 and Figure 2 illustrate, according to a survey of IT users (computer centers), the biggest drawbacks.

Figure 1.

Backup restrictions


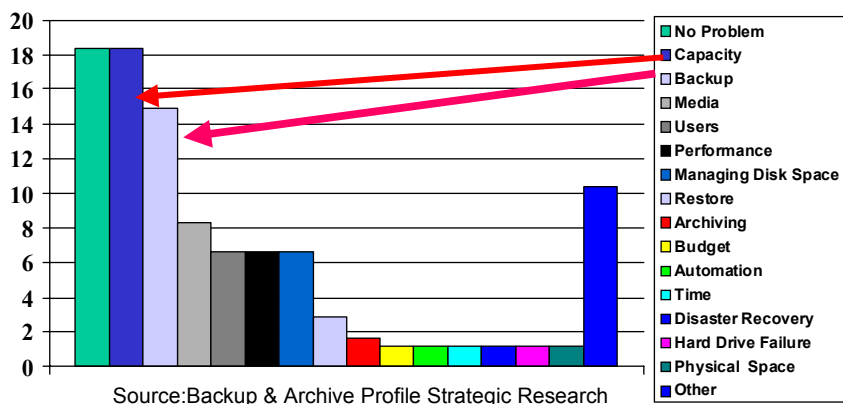


Figure 2.

Storage problems



One of the biggest drawbacks in carrying out backups is the increased run-time needed. Increasing capacity needs for an application results in an increased timeframe that is necessary to transfer the data. The limitations using currently available implementations and older systems result in a difficulty in maintaining the SLAs with the customer.

The roll of handling and the management of the disk system to the increasing data quantities play another important factor.

Data backup techniques

Figure 3 illustrates how data backup is organized under various system configurations. One possibility is a LAN-based backup procedure, where data is transferred over an existing LAN to a backup server. Data from the various disk systems is transferred to the backup server, collected together, and saved as a backup on a save place.


The timeframe of data flow is dependent on the LAN bandwidth, the capacity of the backup server, and the quantities of data to transfer. One technique developed to increase the dataflow is a LAN-free backup. The backup process is started, as usual, from the backup server through the backup agents on the various data servers. The data is then transferred through a high-speed fiber optic cable to the magnetic tape drive. The load on the network is reduced to the administrative information and necessary backup server functions.

Bottlenecks might occur as various servers try to share a single tape drive. The dataflow is now not only limited by the fiber optic cable or the data server but also by the number of tape drives in use. Programs today, like HP OpenView Storage Data Protector, VERITAS NetBackup, IBM TSM, or others, allow multiple parallel backup tasks to individual tape drives, getting the best use of the speed and the capacity of the magnetic tape drives. As tape drives can only be assigned to one server at a time, backups of the individual servers should be done at different times to optimize the magnetic tape drives.


Figure 3 provides a schematic overview of a backup infrastructure.

Figure 3.

LAN based Backup with Media-Server



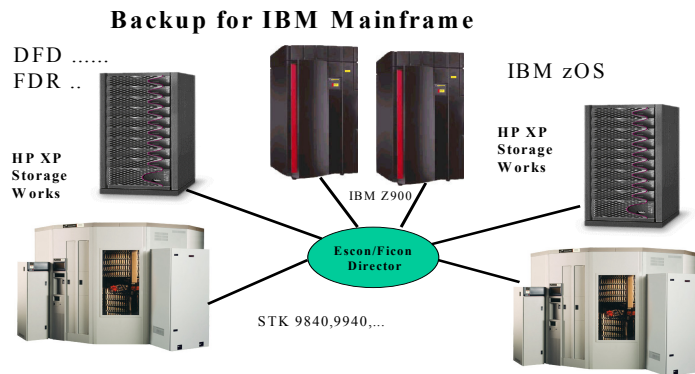
LAN free Backup



Backup in mainframe

Figure 4 shows an IBM environment under z/OS; backup procedures under BS2000 or OSD3 (other mainframe systems) have similar schemes. Figure 4 is a schematic overview of backup on a mainframe.

Figure 4.



The big difference is that ESCON or FICON-Director switches have been long available and permit true sharing of resources between mainframes. This makes it possible for any server on the system to make copies of any disk subsystem in the network, and means that the full bandwidth, allowed through the physical connectors, FICON- or ESCON-channels, can be used. Central control and administration of these processes are normal in this type of environment.

The basic problems described earlier regarding backup and recovery are the same in mainframe environments. The backing up of consistent data must be assured at the respective PIT. The biggest advantage in this environment is that in the long experience gained in this setting, there exists a well-thought-out organization as to how to best address these measures. And it is still possible to improve availability, improving the online availability for the user. New technologies such as duplicating within the same disk subsystem for faster recovery and restart are standard in the mainframe environment. Many users are still using long standard backup and recovery procedures that have worked in practice, but could be optimized with the newer and better technologies currently available.

Optimizing means that fewer resources are needed to reduce the backup and recovery times by optimizing the processes to realize savings in the whole RZ-environment.

But it remains that with the solutions described, mainframe or open system, the availability of the applications is reduced or not possible during the backup process. Reduced availability occurs when competing online copy procedures have started, and as a rule, will reduce the performance for the application. In case of an error, the whole backup process must be restarted. The rules for an online copy process are, in case of error, all data for backup will be reset and then the backups will start again with a new timestamp. The customer must check all data for consistency.

It should be noted that these requirements are necessary when data consistency is a factor. In many computer centers data may be saved but the data is not consistent. For example, in a Windows® file server backup, an incomplete backup has been made, where not all of the files were backed up at the appropriate PIT. A restore procedure may back-copy an open file that does not have an associated application and will result in data loss. This is also true for databanks. It is important that the interface of the applications is suspended during the time necessary for the procedures; the result is that temporarily interrupting the file system assures that a consistent copy is made from the file system, which can be back-copied onto the disk system. Forward recovery is application-specific and, as previously mentioned, must be written and organized in the backup procedures. Standard databank systems—Oracle, DB2, or IMS—have built-in functions that can be activated. These allow for a so-called freeze command (a logical PIT copy) in the databank system, which assures consistent data. It allows timestamps on the transaction logs, so that the data associated to a recovery point is well defined.

This leads to the following questions: How can HP best optimize these processes? How can HP improve the level of service and increase the value for the customer?

The simple answer—with the HP software suite for the HP StorageWorks XP array.

As an addendum, an overview of some of the individual functions for disaster and recovery that are used and maintained is presented; these apply to mainframe as well as to open systems. The software suite and the associated hardware functions offer the possibility of backing up any data from any computer at any time onto any location and being able to save or synchronize the data at any distance. This means that, according to the needs of the customer, HP products can back up any data from any computer at any time, onto one, two, or more locations while maintaining data consistency. To accomplish this, it is necessary to know the customer's needs and to carry out an analysis of the present status. Customers can discuss their long-term and short-term goals, and proposals can be worked out with customers to address their goals and requirements. With HP products, every customer need concerning data mirroring, backup-recovery procedures, and the number of necessary copies is addressed. HP also offers the necessary in-house consultation and can tailor the best solution to help each customer.

Schematic presentation

The following is not meant to be a thorough technical description of how the individual hardware and software functions are used, but descriptions of how they work.

Mirroring in a local subsystem

The basic idea of a local mirror copy is that it works like a mirror; every data movement on the primary page is mirrored (copied) onto a secondary page.

Available systems functions, which can be remotely controlled, make it possible to duplicate data from the primary page, or mirror the data. Multiple copies of this mirror copy can be held on the disk subsystem. The number of mirror copies dictates the physical size of the disk arrays.

It is also important to decide between two procedures, full duplication of a logical volume or updated changes made in the data of a logical volume.

With multiple mirroring it is therefore possible to have, at any time, a copy of the original available (only one BCV is active at one time).

What can be done with these local mirror copies?

For example, making a mirror copy every three hours over the day (8:00 a.m.–16:59 p.m.) results in four copies on the disk subsystems. At 8:00 a.m., the first mirror process is carried out (freeze); the freeze ensures data consistency at the PIT of the mirror copy. Three hours later another mirror copy is made, and the process repeats itself every three hours.

In the duplications made over a nine-hour period—three generations—there are three different but consistent copies of the data generated over this time.

The advantage is that if a logical error in the data occurs, there are three generations with three hours difference available. In most cases, as the time span between the mirroring events and error is short, a minimum time is needed to recover from an error.

It is obviously impractical to continue this schema over a 24-hour period, and one alternative is to limit the number of generations made. When this limit is reached, the oldest generation of data is written over. With this generational procedure it is possible to use any of the available previously mirrored copies. Errors that are older than the mirrored copies available would require a backup/restore from a taped backup, making it important to define the intervals when the mirrored backup copies are transferred to an external medium.

That primary data is protected with RAID techniques; the aim of locally available mirror copies is to improve data consistency at defined PITs. The mirror copy can also be transferred onto tape drives or evaluated without interrupting the online applications.

To meet the requirements of the Basel 2 standards, fast processing times are necessary. New measurements, such as those to prevent money laundering, make large demands on the IT system, as all transactions over \$10,000 must be documented and simultaneously available to all branches of the bank. The increased demands on the IT system have to be supported by the applications and hardware systems.

Mirroring at a second location

Mirroring onto a second location means copying the data from the primary page to a remote site using remote copy techniques with synchronous or asynchronous transfer. This first copy of the original saved locally can be copied again and transferred, making it possible to have various versions of the same data available. Mirroring on a remote location improves disaster recovery measures.

Consideration must be given if additional mirror copies are necessary than those available locally for adequate disaster recovery protection. Otherwise it must be accepted that in the event of a disaster, where data consistency would normally allow for a forward recovery, errors in the logical status of the data may exist.

After such a disaster, the consequence of such a configuration is that the recovery process is unable to use any of the locally saved mirror copies. This would certainly result in a further delay in restarting the system, but take into account that such disasters are rare and will hopefully never happen. These limitations and their consequences should be described in the disaster recovery plan.

In the mainframe environment one can assume that the applications and system programs are recovery-capable and have the ability, in the event of a disaster, of running off a remote site. The technical system prerequisites must still be adhered to. The current standard in the computer center is the Sysplex installation, which with GDPS makes it possible to automate many processes. It does not take the technical hardware changes into account (network, power, and so on). These actions must be described in the disaster recovery plan.

Many of these features are not supported under UNIX® or Microsoft® Windows systems. The application itself must be able to start from a secondary copy, and make the necessary changes to recognize it as the primary copy.

The following table illustrates which available mirroring techniques can be used. Increased application and system availability brings an increase in value for the customer.

HP StorageWorks Business Copy XP	Real-time local mirroring at volume level
Logical Volume Divider	Real-time local mirroring at volume and dataset level
Database Divider	Real-time local mirroring of DB2 databases
Hitachi FlashCopy	Real-time local mirroring and/or snapshots at volume and dataset level
HP StorageWorks Continuous Access XP	Real-time synchronous remote mirroring
HP StorageWorks Continuous Access Extended XP	Real-time asynchronous remote mirroring
HP StorageWorks Continuous Access Journal XP	Real-time asynchronous and multi-site remote mirroring
Hitachi Extended Remote Copy	Host-based asynchronous remote mirroring
Business Continuity Manager	Disaster Recovery Management Software

For more information

- HP StorageWorks Business Copy XP
http://www.hp.com/products1/storage/products/disk_arrays/xpstorgesw/business/index.html
- HP StorageWorks Continuous Access XP, XP Extension and XP Journal
http://www.hp.com/products1/storage/products/disk_arrays/xpstorgesw/continuousaccess/index.html
- HP OpenView Storage Data Protector
<http://h18006.www1.hp.com/products/storage/software/dataprotector/index.html>

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Oracle is a registered US trademark of Oracle Corporation, Redwood City, California. UNIX is a registered trademark of The Open Group.

4AA0-3228ENW, December 2005

