White Paper

IBM and Information Availability (Updated December 2011)

December 2011
Josh Krischer

Contents

IBM and Information Availability

Josh Krischer

Josh Krischer is an expert IT advisor with 42 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Josh Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering mainframes, enterprise servers and storage from 1998 until 2007. During his career at Gartner he was responsible for Data Center, high-end storage-subsystems and disaster recovery techniques. He spoke on these topics and others at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, as well as major vendor events.

1. What is "Information availability"?

Information Availability is the capability of an IT infrastructure to provide service to its users during defined service periods at an acceptable or agreed-upon level. Data availability is crucial for many organizations, large, medium and small. Many businesses today demand as close to 100% availability as possible for the end-user. This means that business applications must run 24 hours a day, 7 days a week, 365 days a year and give end-users round-the-clock access to critical information. Planned and unplanned system outages can have a serious impact on the business operation. The consequences of downtime may be serious and include lost revenue, lost productivity, loss of customers, lower market share, impact on company share price and damage to the company's brand name and reputation. Outages to critical services may result in litigation, potentially with heavy penalties.

The last few years have been characterized by a huge increase in the demand to store electronic documents and by various compliance regulations, such as the Sarbanes-Oxley Act or Securities and Exchange Commission SEC 17a(4). In many litigation cases, electronic documents, such as email, have been requested for disclosure and presented as court evidence. As a result, businesses are required to keep more data longer, in preparation for litigation discovery requests, which usually must be fulfilled on a short notice. The October 2006 amendments to the U.S. Federal Rules of Civil Procedure (FRCP) require organizations to preserve and be prepared to discover any type of electronically stored information (ESI) on any media. Similar regulations exist in other jurisdictions as well, e.g., the German code of commerce (HGB, §146.5.2, 3), requires that business-related documents be available and quickly-accessible throughout the entire legally-mandated storage period.

The extensive use of computing in all businesses creates the need to protect the organization's data and to provide continuous or near continuous IT service. Hardware and software technology progress provides small and midsized businesses with the opportunity to address availability, disaster recovery, and business continuity in similar ways as large organizations.
Today's information availability demands more than hardware Reliability, Availability, and Serviceability (RAS) features; it includes advanced business continuity, data protection and end-to-end infrastructure management.
In addition to hardware and software availability, information availability includes system availability, local and dispersed clustering, dynamic automated fail-over, synchronous and asynchronous remote data mirroring, data protection capabilities, including Continuous Data Protection (CDP) and services to plan, deploy, and maintain these. IBM provides these components and services on all IBM servers' platforms, from the largest System z10 mainframe to the System x BladeCenter, and includes a comprehensive suite of storage systems and features.

Some commonly accepted business continuity tiers are shown in Figure 1 below.

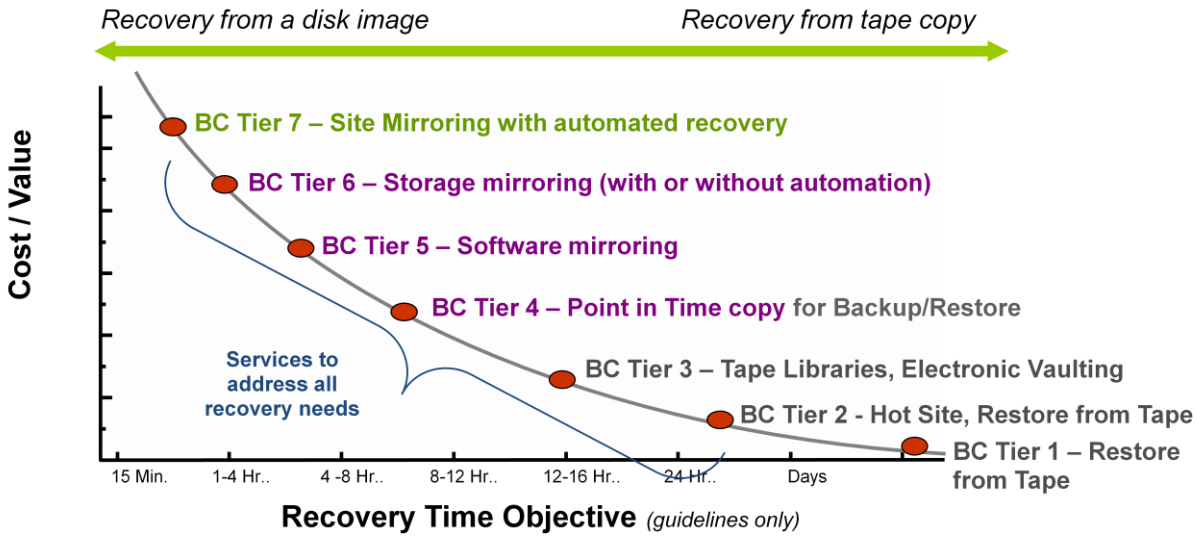

Figure 1: Balancing recovery Time Objective with costs/value (source IBM)

On September, 8th, 2008, IBM announced a framework for organizing and structuring datacenters to meet the evolving needs of corporate enterprises. This framework, known as the Information Infrastructure, is based on four supporting pillars: Information **C**ompliance, Information **A**vailability, Information **R**etention, and Information **S**ecurity (abbreviated as **CARS**). IBM realized this framework by setting Information Infrastructure Training and Certification for Business Partners, Centers of

Excellence, and Business Partner Innovation Centers. It delivers bundled and preloaded solutions, express products for small SMBs, Solution Assurance Reviews with premium support, and flexible, innovative financing.
IBM's Information Availability strategy includes four key elements; business continuity, storage virtualization, data protection, and storage management. This strategy defines steps to ensure that users will have continuous access to critical information and leverages best-practice frameworks for enterprises to achieve the desired level of information resiliency and to mitigate business risks. The IBM Information Availability strategy goal is to maintain continuous and secure access to information, optimize employee productivity and stakeholder satisfaction, and meet service level agreement requirements and both internal and external regulations regarding data retention and delivery.
For many years high availability was one of the defining characteristics of mainframes. However, over the last few years, the virtualization platforms have opened new options to achieve high availability and enhanced business continuity on other platforms as well.

This paper will concentrate on the availability pillar of Information Infrastructure and how IBM products and services fit into this infrastructure.

1.1 Hardware availability

High availability systems are designed to have few[1] or ideally no single points of failure (SPOF) and no single points of repair (SPOR), through the use of redundant components and architectures. Non-disruptive hardware and software upgrades and micro-code updates also play a crucial role in achieving high availability. Some examples of internal availability features in current hardware include:

- Redundant and hot-swappable power supplies, blowers, logic cards
- Error Correction Codes (ECC) on memory
- Memory fencing
- Alternate channels and network paths
- RAID techniques in storage
- Spare disks with automated rebuild.

1 It is common in commercial computer systems to have components which have parts which, in theory, constitute single points of failure. An example may be the power cord in some cases. However, in practice, the probability of having one of these parts fail is so low as to make it not cost-effective to provide redundancy.

Let's take IBM memory ECC for example. IBM developed *Chipkill* or Advanced ECC memory technology specifically for the NASA pathfinder mission to Mars. Standard ECC is able to detect and correct single bit memory errors, which make up the vast majority of memory errors. However, Chipkill memory also has the ability to correct multi-bit memory errors by deploying technique similar to RAID technique used in disk storage subsystems.

The Power 795, in addition to the chipkill technique, also supports Active Memory Mirroring for Hypervisor to prevent downtime in case of a memory failure in the hypervisor area, which otherwise could bring down all partitions on rare occasions. Another high-availability feature, available on IBM mainframes, is automatic activation of a "spare" processor to replace a failing processor. External redundancy is provided by multiple alternate paths, redundant SAN fabric, routers and a redundant network infrastructure. All of these techniques are supported by IBM server and storage products.
For example; IBM DS8000 supports all common used RAID techniques (includes RAID 6) to protect the storage media. The structure of the DS8000 is based on fully redundant server Symmetrical Multi-Processing SMP designed to ensure access to data even in the case of component failure. Each side of the cluster has own cache and the Non Volatile Storage (NVS) of the other cluster side to ensure data integrity protection in the case of a side failure. This "cross connection" protects write data loss in case power loss or other malfunction even for the case that one of the cluster sides suffers from severe malfunction.
The IBM XIV Storage System Model A14 redundancy is designed on the system level and not the component level. It is completely virtualized subsystem where data is divided in 1Mbyte chunks, which are written on all disks in the system. To provide redundancy, each chunk is written twice on different modules, which ensures cache redundancy for write data. Each data module keeps some 'primary' and some 'secondary' blocks. Adding new modules will cause a proportional part of the data to be copied on them. A module failure will cause a reversed action; the "lost" chunks will be copied from the redundant copy on the remaining modules. These operations are completely automated and transparent to the users and the storage administrator. Two other examples are the IBM SAN Volume Controller (SVC) virtualization appliance with it's clustered structure connecting dual, redundant nodes and the TS3500 tape library which supports (optional feature) dual access mechanism to improve performance and availability.

The next step in error recovery is software "retry", to overcome, or to recover from, hardware errors. Software recovery on IBM mainframes, in particular, is very powerful with several layers of retries.

1.2 Clustering – local and remote (stretched or Geo-clusters)

Server redundancy can be provided by local clustering which is supported on all IBM platforms mentioned below. Clusters, initially used to enhance scientific computing, are used today in commercial computing to deliver "scale-out" scalability and higher availability. Equipment redundancy features and local server clustering can protect against server malfunctions but cannot provide protection in the case of local or regional disasters. To protect against local or regional disasters, installations configure systems to exploit remote (or stretched) clustering capabilities and functions. Stretched clusters can span multiple geographic sites to facilitate robust disaster recovery infrastructures. Clusters help to reduce "downtime" and increase availability for both planned and un-planned outages. IBM servers support IBM and third-party clustering software to provide BC Tier 7 for every IBM server platform.

- **System z** supports Parallel Sysplex as a local or remote cluster. Up to 32 local or remote mainframes can participate in a single cluster. System z Parallel Sysplex also works in conjunction with disaster recovery software called Geographically Dispersed Parallel Sysplex (GDPS). GDPS enables automated complete site failover and is an example of a BC Tier 7 solution (See Figure 1). GDPS will be discussed in more detail below.

- **System p** supports PowerHA for AIX and Linux, formerly HACMP (High Availability Clustered Multi Processing) and HACMP/XD (XD - eXtended Distance). PowerHA scales up to 32 nodes running AIX or eight nodes running Linux. It constantly monitors the status of servers, networks and applications to detect failures or performance degradation and responds by automatically restarting a failing application, selecting alternate path to storage and taking care of all network connections in the process.
  System p supports PowerHA SystemMirror for AIX and Linux, formerly HACMP (High Availability Clustered Multi Processing), and HACMP/XD (XD - eXtended Distance). PowerHA XD extends PowerHA's capabilities across geographic sites with Metro-Mirror remote data mirroring and failover using this mirrored data. PowerHA/XD also provides host-based IP-based data mirroring.
  Open HyperSwap, is a new DS8700 and DS8800 function within the Tivoli Storage Productivity Center for Replication. Open HyperSwap is comparable to HyperSwap for System z.
  System p also supports third-party clustering such as the Veritas Cluster Server (VCS) software or EchoCluster from Vision Solutions. Both provide automated failover support on both AIX and Linux platforms. VCS also operates on Hewlett Packard and Sun systems, thereby enabling an installation to install a common automation package in a heterogeneous environment.

PowerVM Live Partition Mobility allows to non-disruptively move a running logical partition (operating system and running applications), from one Power system to another.

- **System i** High Availability Solutions Manager (HASM) is a high availability clustering offering for the i5/OS V6R1 operating system which provides a complete end-to-end hardware-based clustering solution for high availability and disaster recovery operations. HASM with cross-site mirroring (XSM) supports Metro Mirror and Global Mirror on the IBM DS8000.
  Another System i clustering option is iCluster. IBM DataMirror iCluster is based upon IBM i journaling technology. High-speed operational switching automatically detects a primary system failure and invokes failover procedures to switch operations to a backup system. MatchMerge architecture guarantees order of operation and ensures data integrity.
  ISVs such as *noMAX or Cluster1 from Vision Solutions, Inc. offer high availability solutions for System i servers as well.

- **System x** with Windows supports Microsoft Cluster Service (MSCS). MSCS clusters do not replicate data and only provide failover within a single site. Storage control unit or host replication can be used. GeoCluster extends the capabilities of Microsoft Cluster Service (MSCS) to create a stretched cluster. There are also third-party solutions such as Veritas Cluster Server 5.1 and above or VMware High Availability clustering (VMware HA), formerly known as Distributed Availability Services or DAS, which provides high availability to virtual machines through automatic failover within a cluster. Microsoft Hyper-V R2 introduced Cluster Shared Volume« (CSV), which enables multiple virtual machines to share a single shared storage volume and allows migration of these VMs among the different hosts in the cluster.

- **System x virtualization availability features for windows and Linux.** In the last few years the developments of the x86 virtualization hypervisors has brought some of the most significant IT enhancements by delivering features to improve the availability of these platforms. It was pioneered by VMware's Vmotion, which enables non-disruptive migration of running virtual machines from one physical server to another with complete transaction integrity. Both machines require access to shared storage.
  VMware's competition answered with a comparable feature called XenMotion in Citrix Xen Server, Live Migration of Microsoft Hyper-V and Red Hat KVM's Live Migration, respectively.

- The previous paragraph lists site high-availability features of different virtualization platforms; however, in order to deploy disaster recovery at dispersed locations different techniques are used, for example, VMware's Site Recovery Manager (SRM), support planning, discovery, testing, and automated failover. Citrix' XenServer delivers a similar disaster recovery solution. SRM typically fails over an entire site rather than individual components or VMs and requires a "fail-over process", which takes approximately 15 to 20 minutes to complete. A new feature, VMware vSphere Replication, of the VMware vCenter Site Recovery Manager 5 is a built-in hypervisor-based replication product. vSphere Replication (similar to any other host-based replication) enables customers to use heterogeneous storage across sites, which increases flexibility and may reduce costs. Despite this fact, VMware continues to advise large, business-critical environments to use storage-based, while VMware vSphere Replication is recommended for smaller cost-efficient sites.
  Another feature of the VSphere 5 is Metro vMotion (announced August 2011) support over the network, which eliminates the fail-over delay as needed in SRM.

  A dispersed disaster recovery infrastructure requires data replication between the main and the recovery sites. It can be done synchronously or asynchronously through the use of major storage subsystems features. There are also several host-based asynchronous replication applications such as IBM Softek Replicator - a multi-system asynchronous remote copy using TCP/IP to connect to the recovery site or third-party products such as CA/XOsoft, DoubleTake Software (acquired by Vision Solutions), Neverfail, Open-E DSS, etc.

2. Data protection and migration

2.1 Backup/Restore

Traditional data protection techniques that use tape media for backup and recovery are often adequate for less time sensitive information and are considerably less expensive. Backup is a crucial part of availability because it stores the latest logical image of the data. Some customers using remote data replication lowered the importance of traditional backup and recovery using tape storage and have found that this may expose data to corruption or human mistakes as errors on the primary instance of data are propagated to all mirrored local or remote active mirror copies. IBM provides robust data protection function with Tivoli Storage Manager (TSM). TSM is composed of a family of offerings provide centralized, automated data protection to store and manage a range of data types such as backup and archive for applications, and databases with space and data retention management.

Tivoli (and other backup applications) also utilizes the IBM System Storage ProtecTIER data deduplication (a.k.a. Diligent) offering, that can significantly reduce the amount of disk storage required for active backup information. On July, 14th, 2009 IBM announced native IP-based replication capability for IBM System Storage™ TS7650 ProtecTIER Gateways and Appliances. This new enhancement functionality replicates backup data on virtual tape cartridges to remote location. It allows flexibility to replicate or move all or some virtual tape cartridges at preset time and at required priority. ProtecTIER replicate only the deduplicated data which can be up to 25 times smaller than the raw data which means that the required bandwidth for the transmission can be significantly lower and cheaper.
In addition, IBM provides a comprehensive family to tape storage offerings that work with TSM and other popular data protection software. The tape systems span the gamut from entry single tape drive LTO offerings to the TS3500 tape library which can store over 6800 LTO cartridges. The LTO Generation 4 and TS1120 and TS1130 tape drives and associated libraries also offer encryption capabilities to provide for safe, secure storage of sensitive information.

2.2 Local Point-in-Time (PiT) copies or Flashcopies

PiT is an important feature in maintaining applications' high availability (BC Tier 4 on figure 1). Many operational processes which are often disruptive, such as backup, can be performed on the "frozen" copy of the primary data allowing applications to continue with minimal disruptions. IBM FlashCopy SE ("Space Efficient") is an incremental snapshot-type local replication to complement the "full volume" FlashCopy feature. A "full volume" FlashCopy requires allocation of capacity equal to or greater than the source volume regardless of the number of modifications. The "space efficient" FlashCopy only consumes capacity for modifications, and therefore, can require less storage capacity in comparison to the "full Volume" technique. Users benefit from better storage utilization and lower TCO. Another possibility enabled by SE's potentially smaller-capacity consumption is the ability to create several point-in-time copies (time shifted), with the ability for fast data recovery. In case of data corruption or operator mistake, the recovery can be performed by returning to the last uncorrupted copy. Snapshot capabilities are required for several near-Continuous Data Protection (CDP) techniques. IBM DS8000 supports both types of flash copies. All other IBM storage subsystems support at least one technique of local PIT copy.

2.3 Continuous Data Protection (CDP)

SNIA formally defines CDP as *"a methodology that continuously captures or tracks data modifications and stores changes independent of the primary data, enabling recovery points from any point in the past."* Continuous data protection (CDP) is

sometimes referred to as "continuous backup" because it stores records or “logs” in real time. CDP products that maintain running logs of individual transactions can be restored to within milliseconds of a failure. CDP creates a running journal of storage activity, with a new entry generated each time a change occurs to the system. In the case of data corruption of a data base for example, the administrator can take the last snapshot of the system just prior to the event and then restore the database to the very point prior to an event using the logs. CDP is not intended to replace existing backup tactics, but instead provide a new layer of protection that can improve IT service to enterprise users, by speeding the time to recovery. It is a supplemental technique to recover from logical data corruption and accidental or malicious deletion of critical data but it not provides real disaster recovery protection.

On April 21, 2008 IBM acquired a company that specializes in continuous data protection and nearly instant data and application recovery software for enterprises and remote/branch offices. FilesX Xpress Restore software enables global enterprises to have full protection and easy recovery of data for mission-critical Windows applications, in the data center and in remote branch offices, such as Microsoft® Exchange, SQL, Lotus NotesTM, Oracle and SAP. The FilesX technology is now integrated and branded in the IBM Tivoli Storage Manager (TSM) family of products.

2.4 Data Migrations

One of the most disruptive tasks is installing new storage subsystems and migrating the data. It is usually also a time consuming and sometimes risky task. Data migrations are required also during tuning activities or for building tiered storage infrastructure. The IBM SAN Volume Controller or the StoreWize v7000 storage subsystem can improve application availability by performing dynamic data migrations. This function non-disruptively moves data from one storage subsystem to another. This ability can help administrators reallocate and scale storage capacity without disrupting applications; it can help contribute to increased availability of applications and better SLAs for customers.
Other migration options are offered by IBM Global Technology Services (IGTS), using the software migration tools such as Softek Transparent Data Migration Facility (TDMF) and Softek Logical Data Migration Facility (LDMF). IGTS can assist users with project planning and management and can provide technical assistance of experienced services professionals.
All these options support data migrations between storage subsystems of different vendors.

3. Availability and Business Continuity Planning, Risk Analysis, Business Impact Analysis and Business Continuity Management

For many companies, a heightened awareness of the need to have an up-to-date tested Business Continuance / Disaster Recovery plan was partially trigged by the tragic events of September 11th, Hurricane Katrina and, the Northeast Blackout of August 2003 (a massive power outage that occurred throughout parts of the Northeastern and Midwestern US, and Ontario, Canada). These examples of extreme regional disasters (although sometimes called “global disasters”) are relatively rare in comparison to “local disasters”, such as power grid failures, fires, building water damage, flooding, and extreme weather conditions. Companies are also becoming increasingly aware that BC/DR plans are necessary to enable compliance with existing and emerging government regulations.

There are no “golden rules” about how to deploy the "best" recovery strategy for a company. The selection of the most suitable techniques depends on a number of issues, including: the geographic locations of the sites, the possible natural disasters in these locations and the nature of the business-critical applications. Executing a business impact analysis (BIA) study will help determine how much data the institution can afford to recover or lose what the costs are to recover the data, and what the business damages are of not being in operation for some period of time. Upon completion of the BIA study, the company is in a position to make informed trade-offs: how much are they willing to spend to prepare for a particular event vs. the impact and probability of occurrence of a particular event. IBM Global Technology Services can conduct or merely lead a BIA study having accumulated knowledge and experience from hundreds of such engagements.

The experience collected from many disasters has caused a shift in disaster recovery planning from merely being able to recover IT-related aspects in the event of a disaster to corporate-wide business continuity management, which in addition to the IT infrastructure, also includes other critical resources such as:

- housing for company personnel in the event of a regional disaster,
- internal and external communication,
- travel to the recovery site, and
- alternate facilities which act as a recovery site.

Ensuring business continuity plans are complete, up-to-date and tested is an ongoing process. Having effective change management procedures which are linked to the

Business Continuance Plan is critical to ensuring the BC/DR plans remain up-to-date and function properly.
IBM Tivoli Business Continuity Process Manager V7.1 (BCPM) helps automate the key processes required to identify and recover critical business systems as quickly as possible after an outage occurs. Hence, BCPM can help reduce the cost to the business due to the unavailability critical applications. The pre-tested automated processes are ITIL[2] compliant and can be handled by less-skilled operators. Management is informed through automatic alerting, thereby enabling them to monitor and manage execution of the recovery plan.

3.1 Different techniques of remote data replications, data consistency

Working in conjunction with server clustering, storage systems provide a number of proven techniques for supporting Business Continuity. There are two basic techniques for remote data replication: synchronous mirroring and asynchronous mirroring. The synchronous mirroring function is performed by a storage control unit. Asynchronous mirroring can be done by the control unit, by software that runs on a host or by special appliances that are part of the Storage Area Network (SAN) such as a virtualization device. Control-unit-based techniques require similar hardware at both the Primary and Secondary sites. The choice of which replication technique to use is determined by the distance between the Primary and Secondary sites, the available infrastructure for the remote data links, i.e. dark fiber or a common carrier, budget, and the business recovery requirements.

There are two principal metrics which can quantify an organization's business recovery requirements: Recovery Point Objective (RPO) and Recovery Time Objective (RTO). RPO is the point in time that marks the end of the period during which data can still be recovered after a disaster event and is a direct measure of potential data loss – higher RPO times indicates potentially more transactions will be lost. Another important factor is the Recovery Time Objective (RTO), which is the length of time that transpires between when a disaster occurs and when the business process must be back in production. Usually synchronous remote copy has shortest RPO, in comparison to asynchronous, which usually collects and transfers "groups" of

2 Information Technology Infrastructure Library (ITIL) is summary of a best practices framework for IT projects and processes. It was developed by the Central Computer and Telecommunications Agency (CCTA) of the United Kingdom (UK) government in the late 1980s and used by many organizations worldwide.

transactions. Shortening the RPO and RTO is in inverse relation to the budget required to achieve the desired RPO and RTO.

A very popular initial business requirement for planning is; “I do not want to lose a single transaction” which means RPO=0 and requires synchronous remote copy. However, in reality, RPO=0 cannot be achieved and it is most likely that even in the best case the last modification(s) will be lost. Ensuring minimal data losses is crucial for some applications, but a much more important issue is assuring data consistency, write data order and data integrity at the recovery site, i.e., it is better to be consistent and missing a few records than having them all but non-consistent. If the data is not consistent at the recovery site, a time-consuming backup is usually required, which may take days. Also, resolving data conflicts and reconciling the status of key information may cause a lengthy recovery time.

All IBM storage subsystems support array-based remote copy techniques. IBM System z supports host-based asynchronous remote copy - z/OS Global Mirror (formerly XRC). Other IBM server platforms support third party asynchronous host-based replications such as Double-Take for Windows and Linux. IBM storage subsystems support the following remote copy techniques (BC Tier 6 on figure1):

- The high-end enterprise DS8000 and DS6000 series support synchronous mirroring, known as IBM Metro Mirror (formerly PPRC) and asynchronous mirroring, known as IBM Global Mirror.
- The IBM XIV Storage System supports synchronous and asynchronous mirroring.
- The mid-range DS4000 and DS5000 series support synchronous and asynchronous mirroring. The Global Mirror asynchronous mirroring function supports write-order consistency.
- IBM N series storage (OEMed from NetApp) supports asynchronous mirroring, known as SnapMirror, over LAN or WAN using IP protocol and SyncMirror, which is a synchronous remote copy. The N series supports Fabric Attached MetroCluster, which is a combination of array-based clustering with synchronous mirroring. The two separated chassis contain one controller each, and are connected in active/active stretched MetroCluster configuration using the FC-VI card to enable heartbeat connectivity between the controllers. Because of the active/active configuration there is no need to create failover scripts. In case of a disaster in the primary site, services are restored at the secondary site with an automated single command in a matter of seconds, without failover scripts or restart procedures.

- The IBM DR550 model 2 archiving system supports Global Mirror Write-order consistency and synchronous mirroring.
- The IBM SAN Volume Controller (SVC) and the StoreWize v7000 support both synchronous Metro Mirror and asynchronous Global Mirror between heterogeneous storage subsystems from IBM and other vendors. In October 2011 IBM announced the ability to replicate between SVC and the The SVC supports automated failover similar to the MetroCluster of the N series. V7000 or vice-verse.
- IBMs Virtualization Engine for Tape "Grid" function sometimes called "tape mirroring", writes the data to a local IBM TS7740 and then automatically copies it to a remote mirror.

3.2 Geographical Dispersed Parallel Sysplex (GDPS)

Multiple sites and remote data mirroring are only a part of a well-structured disaster recovery infrastructure. For very high availability (Tier 7) applications, a second essential component is software to control the fail-over and to provide end-to-end application availability in case of a planned outage or an unplanned disaster. IBM's Geographically Dispersed Parallel Sysplex (GDPS) for System z is a multi-site application-availability solution, with fast recovery time and highly-automated control. It manages application availability in and across sites for both planned maintenance and unplanned situations, such as site failure or full-blown disaster. GDPS was initially designed for mainframe z/OS systems, but with continuous development was later enhanced to support selected other systems platforms, as well. GDPS can now manage other System z production operating systems, such as, Linux for System z, z/VM, and z/VSE.
At the end of June 2011 IBM launched GDPS 3.8, which extended the zSeries' business resiliency strength via significant enhancements to GDPS: an active/active configuration (in addition to the active/standby, which is commonly used). The GDPS active/active continuous availability is the next generation of GDPS and a fundamental paradigm shift from a failover model to a near-continuous availability model. IBM GDPS active/active continuous availability is a solution for organizations using two sites separated by unlimited distances, running the same applications, and having the same data with cross-site workload monitoring, data replication, and balancing. In IBM tests a planned active/active site switch took 20 second compared to 1-2 hours, and an unplanned workload switch took 150 second compared to 1 hour with the previous version.

GDPS is a licensable product that a customer deploys thru IBM Global Technology

Services. Once the installation and testing are complete, the customer is self-sufficient in operation, modifications, configuration changes, updates etc.

3.3 GDPS for other platforms

Many applications (such as SAP ERP) may use multiple platforms, with a mainframe-resident database and the application on another platform. In such cases, mainframe disaster recovery is not sufficient to provide robust end-to-end business continuity. To answer these issues, GDPS evolved from mainframe disaster recovery automation to data center automation. The GDPS/PPRC Open LUN management capability allows PPRC replication for both System z and other data platforms from a single point of control, providing a complete disaster recovery solution for applications that span heterogeneous systems.
The IBM/Tivoli software product xDR expands GDPS/PPRC functionality to Linux for System z and provides disaster recovery for users with distributed hybrid applications that span z/OS and Linux for System z (e.g., SAP ERP, WebSphere, etc.).

3.4 HyperSwap function

The HyperSwap function is probably the most important business continuity and availability improvement for IBM mainframes. While disasters seldom occur in reality, disk subsystem failures are far more likely to happen. In current integrated and complex application environments − assuming a highly available data-sharing Parallel Sysplex environment− disk becomes a single point of failure for the entire Sysplex. The HyperSwap function, which is used by multiple GDPS solutions, is controlled by GDPS automation. Use of HyperSwap can eliminate an outage caused by planned maintenance or disk failure by reducing the time needed to switch disks between sites to a matter of seconds and allowing the primary site to use the secondary site disk storage subsystems.
Basic HyperSwap between two local installed storage subsystems to provide automated fail-over for planned or un-planned outages can be deployed with z/OS only, without multi-side GDPS.

Open HyperSwap, is a new DS8700 and DS8800 function within the Tivoli Storage Productivity Center for Replication. Open HyperSwap (comparable to HyperSwap for System z) is an important feature improving the continuous availability of AIX by managing a set of planned and unplanned disk system outages for Metro Mirror PPRC-capable disk systems. It provides an ability to swap IBM DS8000 volumes (from the primary storage system to the secondary storage system) in seconds, reacting to operator command or in automatic (no operator interaction) fashion upon a storage system failure. It is designed to scale to multi-thousands of volumes. This

non-disruptive feature, which allows applications to keep using the same device addresses, is configured and managed by Tivoli Storage Productivity Center for Replication V4.2.

3.5 Capacity Back-UP (CBU)

A best practice in disaster recovery planning is to purchase provision of additional processors, which can be activated in an emergency to compensate for lost capacity. This option enables business-critical applications to run with appropriate service levels without up-front investment for "spare" capacity. GDPS can automatically activate the CBUs in case of a disaster. CBU options are available also for the IBM Power Systems.

4. IBM Business Continuity and Resiliency Services (BCRS)

IBM offers a rich suite of availability and business continuity products. However, the main challenge for many organizations is the lack of available experienced IT staff to develop the corporate-wide recovery plan, deploy the appropriate BC/DR recovery products, and then test the overall plan. IBM Business Continuity and Resiliency Services (BCRS) from IBM Global Technology Services can be employed to design, manage and deploy a customized, integrated, end-to-end continuity program. The solutions are designed to be integrated into the existing corporate organization, and meet relevant government and industry regulations.

IBM has over 1,300 professionals dedicated to Business Continuity and Resiliency Services (BCRS) located in over 70 countries. Additionally, BCRS maintains 154 recovery centers across 55 countries around the world to help facilitate businesses engaged in disaster recovery.

BCRS supports a wide variety of third-party hardware and software components, in addition to IBM products. BCRS can be employed to perform a risk analysis and work with the organization's users to perform a Business Impact Analysis (BIA). The BIA will identify critical business processes, components, platforms and applications, and analyze individual recovery requirements. It can also provide alternate recovery services in the event of a major disaster, including hot site provisioning, mobile data centers, and staff. BCRS's vendor- and platform-neutral approach allows enterprises to preserve investment in third-party products and existing infrastructure while finding the most efficient route to achieve the appropriate level of protection for the business.

IBM Information Protection Services (Arsenal) can help customers provide physical protection for critical data.

5. Summary and conclusions

Constant availability and business continuity are requirements of most organizations. Many availability features are now standard and "built-into" modern storage hardware and software. More advanced options are sold or licensed separately. One of the most complicated tasks an enterprise faces is building a business continuity infrastructure. Planning for business continuity and disaster recovery planning is not a trivial issue and can be the most important task of the project. Bad assumptions in the planning process are usually very difficult to correct without introducing high financial and operational costs. Business impact analysis will help to get the appropriate funding; risk impact analysis will determine the distances, technologies, and costs. All of these may create challenges that exceed the ability of a business, education or government organization to complete.

IBM's Information Infrastructure framework is helping define the service delivery strategies and functional capabilities of evolving corporate enterprise data centers. IBM's advantages are its rich server and storage portfolio, Tivoli automation suite of products, advanced virtualization, GTS services, and the synergies between IBM hardware and software. IBM is a global corporation and proven infrastructure provider for all platforms, and is recognized for its innovation and advanced technology; the industry's broadest and deepest services; its many flexible financing options; financial strength; and unshakeable viability. IBM can help large and small organizations achieve their appropriate level of business continuity and disaster protection.