TECHNICAL WHITE PAPER

bmcsoftware
ACTIVATE BUSINESS WITH THE POWER OF I.T.™

Remote Recovery — Advanced Technology Solutions for z/OS Recovery

By Rick Weaver, Product Manager,
Mainframe Database Recovery, BMC Software

Table of Contents

Overview 1

Criteria for Advanced Technology 2
> Recovery point objective 2
> Recovery time objective 2
> Recovery geography objective 2
> Data integrity/synchronization 2
> Availability 3
> Performance 3
> Cost 3
> Risk 3

Advanced technology recovery 4
> Advanced technology recovery vendors and products 4
> Which technique is best for my installation? 4
Remote logging 5
Data Propagation 5
Remote disk shadowing 6
Remote disk mirroring 6

Summary 7

Overview

Reducing or eliminating IT recovery impact on business operations is an increasingly important goal for corporations today. The prospect of making databases and files unavailable for hours or days while recoveries occur is unacceptable. Companies are looking for solutions that will maximize availability of business-critical data in a global environment, yet still allow for consistent and rapid recovery of the critical data.

Business continuity for IT has two major areas of concern — local high availability, and remote site disaster recovery. Of the two areas, disaster recovery typically gets the most attention (and funding). Many vendors are interested in addressing this challenge.

The pros and cons of the solutions must be balanced against the corporate requirements. We can break the requirements down into the following criteria:

- Recovery point objective (data loss)
- Recovery time objective
- Recovery geography objective
- Data integrity/synchronization (consistency)
- Availability impact (local)
- Performance impact (local)
- Cost and benefit analysis
- Risk of solution failure

The relative importance of these criteria differs from company to company, and may be modified based on changing business needs.

The solutions available today include offerings from several vendors in the z/OS marketplace. The solutions fall into four broad categories:

- Remote vaulting or journaling
- Data propagation (software shadowing)
- Remote disk shadowing (asynchronous replication)
- Remote disk mirroring (synchronous replication)

Criteria for Advanced Technology

The criteria noted earlier need some further definition and clarification. Again, the sequence here does not imply prioritization. Each company will have to arrange the criteria to meet its needs, and some of the criteria may conflict. Most companies find that when they sequence the criteria and explore the impact, a new number one becomes apparent.

Recovery point objective

The objective of a backup is to allow for recovery in the event of a system failure. The backup strategy should incorporate application knowledge to ensure the business can continue processing against the recovered files and databases. After a restore, the business data looks as it did at the time of the backup. If that backup is a week old, it's as if the last week of processing never occurred. Even with a paper trail, or retransmission of an electronic business-to-business feed, it's unlikely the data will accurately reflect the current business status.

The recovery point objective (RPO) identifies the acceptable data loss in the event of a total site failure. It may be much more aggressive than the current backup strategy allows, and in fact an entire class of advanced technology products exists to address the data loss associated with recovery to an "old" backup.

Recovery time objective

When recovery is required operators are faced with gathering lots of media (typically tapes or cartridges) and processing them through restore or recovery jobs. In a database environment, there may be additional processing to apply any data changes logged since the last backup. There are procedures for shipping and storage (vaulting) of the tape assets for local and disaster recovery and presumably procedures for gathering and processing the media at the time of failure.

The recovery time objective (RTO) is the business goal for application resumption, measured from the time of the outage. It is not restricted to a disaster recovery discussion. There are local recovery RTOs as well. Some of the components of recovery time include:

- Problem analysis
- Recovery declaration
- Transporting/shipping/gathering assets (including personnel)
- Recovery processing (including database log apply)
- Application availability
- Reprocessing missed or lost transactions

In most production environments, the RTO for disaster recovery is stated in days, but in fact most companies would be hard pressed to resume business within a week after a site failure.

For local site recoveries, the RTO is typically stated in hours. Many companies have experienced extended outages due to lack of expertise or failure of resources and assets during recovery.

Recovery geography objective

To protect the corporate data asset from destruction, each company should perform some analysis to determine the main threat to the facility. For many companies, that threat is a wide-scale event such as a hurricane, catastrophic flood, or regional power failure. When disasters strike, it is important that the recovery facility be remote enough to protect it from the original disaster. Many companies have a second computer center, or they subscribe to a service that is hundreds or thousands of miles away from the production facility.

The recovery geography objective articulates the need to house a recovery facility an extended distance away from the production facility to lessen the likelihood of one disaster affecting both production and recovery facilities.

Data integrity/synchronization

When establishing a recovery environment — advanced technology-based or not — it is critical that all the pertinent files for an application be consistent after the recovery. For advanced technology recovery solutions, the requirement for a consistent recovery point is critical and hard to obtain. The fact that changes are being captured and transmitted while the transaction may still be executing can create some horrendous recovery headaches. It can be very difficult or even impossible to guarantee data integrity for the files and databases involved in the recovery.

A recovery from a fuzzy or dirty copy is probably worse than no recovery at all. Applications evolve over time, and develop relationships with other applications in which data is shared. If a file restore or database recovery corrupts the data, it affects more than just that application. It may affect several other applications, and the effects may be very subtle. Instead of getting a transaction failure with a message stating inconsistency, the problem may go undetected for months until a sharp-eyed customer notices a mistake in his records.

Availability

Depending on the nature of the business, the solutions discussed in this paper may seem outrageously expensive, or they may seem like a bargain. The key issue is the impact on your business and loss of availability of your application data. In a global production environment there is no nighttime batch window. Any time is prime time, and the "doors" are always open for business. Even a brief outage may be unacceptable.

Performance

Occasionally a solution is implemented to resolve one problem, only to cause another. Some of the advanced technology recovery solution may cause additional overhead, either to processor resources or to disk controllers. Generally, performance impact from advanced technology recovery processing is a major consideration in choosing a solution.

Cost

All of the advanced technology recovery solutions require additional resources, and some are very expensive. Customer costs entail the following components:

> Hardware (additional disk or cache)
> Software (e.g., EMC SRDF, IBM® XRC)
> Network (dedicated telecommunications links)
> Facilities (to house the remote assets)
> Operations (or services)

Even with cheaper disk, it's daunting to consider doubling (or tripling, or more) the production disk requirements to satisfy recovery requirements. The hardware disk-mirroring vendors have several sales points regarding the difference in their technology versus their competitors' technology. For now, consider the fact that just doubling the disk costs isn't the end of the cost discussion — even the disk-controller-based hardware techniques require some software (which is an extra charge) and will certainly require additional training for the operations staff. Some of the hardware components are very heavy, and you may have to reinforce the facility to accommodate them.

Because you'll be moving and mirroring large volumes of data, you'll want to ensure ample bandwidth. For many of the solutions that means ESCON connectivity. ESCON has a limited range (about 66 km, or 30 miles). Some customers demand more distance between their production and recovery sites, which results in the use of high-bandwidth connectivity.

Risk

One of the biggest risks of the advanced technology recovery solution family is the false sense of security it can foster. One might think the investment is that the technology completely safeguards the business from all failures. This is a fallacy. All of the solutions require massive data movement, and that process can fail. There is the risk of what is referred to as a "rolling disaster," which means that a disaster doesn't necessarily happen all at once — components may fail over a period of time, causing inconsistent data at the remote site. The link to the remote site could fail, rendering the remote-site version of the data obsolete. When a disaster occurs and the remote site is activated, some hardware configuration changes may be required. There are three types of disasters you should consider when determining the risk:

> Mega Disaster — This would be a total production site failure. Keep in mind it may not happen all at once, but rather in stages
> Mini Disaster — A link failure, a disk device failure or a software glitch. You probably won't invoke the disaster recovery process to resolve this type of failure
> Hack — A logical failure, caused by faulty programming or malicious data manipulation. You probably won't invoke the disaster recovery process to resolve this type of failure

No matter which of the following solutions you choose, you still risk data corruption and recovery failure at both the production and recovery sites. You still need valid restore and recovery procedures in place to rebuild your applications in case the advanced technology recovery solution you employ fails.

Advanced Technology Recovery

The technology available to address z/OS disaster recovery is mature, with some of the solutions dating back to the late 1980s. Some form of additional hardware is required for all of these solutions — additional disk, cache, or bandwidth. The solutions transmit data, either partial or complete, from the production site to a remote recovery site. Both sites must be available 24x7, with full operations and support requirements.

Advanced technology recovery vendors and products

The list of advanced technology recovery vendors is relatively stable. The vendors include IBM, EMC, E-Net, StorageTek (STK), and Hitachi Data Systems (HDS). Because this is a competitive market space, the products are continually enhanced with new features. This list of vendors and products is current as of Fall 2006.

Remote logging captures logged changes to the database and transmits to the remote site. The Remote Recovery Data Facility (RRDF) product from E-Net RRDF supports IMS, DB2, CICS, IDMS, ADABAS, and other database management systems.

Data propagation captures logged changes to the database and transmits to the remote site, and then applies the changes to a shadow database. The vendor products include:

> E-Net RRDF with LogApply for VSAM

> E-Net RRDF with LogApply for DB2 (up through DB2 V8 Compatibility mode)

> E-Net RRDF with ENET2 for IDMS

> E-Net RRDF coupled with high-speed recovery utilities such as RECOVERY PLUS for IMS from BMC Software

> E-Net Enterprise Data Replicator (DB2 V7 and V8 only at this time)

Remote disk shadowing captures changes to the local production disk and transmits to a remote site disk with some lag time. The vendors include:

> EMC SRDF/A (Symmetrix Remote Data Facility) Asynchronous mode

> EMC SRDF/AR (Automated Replication for Multi-Hop scenario)

> EMC SRDF/Star (combination of SRDF/S and SRDF/A)

> IBM Global Copy (previously known as PPRC Extended Distance)

> IBM Global Mirror (previously known as PPRC Asynchronous)

> IBM XRC (Extended Remote Copy) — emulated by Hitachi Data Systems (HDS) HXRC

> IBM GDPS (Geographically Dispersed Parallel SYSPLEX) on XRC

> HDS TrueCopy Asynchronous mode

> STK DataReplicator Asynchronous mode

Remote disk mirroring captures changes to the disk and immediately transmits to the remote site in the same work unit. Vendor products include:

> EMC SRDF/S (Symmetrix Remote Data Facility) in Synchronous mode

> IBM Metro Mirror (previously known as PPRC Synchronous). The other storage vendors (EMC, STK, and HDS) can emulate this mode

> IBM GDPS (Geographically Dispersed Parallel SYSPLEX) on PPRC

> HDS TrueCopy Synchronous mode

> STK DataReplicator Synchronous mode

> STK PPRC emulation and PowerPPRC

Which technique is best for my installation?

How do you decide which is the right choice? How can you differentiate between the various techniques? You must prioritize and balance the criteria against the benefits to your processing environment.

The relative pros and cons of each of the solutions follow. In some cases, you'll see a criterion as both a pro and a con because the criterion may have a favorable affect in one sense, but an unfavorable affect in another. (For example, in remote logging, performance at the production site is somewhat affected due to the extra processing required to support the solution, and that's a negative in some eyes. On the other hand, the solution is not distance-sensitive as many of the disk techniques are, and that's a benefit relative to them.)

Generally speaking, the vendors and customers have focused on asynchronous replication (remote logging, data propagation, remote disk shadowing) as valid techniques for disaster recovery. These techniques balance the local performance impact with the remote site network and distance costs. Most of the recent technology advances for these techniques have been reduced data loss while maintaining the remote site data consistency in the event of a rolling disaster (a disaster that initially affects part of the environment, but then cascades to involve the entire production site).

Remote logging

The following table lists the advantages of remote logging.

Recovery point (for databases)	Data loss is minimal (probably just in-flight transactions).
Availability impact	None, other than normal backup/copy processing.
Risk	This solution handles the three types of disasters well. **Mega disasters** (total center outage) are what all of these products are built to support. This solution gives you visibility to the last in-flight transactions, so you can recover or back out any units of work at the remote site and resume business operations quickly. In the event of a **mini disaster** like a link failure, there are buffers on the production site to queue the changes. A gap analysis component ensures that the remote-site log is complete. For **hack** (logical) failures, some type of process at the production site would result in correcting the data. Because those changes would be captured and transmitted to the remote site, it would also reflect the corrected data.
Data integrity (database to database)	Because all logging for related applications shares the same task, the records are streamed across the network in sequence. Your recovery point at the remote site would be the equivalent of a local power failure — at startup, some databases would be forward-recovered, some would be backed out, and all would be consistent.
Performance	This solution is not distance sensitive. The remote site can be thousands of miles away with little or no additional impact on the production applications.
Cost	This technique is probably the least expensive to implement when considering all of the cost components. It is especially frugal with network bandwidth and remote site storage.

The following table lists the disadvantages of remote logging.

Data integrity (database to sequential file)	The solution doesn't address sequential files. There is no way to synchronize them with the database, so they will have to be rebuilt at the remote site. There is also no support for Distributed Systems databases — it is strictly a mainframe to mainframe solution.
Recovery time	More log data means longer recovery. You're still recovering databases from copies taken hours or days before the disaster and applying all logged changes.
Performance	There are tasks running on both processors, which can add overhead. The more database activity there is the more overhead is produced.
Cost	Hardware — some processor, disk and tape/cart at the remote site (but not large volumes). Software — RRDF licenses on both processors. Bandwidth — a couple of T1 (1.5 MBS) is enough, but it depends on log traffic. Facilities — can be conscripted by the disaster recovery hot site providers (e.g., Sungard or IBM BRS). Operations — remote site requires 24x7 operations staff; it's now an extension of production.

Data Propagation

The following table lists the advantages of data propagation.

Recovery point (for databases)	Data loss is minimal, probably more than just in-flight transactions, but less than going back to last night's backup.
Recovery time	It takes a few hours to get back the last changes since the last apply process. E-Net claims EDR can support near instantaneous cut over, resulting in very little downtime.
Performance	It is not distance sensitive, the remote site can be thousands of miles away.
Availability impact	None (other than normal backup/copy processing).

The following table lists the disadvantages of data propagation.

Data integrity (database to sequential file)	The solution doesn't address sequential files. There is no way to synchronize them with the database, and they must be rebuilt at the remote site. Distributed systems data is not supported.
Performance	There are tasks running on both processors, which can add overhead. More database activity produces more overhead.
Risk	The solution handles the three types of disasters well. You still need a process to fix the logical errors at production, but the logical recovery will be propagated to the remote site and reflected there.
Cost	Hardware — more processor, lots more disk (double or more), and tape/cart at the remote site. Software — RRDF + Log Apply (or EDR) licenses on both processors. Bandwidth — a couple of T1 (1.5 MBS) is enough, but it depends on log traffic. Facilities — can be conscripted by the disaster recovery hot site providers (e.g., Sungard or IBM BRS). Operations — remote site requires 24x7 operations staff. It's now an extension of production.

Remote disk shadowing

The following table lists the advantages of remote disk shadowing.

Recovery point	Data loss is minimal — probably just in-flight transactions or a few I/O operations.
Recovery time	No recovery is needed, but the remote disk must be attached to a production-sized processor before applications can process the data.
Data integrity	These solutions address sequential files and databases, and some can address mainframe and distributed system data at the same time. The vendors have invested a lot of resources into resolving data consistency at the remote site in the event of total or rolling disasters.
Availability impact	None, other than normal backup/copy processing.

The following table lists the disadvantages of remote disk shadowing.

Performance	Extra process resources are required to support data sequencing and consistency.
Risk	If you can't guarantee data integrity, why bother with the solution? It's increasingly difficult to re-create the transaction to fix the data once it has been restored inconsistently. Even if you have a technique implemented that has tolerable advantages for the needed disadvantages, you still need a viable backup and recovery plan, both for local production and for remote-site production in case the advanced technology recovery solution is unavailable at the time of disaster. Disasters don't always go as planned, and you need a fallback support position.
Cost	Hardware — more processor for some solutions, lots more disks (double, triple, or more). Software — SRDF on EMC, XRC and SDM (System Data Mover) for IBM, etc. Bandwidth — Many T3s; one customer had 18 T3s for shadowing a few terabytes of data. Facilities — conditioned floor space, can be conscripted by the disaster recovery hot site providers. Operations — remote site requires 24x7 operations staff. It's now an extension of production.

Remote disk mirroring

The following table lists the advantages of remote disk mirroring.

Recovery point	No data loss.
Recovery time	Zero, but must attach the disk to a processor.
Data integrity	Ensured for all data on the mirrored volumes. Disk mirroring technology can support mainframe and distributed system databases and files.
Availability impact	None, other than normal backup/copy processing.

The following table lists the disadvantages of remote disk mirroring.

Performance	Synchronous disk mirroring is extremely distance sensitive. It really must be done with ESCON range (66 km, 30 miles).
Risk	30 miles doesn't get outside the impact zone of many common disasters. Why put all the time and effort into a recovery solution that may go up in smoke (or under water) with the site it's protecting?
Cost	Hardware — lots more disk (double, triple, or more). Software — e.g. EMC SRDF/S or IBM Metro Mirror. Bandwidth — a lot — probably OC3 or more. Facilities — conditioned floor space, within 63 km. Operations — remote site requires 24x7 operations staff. It's now an extension of production.

Summary

This paper has attempted to identify the major vendors and product offerings that address disaster recovery. Because the difference between some of the solutions is slight, consider the price and vendor relationship. After perusing this paper, you should have the basic understanding required to make an informed decision.

The disaster recovery solutions documented here solve one problem — total site failure due to disaster. For most events, the impact is local and affects a small number of applications and databases. A disaster would not be declared in the event of a volume failure or a bad program update. Because the solutions cited here do not protect from those sorts of failures, you still need valid backup and recovery processes and products to ensure business continuity.

ACTIVATE BUSINESS WITH THE POWER OF I.T.™

About BMC Software

BMC Software delivers the solutions IT needs to increase business value through better management of technology and IT processes. Our industry-leading Business Service Management solutions help you reduce cost, lower risk of business disruption, and benefit from an IT infrastructure built to support business growth and flexibility. Only BMC provides best practice IT processes, automated technology management, and award-winning BMC® Atrium™ technologies that offer a shared view into how IT services support business priorities. Known for enterprise solutions that span mainframe, distributed systems, and end-user devices, BMC also delivers solutions that address the unique challenges of the midsized business. Founded in 1980, BMC has offices worldwide and fiscal 2006 revenues of more than $1.49 billion. Activate your business with the power of IT. For more information, visit www.bmc.com.

67315