



Disaster recovery solutions for System p servers
and AIX 5L
February 2006

Alan M. Wilcox
IBM Corporation

Contents

Abstract	3
Introduction	5
Disaster Recovery Model	6
Tivoli Storage Manager on System p Disaster Recovery real time solutions	
Disaster Recovery – Software Products Available on System p ...	8
Oracle and DB2 HACMP and HACMP/XD	
Summary	11
Appendix	12
A. Disaster Recovery Planning Model (expanded) B. HACMP and HACMP/XD Mirroring and Data replication options C. Table of Networking Technologies and Data transfer rates	
References	17

Abstract

This paper provides the reader with the basic guidelines on how to identify, discuss, and qualify Disaster Recovery opportunities for IBM clients running AIX 5L™ applications on System p™ servers. The discussion that follows focuses on the key elements of a complete disaster recovery plan for a business enterprise, and the contributing roles and features of System p software products. The following table summarizes System p solutions with the potential to address two main types of outages that need to be considered in Disaster Recovery planning:

- Complete loss of the primary data center site, caused by one or more catastrophic events. The site will be unavailable for an extended period of time.
- Loss of some critical resource in the data center for a long enough duration to impact ‘top tier’ business-critical applications, and therefore, the business as a whole. The data center itself remains physically intact, but due to a component failure, or human error, critical capacity is lost, or access to the enterprise network is not available.

Scenario	Option	System p solution
<p>Primary site of operations will be unavailable for an extended time, or permanently unavailable.</p>	<p>Transfer all operations / applications to an alternate site as soon as possible.</p> <p>Applications which can tolerate longer outage times will be restarted via restore from backup tapes.</p> <p>Transfer business-critical [time sensitive] application processing to an alternate site. This must be done in significantly less time, and with significantly less data loss than could be achieved by restoring data from backup tapes</p> <p>When the primary site is again available, transfer critical application processing back: Reestablish the primary site as the dominant site</p>	<p>For example : IBM Business Continuity and Recovery Services (BCRS)</p> <p>Tivoli® Storage Mgr (TSM)</p> <p>HACMP™/XD</p> <p>HACMP/XD</p>
<p>A temporary outage at the primary site has occurred. The site is physically intact. Some key resources needed for all data center operations are unavailable. (EG: loss of power, or network for 4 to 24 hours)</p>	<p>Applications which can tolerate longer outage times will likely be back online at the primary site in less time than it would take to restart them at an alternate location by restoring data from backups.</p> <p>Transfer business-critical application processing to an alternate site. This must be done in significantly less time, and with significantly less data loss than could be</p>	<p>Resolve temporary problem at primary site</p> <p>HACMP/XD</p>

	<p>achieved by restoring data from backup tapes</p> <p>When the primary site is again available, transfer critical application processing back: Reestablish the primary site as the dominant site</p>	HACMP/XD
--	---	----------

This chart suggests that *differentiation* of business applications which will severely impact business operations, even if they are down for a limited amount of time, is a key early step in putting together a comprehensive Disaster Recovery plan.

The chart really implies TWO major issues that must be considered when formulating a comprehensive Disaster Recovery plan for a client enterprise.

- 1) The DR strategy is developed based on the business requirements for recovery. The chart speaks directly to this. Normally, the client exercises full ownership of this.
- 2) WHO/WHAT organization will be responsible for implementation of the DR plan, as well as its day to day management?

This author has seen a clear separation of clients along the following lines:

- A) Clients who already have, or are planning hot site operations, which they themselves will staff and manage.
- B) Clients who wish to outsource the time and material cost of implementing and managing a DR solution to an experienced DR vendor, such as IBM BCRS.

Obviously, different clients will choose either A or B, based on their particular financial criteria that (they believe) will lead to the best 'fit', most cost effective method for implementing their DR strategy. A case in point, from actual experience, is presented in the Summary section on page 11.

Introduction

Disaster Recovery refers to the practice of providing an alternative, 'backup', or 'secondary' site from which to resume the business I/T operations, until the primary site can be put back in service. The 'Disaster', in 'Disaster Recovery', refers to the loss of business I/T processing operations, due to the occurrence of any catastrophic event which renders the primary, or normal processing site unavailable. 'Catastrophic events' are either natural or man made occurrences, such as tornadoes, fires, floods, or terrorist acts.

Disaster Recovery options can be visualized in a pyramid model. The least expensive, and most time consuming DR solutions would be found at the bottom levels of the pyramid. At these levels, 100% recovery of data lost in a disaster at the primary site is not possible. As you proceed to the top, each successive "tier" of the pyramid requires a greater investment in software, communications facilities, transportation equipment and System p server and storage hardware. The return for this higher investment in services and equipment is the potential of a significantly shorter recovery time, with minimal loss of data, compared to the preceding tier. The top of our pyramid model would be reserved for those mission-critical applications with effectively zero data loss, as well as very rapid restoration of operations, following a disaster.

Disaster Recovery Model

In this paper, we'll present such a model for Disaster Recovery, and then adapt it to the System p/AIX 5L platform. Here is a brief summary of the model, as presented by IBM at Share in 1992. We'll add more details to this view later, after we describe the Disaster Recovery features of key some software products that run on the System p/AIX 5L platform.

Figure I. The Disaster Recovery Planning Model

(Presented by IBM, Share Users Conference 1992).

Tier 6 – Zero data loss-
Recovery Time in minutes.
100% data recovery, minus data in transit at the time of the disaster.

Tier 5 – Two site two phase commit –
Recovery time from minutes to several hours. 100% data recovery
not possible if transmission is limited to inactive data (log shipping).

Tier 4 – Continuous Electronic vaulting between active sites
Recovery time from hours to days.
100% data recovery is not possible, even with online backups.

Tier 3 – Electronic vaulting of some backup data
Recovery time from hours to days.
100% data recovery is not possible

Tier 2 – Offsite vaulting with a hot site
Recovery Time in days or weeks
100% data recovery is not possible

Tier 1 – Offsite vaulting of backup data, by manual process (courier)
Recovery time in days or weeks
100% data recovery is not possible

Tier 0 – No Disaster Recovery protection – no off site data
There is no recovery from a primary site disaster

The lower levels of the model deal with recovery by restoring data from backup tape volumes that have been captured while the primary site is still in operation. Over time, these volumes are catalogued or archived, in order to be made available for restore operations at the secondary site, so that business operations could resume there within some 'recovery time', in hours or days. 'Recovery time' would have been previously negotiated in to a Service Level agreement between the data center operators and the end users. If the secondary site is a 'cold', or even a 'warm' site, additional equipment and staff will have to be provisioned to the site as part of the DR recovery of the business. Alternatively, a 'hot site' provides a 'ready to go' physical infrastructure and staff. Hot sites are an ideal solution when the client already owns, or is planning, a second datacenter. Additional activity, such as data restores, may be needed to bring up the business operations at the hot site.

Our model will make a distinction for the 'Tiers' of Disaster Recovery for which a 'hot site' is required. Other distinctions between the tiers of the Model are how quickly the client needs to recover the data, the 'Recovery Time Objective' (RTO), how quickly they need to recover the services provided by their environment, and how much data they cannot afford to lose, aka the 'Recovery Point Objective' (RPO). DR solutions are cost effective solutions which balance the client investment required to provide a given recovery level, or 'Tier' of the pyramid, against the cost to the business in lost revenue. The lost revenue

cost is estimated based on outage time, type of application, and Industry. Table I provides an ‘Average Cost of down time chart by Industry’, based on 2002 dollars.

Table I. Average cost of down time for selected US industries in millions of dollars (2002)

Industry	Average down time cost per hour
Brokerage Operations	\$6.5 M
Energy	\$2.8 M
Credit Card	\$2.6 M
Telecommunications	\$2 M
Manufacturing	\$1.6 M
Finance/Banking	\$1.4 M
Information Technology	\$1.3 M
Retail	\$1.1 M
Pharmaceuticals	\$1 M

IBM Tivoli Storage Manager on System p

IBM Tivoli Storage Manager (TSM) is used by many clients on the System p platform. It provides a number of data management features, for backup and recovery operations that can form the building blocks of a Disaster Recovery Plan. Among its many impressive features is the ability [of a single TSM server] to manage multiple copies of the same backup data within different ‘storage pools’ – groups of disk or tape storage volumes that can be maintained in multiple locations simultaneously, via the use of TCP/IP connections. Electronic vaulting procedures can be designed with TSM, so that the backup data is automatically and continuously provided to the secondary site. Multiple TSM servers at different sites can also be coupled together. This enables a recovery site TSM server to be immediately available, to commence restoration operations. Restarting business operations at a secondary site by restoring data from backup tapes could be expected to take days, or even weeks before all the applications are 100% available again. This author has worked with many System p/TSM clients who have this type of Disaster Recovery plan in place.

Clients in this category will not need extensive briefing on the DR Model; however, these clients may still be a candidate for a real time business recovery solution.

This type of System p client could be your first candidate profile for the following questions:
What if restoring data from backup tapes takes longer than the agreed time limit for the application’s recovery and restart? How much real time data is potentially lost between the time of the last backup and the occurrence of a disaster? What’s the business impact of this?

Disaster Recovery – real time solutions

Real time Disaster Recovery solutions focus on *transaction level* recovery. They sit at the upper two tiers of our DR model, and are intended for those critical business applications for which an outage of more than a few hours would have a severe impact on business revenue.

Real time Disaster Recovery solutions strive for recovery of nearly 100% of the live transaction data when the primary site is rendered unavailable. In addition, they are designed with the intent of restarting business operations at the secondary site as soon as possible, normally, in a few minutes. This does not include any time required for the application, once restarted at the ‘hot site’, to perform transaction level recovery or roll back operations.

To design such a solution, the application data must be “mirrored”. Data mirroring is a generic term that refers to the replication of data from an active disk volume at the application’s primary site, to a receiving disk volume at the secondary site, using a communications protocol and communication mode.

Synchronous transmission modes insure data equivalence, but impose a network performance penalty on the primary site application, which may affect local application response time, as well as limit the distance between the sites. Asynchronous transmission mode allows the data replication at the secondary site to be 'decoupled', so that primary site application response time is not impacted. Asynchronous transmission is commonly selected, with the exposure that the secondary site's version of the data may be out of sync with the primary site by a few minutes or more. This 'lag' represents data that would be unrecoverable in the event of a disaster at the primary site.

- The communications protocol, mode, and transmission medium of choice are dependent on
- (a) The proximity of the client's primary, and secondary/ 'hot site' facilities; and
 - (b) How closely synchronized the local and remote copies of the data need to be, depending on whether or not 100% of the data needs to be recovered;
 - (c) Data storage and network communications facilities the client may already be committed to.

TABLE II. Facility locations and Communications protocol options

Facility locations	Communications protocol options
Campus Wide – Resources in multiple buildings in the same local area network	- Local area network owned and managed by the client
Metro Wide – Separate data centers within the same metro area - local network provider	<ul style="list-style-type: none"> - Metro Area network (MAN) using TCP/IP - Storage Area Network using IBM ESS PPRC or EMC SRDF Transmission mode may be synchronous or asynchronous. - MAN or SAN access must be leased from a local TC provider
Unlimited – data centers in different geographies; State or country.- Network access must be leased from a Telecommunications provider	<ul style="list-style-type: none"> - Wide area network connection (WAN) using TCP/IP - Storage are network using IBM ESS PPRC or EMC SRDF Transmission mode will be asynchronous, or a synchronous method which can overlap multiple writes to the secondary site.. - WAN or SAN access must be leased from a global TC provider.

Facility locations and telecommunications solutions are among the first parameters that need to be considered and finalized as part of building the DR plan. Communication path redundancy, bandwidth requirements, and application data volumes to be replicated are all critical factors that can have a large impact on the networking component of the DR solution cost. In Appendix C, data transfer times for payloads varying in size from 10 gigabytes to 100 terabytes are compared using a representative range of communications protocols and network bandwidths.

Disaster Recovery – Software Products Available on System p

What kinds of software products are available on System p servers to manage these solutions? How do they support the communications protocols and modes we have been discussing?

Oracle and DB2

Oracle and DB2® are relational database products that provide database ‘point in time’ recovery. Log buffers are kept in memory for current transactions, and flushed to disk based on configuration parameters. This gives Oracle or DB2 based applications the potential capability to recover the environment to any previous point in time.

Oracle and DB2, both support ‘log shipping’. This is the transmission of closed and archived database log files to a database at the secondary site location which is in standby mode. The standby database is only available to accept and apply the archived logs from the primary site; it is not available for active user queries. Detection of a primary site failure is not provided by the software at the time of this writing. Site failover detection, and the resulting decision, is a manual process. Once the decision has been made to restart operations at the secondary site, the standby database can become an active database in seconds. This approach does not capture the current active [DB2 or Oracle] log buffers. This means active transaction data is at risk, and recovery up to the last transaction is not possible.

To minimize the loss of current log transaction data, users of DB2 version V 8 (Fix Pack 7+) may elect the ‘High Availability and Disaster Recovery’ (HADR) feature. This feature, announced in 4Q2004, enables an active DB2 database to replicate its current log buffers to a secondary site using one of the communication options we have discussed previously in Table II. This feature could be combined with log shipping to provide transaction level recovery at the secondary site, for DB2 applications. ‘Heart beat’ (for failure detection and recovery of the IP link) between the sites is not provided, so the decision to transfer operations to the secondary site is a manual one.

HACMP and HACMP/XD

HACMP is IBM’s premier high-availability software solution for providing server redundancy and enhancing application availability between System p servers that are ‘local’ to one another. Within a campus environment, data can be mirrored among member nodes of a HACMP cluster, using an IBM ESS PPRC SAN connection, or AIX 5L LVM split site mirroring. HACMP provides ‘heart beat’ between the member nodes, as well as monitoring of the applications. This provides automated application fall over and restart capability in the event of an unplanned outage on the primary node.

Where metro area, or unlimited distance solutions are needed, HACMP Extended Distance (HACMP/XD), a product feature of HACMP, supports geographic mirroring¹ over TCP/IP MAN or WAN connections, and also supports SAN data mirroring over IBM ESS PPRC connections. HACMP/XD provides automated site failure detection, automated application fall over and restart, and, once the primary site is restored,

automated fallback capability. HACMP/XD can work with any AIX 5L journaled filesystem or raw logical volume data, belonging to any application or database that is supported on the System p/AIX 5L platform. A recently announced feature of HACMP/XD, the Geographic Logical Volume Manager (GLVM), enables HACMP/XD to manage the local and remote copies of the data as AIX 5L LVM mirror copies, in volume groups that span across the sites. A summary of HACMP/XD data replication options is provided in Appendix B, on page 14.

This type of solution is transparent to the underlying type of data and application, provides heart beat and application monitoring between sites, and is configurable to provide automated fall over and fall back capability, which makes HACMP/XD a tier 6 [top] ranking in our DR model for System p.

1

Geographic mirroring over IP¹ refers to the mirroring methods that originated in the HAGEO product offering. They continue to be offered with the GLVM option of HACMP/XD and are independent of the underlying disk storage.

Summary

The following table is a summary of DR solution options we have discussed. For each tier of the pyramid, potential recovery time and recovery point, backup site requirements, and suggested System p solutions are presented. HACMP and HACMP/XD data replication options are described in more detail in Appendix B on page 15.

Table III. DR solutions summary for System p servers based on IBM DR Planning model. (Fig I)

DR planning model reference	pSeries solution	100% recovery of application data is possible ?	Automatic detection of site failure ?	Facility locations supported (Table II)	Communications modes/ protocols (Table II)
Tier 6 – Zero data loss. Recovery up to application restart in minutes.	HACMP with data replication option: - split site LVM mirroring HACMP/XD with data replication options: - HAGEO mirroring - GLVM mirroring - PPRC mirroring	Yes – minus data in transit at time of disaster	Yes. Failover and restart of applications, and failback are automated	All – Consider standard HACMP for campus solution	MAN or WAN – All SAN – IBM ESS PPRC
Tier 5 – Two site two phase commit. Recovery time varies from minutes to hours.	Oracle or DB2 log shipping to a remote standby database	No – does not include active log data	No	All	All
	Oracle or DB2 active data replication to a remote database	Yes	No	All	All
	DB2 log shipping + DB2 HADR	Yes	No	All	All
Tier 4 – Continuous electronic vaulting of backup data between active sites. Active data management at each site is provided.	TSM with copy pool duplexing between sites, and active TSM servers at each site.	No Recovery in days or weeks. Must restore from backups.	No	All	All
Tier 3 – Electronic vaulting of critical backup data to a hot site. The hot site is not activated until a disaster occurs.	TSM with copy pool duplexing to the hot site. TSM server at active site only.	No Recovery in days or weeks	No	All	All
Tier 2 – Off site vaulting with a hot site. Backup data is transported to the hot site manually. The hot site is staffed and equipped but not active, until a disaster occurs.	TSM with multiple local storage pools on disk and tape at active site.	No Recovery in days or weeks	No	N/A	N/A
Tier 1 – Offsite vaulting of backup data by courier. A third party vendor collects the	TSM with multiple local storage pools on disk and tape at active site.	No Recovery in days or weeks	No	N/A	N/A

data at regular intervals and stores it in their facility. When a disaster occurs : a) a hot site must be prepared b) backup data must be transported					
Tier 0 – No Disaster Recovery plan or protection.	Local backup solution may be in place, but no offsite data storage	No DR recovery – site and data are lost	N/A	N/A	N/A

As alluded to in the Abstract to this paper, on pages 3 and 4, overall cost and day to day maintenance of the chosen DR solution may result in the client's selection of a less expensive DR plan; as long as it meets the requirements of the business recovery strategy they have chosen. This author has very recently worked with a pSeries client with a significant amount of TSM, and single site HACMP experience. Implementation of their DR strategy was outsourced to IBM BCRS. To fulfill the client's requirements, BCRS has chosen to deploy IBM Enterprise Storage Server™ features, Peer to Peer Remote copy, for replication between the primary and DR sites, and local FlashCopy® on to production replica volumes at the DR site. This approach is not expected to promote as rapid of a recovery time as would an HACMP/XD solution. The local FlashCopy at the DR site must be performed at fixed intervals, and requires an interruption of service, albeit a brief one. In the event of an outage at the primary site, this DR environment will only be as current as the most recent flash operation, which may have occurred an hour or more earlier. Yet, BCRS chose against deployment of HACMP/XD in this specific case, upon agreement from the client that their recovery criteria could be met using a combination of these ESS features and locally written scripts, at [what was expected to be] a lower implementation cost.

Appendix

A. The Disaster Recovery Planning Model (Expanded)

(Presented by IBM, Share Users Conference 1992).

Tier 6 – Zero data loss*

Recovery Time in minutes.

*** 100% data recovery, minus data in transit at the time of the disaster.**

Encompasses zero loss of data, up to the current transaction's in flight data, and a nearly immediate, automatic transfer to the secondary platform. Data is considered lost if a transaction has commenced (for example, a user hits the Enter key to initiate an update), but the request has not been satisfied. Tier 6 is the ultimate level of Disaster Recovery. Local and remote copies of all data are kept updated via duplicated online storage at each site.

Tier 6 applications will require the highest availability, and fastest recovery time of tiers 0 - 6. Typically, there is more than one proprietary data base in the transaction path, so the capability to manage different kinds of application data between the sites is required .

Tier 5 – Two site two phase commit –

Recovery time from minutes to several hours. 100% data recovery not possible if transmission is limited to inactive data (log shipping).

This tier encompasses all the requirements of tier 4, plus the “mirroring” of selective data between applications at both sites. This is the first tier at which application data is transmitted to the hot site. The applications are synchronized between sites, for the type of data being transmitted.

Practical examples include data replication between Oracle or DB2 databases, or database ‘log shipping’. Log shipping refers to the transmission of closed database logs from the primary site database to a standby database at the secondary site where they are applied.

In standby mode, the hot site database is not available for user activity. Once the standby database is activated, it will contain all the transaction activity of the primary site database, minus the activity captured in the current active database log files.

Tier 4 – Continuous Electronic vaulting between active sites –

Recovery time from hours to days.

100% data recovery is not possible, even with online backups.

Encompasses Tiers 0,1,2 and 3. In addition, critical applications are running at both sites.

This tier introduces the requirement for active data management at the hot site. With an active TSM server at each site, bi directional recovery is now possible.

Each site acts as an electronic repository for the other's critical data. High bandwidth network connections are used to transfer backup data between sites, and store it in (the partner site's) tape libraries, managed by local servers running TSM, or an equivalent backup and recovery product.

Either site may be configurable to assume the part, or all of the work load of the other site, but this would require significant pre planning, testing, and manual intervention. An equal amount of planning and preparation would be expected in order to restart operations at the original site, following its reconstruction.

Tier 3 – Electronic vaulting * of backup data

Recovery time from hours to days.

100% data recovery is not possible

Encompasses Tiers 0 + 1 + 2, plus adds "electronic vaulting" of a subset of critical data from the primary site to a secure site. The secure site may be operated by a third party. There may also be electronic vaulting of data from the primary site to a hot site, but active Data management at the hot site, such as an active TSM server, is not provided.

* "Electronic Vaulting" provides electronic transmission of backup data over a network from the primary site, directly to the secure site. A TSM server at the primary site maintains nearly equivalent content of the production data through concurrent management of local and remote TSM copy storage pools located at each of the sites..

A practical recovery design at Tier 3 might call for recovery of all business critical data, and the most critical applications, through the restore of the TSM managed data that is already at the hot site. In other words, get the most critical applications back on line at the hot site while some less critical back end systems (for example, analysis functions) are still being recovered from TSM backup data that was retrieved from the offsite vault via PTAM.

Since the client will want to justify the additional investment in "electronic vaulting", recovery time for business critical applications at this tier should be within two to three days.

Tier 2 – Offsite vaulting with a hot site

Recovery Time in days or weeks

100% data recovery is not possible

Encompasses Tier 0 + 1, plus adds provision of a "hot site" at which to perform recovery.

In this tier, the DR plan is expanded to include an alternate facility which provides sufficient power, network capacity, and replacement computers and disk storage to process the work loads of the primary site, once the data has been restored to an "RPO" or "Recovery Point Objective". If a disaster occurs, the backup data must be first brought to the hot site using the courier services.

Tier 1 – Offsite vaulting of backup data, by manual process (courier)

Recovery time in days or weeks

100% data recovery is not possible

Encompasses Tier 0 plus manual transport of backup tapes to an offsite vault

There is a Disaster recovery plan, whose primary component is inclusion of offsite vaulting. This involves regular transport of backup data (tapes) to a secure facility, often run by a third party. The transport mechanism is by courier (truck or van), and hence is referred to as the "Pickup Truck Access Method" (PTAM). Some recovery requirements have been identified, but a backup site at which to recover and resume operations, and the requisite hardware, are not provisioned unless a disaster occurs.

Tier 0 – No Disaster Recovery protection – no off site data

There is no recovery from a primary site disaster

No DR plan - critical and non critical data is backed up and stored locally.

Estimated recovery time - Days or weeks - Including time to reconstruct facilities, order, receive, install and configure replacement equipment, restore data, etc.

B. HACMP and HACMP/XD data replication options

The base HACMP licensed program product supports split site LVM mirroring.

- LVM split site mirroring:
 - Actual disks from local and remote storage systems are visible to two servers on separate sites via SAN or similar technology
 - VG's, LV's are created as normal and mirroring is performed via AIX 5L from the local site to the remote site.
 - HACMP has tools that verify mirrors of LV's are on separate disks

HACMP/XD provides three different technologies in support of data replication to the DR site:

- PPRC
- GLVM
- HAGEO.

Each of these methods uses a very different underlying technology to accomplish replication, although they provide availability in much the same way:

- The PPRC alternative provides synchronous replication and uses the IBM Enterprise Storage Server (ESS), or DS8000 PPRC function
- GLVM provides synchronous replication and extends the AIX Logical Volume Manager to remote site support, with the concept of the RPV - remote physical volume:
 - The only connection needed between local and remote sites is an IP network (no SAN)
 - GLVM tools create a virtual disk (RPV) on the primary site server that maps to a real disk, LVs and VG, at the remote site server location
 - The GLVM device driver takes care of sending the data from the primary site HACMP/XD node where the IO is occurring, to the remote location, actual drive, and LV on the remote HACMP/XD server.
 - VG's, LV's are created as normal and mirroring of data to the remote site is performed via the primary site server's GLVM device driver by transferring data through the IP network
 - HACMP has tools that verify mirrors of LV's are on separate disks
 - The primary site HACMP/XD server's AIX 5L LVM has hooks that let it know that the virtual drives are virtual (RPVs), so that it reads from the local mirror. For write activity, both the local disk device and the corresponding RPV device (mirror) are written to.
 - At this point in time, GLVM only writes in synchronous mode.
- HAGEO provides synchronous, asynchronous and MWC replication and operates outside of the LVM using special kernel code, and UDP or TCP protocols to keep a copy of the data at the remote site. MWC, or Mirror Write Consistency replication is similar to synchronous replication, except that MWC allows the overlap of multiple writes to the remote site, improving performance.

C. Table of Networking Technologies and Data Transfer rates

(Data Xfer times assume that the network operates at 70% of its theoretical capacity)

Network Technology	protocol	Bandwidth gigabits/sec	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)
Local Area Networks (LAN)s			10 Gigabytes	100 GB	1 Terabyte	10 TB	100TB
10 Mbps Ethernet	TCP/IP	.01	190 Minutes	1900 minutes	19000 minutes	190000 min	1900000 min
100 Mbps Ethernet	TCP/IP	.1	19	190	1900	19000	190000
Gigabit Ethernet	TCP/IP	1	1.90	19	190	1900	19000
Stg Area Network	FCP –SCSI3	2	.85	8.5	85	850	8500
Stg Area Network	FCP –SCSI3	10	.19	1.9	19	190	1900
Network Technology	protocol	Bandwidth gigabits/sec	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)	Data Xfer Time (min)
Wide Area Networks/ Optical Networks			10 Gigabytes	100 GB	1 Terabyte	10 TB	100TB
T1	TCP/IP	.0015	1270 Minutes	12700 minutes	127000 min	1270000 min	12700000 min
T3	TCP/IP	.0447	43	430	4300	43000	430000
OC3/STS3	ATM	.1552	12.3	123	1230	12300	123000
OC192	ATM	10	.19	1.9	19	190	1900
SAN+DWDM	FCP+DWDM	200	.014	.14	1.4	14	96

References:

(1) “Disaster Recovery Strategies with Tivoli Storage Management” [IBM Redbook] SG24-6844. November 2002

(2) “Geomirror performance for HAGEO and GEORM: A Commentary”
IBM White Paper by John Easton and Simon Marchese. Copyright 1999.
<http://www.ibm.com/servers/eserver/pseries/software/whitepapers/gmdperf.html>

(3) HACMP XD for Geographic LVM: Planning and Administration v5.2 SA23-1338

(4) IBM High Availability Cluster Multiprocessing for AIX 5L V5.3 (HACMP V5.3)
<http://www.ibm.com/systems/p/ha/>



© IBM Corporation 2006

IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
February 2006
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the e-business logo, @server, AIX 5L, HACMP, System p, Tivoli are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at:
<http://www.ibm.com/legal/copytrade.shtml>.

Other company, product, and service names may be trademarks or service marks of others.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

When referring to storage capacity, 1TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at:
<http://www.ibm.com>.

The IBM System p home page on the Internet can be found at: <http://www.ibm.com/systems/p>.