

Beyond Disaster Recovery: Taking Your Enterprise from High Availability to Continuous Availability

Karen Durward
IBM

August 11, 2011
Session 9666

Important Disclaimer



© Copyright IBM Corporation 2011. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

The information contained in this presentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided “as is” without warranty of any kind, express or implied. In addition, this information is based on IBM’s current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other documentation. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of any agreement or license governing the use of IBM products and/or software.

IBM, the IBM logo, ibm.com, InfoSphere, DataStage, MetaStage, QualityStage, Information Agenda, and Information on Demand are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

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



Material Relating to Future IBM Products, Features, or Functionality



- ***Disclaimer:***

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.




Shifting to Continuous Availability



- A little bit of history...
- Drivers
- Requirements
- Concepts
- Configurations


Availability and the IBM Mainframe

1964




**Single
System**

1982-1990



**Multiple
partitions/images**


1995



**Parallel
Sysplex**

Availability and the IBM Mainframe

1998



Parallel Sysplex

Synchronous Mirroring

HyperSwap for primary disk failures

Near-continuous Availability

Limitation:

Performance at distance

**Geographically Dispersed Parallel Sysplex
Peer to Peer Remote Copy
(GDPS/PPRC)**

What are customers doing today ?




Continuous Availability
of Data
within a Data Center

GDPS/HyperSwap Mgr
RPO = 0 / RTO = 0

Single Data Center
Applications remain active

Continuous access to data
in the event of a storage
subsystem outage



GDPS / HyperSwap Manager

- Synchronizing copies of data within a data center using a primary and secondary storage control unit
- Provides continuous data availability by protecting against a storage outage within a data center

ISSUE:


- No protection given for a data center level failure!

NOTE:

Recovery Point Objective (RPO) measures data copy latency

Recovery Time Objective (RTO) measures time to “system” recovery

What are customers doing today ?



GDPS/PPRC

- Built on a multi-site Parallel Sysplex and synchronous disk replication
- Provides
 - Metro-area Disaster Recovery
 - Limited to ~120 miles (200 km)
 - RTO up to an hour to start applications on the remote site
 - Metro-area Continuous Availability
 - RTO of Zero but limited to ~20 miles (35 km) apart using fiber
 - Signal latency impacts OLTP thru-put and batch duration
- **Issue:**
Insufficient site separation for some workloads!

What are customers doing today ?




Continuous Availability of Data within a Data Center

GDPS/HyperSwap Mgr
RPO = 0 / RTO = 0

Single Data Center
Applications remain active

Continuous access to data in the event of a storage subsystem outage




Continuous Availability w/ Disaster Recovery within a Metropolitan Region

GDPS/PPRC
RPO = 0 / RTO <1hr (>20 km)
RPO = 0 / RTO = 0 (<20 km)

Two Data Centers
Systems remain active

Multi-site workloads can withstand site and/or storage failures




Disaster Recovery at Extended Distance

GDPS/GM & GDPS/XRC
RPO secs / RTO <1 hr

Two Data Centers
Rapid Systems Disaster Recovery with "seconds" of Data Loss

Disaster recovery for out of region interruptions



GDPS/XRC and GDPS/GM

- Added asynchronous disk replication
- Unlimited distance for Disaster Recovery solutions
- Requires the failed site's workload to be restarted in the recovery site resulting in RTO of up to an hour

Issue:

Can NOT achieve an RTO of seconds

What are customers doing today ?



Continuous Availability of Data within a Data Center	Continuous Availability w/ Disaster Recovery within a Metropolitan Region	Disaster Recovery at Extended Distance	Regional Continuous Availability w/ Disaster Recovery @ Extended Distance
GDPS/HyperSwap Mgr RPO = 0 / RTO = 0	GDPS/PPRC RPO = 0 / RTO <1hr (>20 km) RPO = 0 / RTO = 0 (<20 km)	GDPS/GM & GDPS/XRC RPO secs / RTO <1 hr	GDPS/MGM & GDPS/MzGM
Single Data Center Applications remain active Continuous access to data in the event of a storage subsystem outage	Two Data Centers Systems remain active Multi-site workloads can withstand site and/or storage failures	Two Data Centers Rapid Systems Disaster Recovery with "seconds" of Data Loss Disaster recovery for out of region interruptions	Three Data Centers High availability for site disasters Disaster recovery for regional disasters

Hybrid Solution



GDPS/PPRC

- Provides a continuous availability environment by using two sites within a metro region
 - Protects against single site outages

GDPS/XRC

- Uses a third site to provide Disaster Recovery for both primary sites
 - Protects against a regional disaster

Issue:

Still can not achieve an RTO of seconds during a regional outage

Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System

[Docket No. R-1128] (April 7, 2003)



1. Identify clearing and settlement activities in support of critical financial markets
2. Determine appropriate recovery and resumption objectives for clearing and settlement activities in support of critical markets

...core clearing and settlement organizations should develop the capacity to *recover and resume* clearing and settlement activities within the business day on which the disruption occurs with the overall goal of achieving recovery and resumption *within two hours* after an event

3. Maintain *sufficient geographically dispersed resources* to meet recovery and resumption objectives
 - Back-up arrangements should be as far away from the primary site as necessary to avoid being subject to the same set of risks as the primary location.
 - The effectiveness of back-up arrangements in recovering from a wide-scale disruption should be confirmed through testing

4. *Routinely use or test* recovery and resumption arrangements.


One of the lessons learned from September 11, 2001 - is that testing of business recovery arrangements should be expanded



How much interruption can your business tolerate?

Ensuring Business Continuity:

Standby



Active/Active

- Disaster Recovery
 - Restore business after an unplanned outage
- High-Availability
 - Meet Service Availability objectives
 - 99.9% availability
 - 8.8 hours of down-time a year
- Continuous Availability
 - No downtime (planned or not)

Cost of 1 hour of downtime during core business hours

Cost of Downtime by Industry	
Industry Sector	Loss per Hour
Financial	\$8,213,470
Telecommunications	\$4,611,604
Information Technology	\$3,316,058
Insurance	\$2,582,382
Pharmaceuticals	\$2,058,710
Energy	\$1,468,798
Transportation	\$1,463,128
Banking	\$1,145,129
Chemicals	\$1,071,404
Consumer Products	\$989,795

Source: Robert Frances Group 2006, "Picking up the value of PKI: Leveraging z/OS for Improving Manageability, Reliability, and Total Cost of Ownership of PKI and Digital Certificates."

Global Enterprises that operate across time-zones no longer have any 'off-hours' window, making Continuous Availability a requirement!

Disruptions affect more than the bottom line



September 9, 2008

London Stock Exchange Paralyzed by Glitch

THE WALL STREET JOURNAL.

August 4, 2010

Singapore Censures DBS Bank For System Outage On July, 5 2010

GOVMONITOR
Public Sector News & Information

September 6, 2010

Virginia Grapples with IT Outage

InformationWeek

April 21, 2011

Failure at Amazon data center takes down sites across Internet; millions of Web users affected,

AP Associated Press

Enormous impact on the business

- Downtime costs can equal up to 16 percent of revenue ¹
- 4 hours of downtime severely damaging for 32 percent of organizations, ²
- Data is growing at explosive rates –
growing from 161EB in 2007 to 988EB in 2010³
- Regulatory compliance implications as fines associated with downtime are common in some industries
- Downtime ranges from 300–1,200 hours per year, depending on industry¹

¹ Infonetics Research, *The Costs of Enterprise Downtime: North American Vertical Markets 2005*, Rob Dearborn and others, January 2005.

² Continuity Central, "Business Continuity Unwrapped," 2006, <http://www.continuitycentral.com/feature0358.htm>

³ The Expanding Digital Universe: A Forecast of Worldwide Information Growth Through 2010, IDC white paper #206171, March 2007

Customer Requirements

RTO near zero, Replace roll-your-own, Leverage all resources





Shift focus from failover to nearly-continuous availability


- “Recover my business rather than my platform technology”
 - Multi-sysplex, multi-platform solution
 - No application changes
 - Access data from any site with unlimited distance between sites
 - Provide application level granularity rather than the current “all-or-nothing” model
 - Some workloads may require immediate access from every site
 - Some workloads may only need to update other sites every 24 hours
- Minimize costs and Optimize resource utilization
 - Automated recovery processes (similar to GDPS technology today)
Minimizing operator learning curve
 - Provide workload distribution between sites
 - Dynamically select sites based on their ability to handle workload
 - Route around failed sites

Active/Active concepts

San Jose



London



Two or more sites, separated by unlimited distances, running the same applications & having the same data to provide:


- Cross-site Workload Balancing
- **Continuous** Availability
- Disaster Recovery

Workload Distributor:


Workloads are managed by a client and routed to one of many replicas, depending upon workload weight and latency constraints, extending workload balancing to SYSPLEXs across multiple sites!

Active/Active concepts

San Jose



Replication



London

Workload Distributor

Transactions

Two or more sites, separated by unlimited distances, running the same applications & having the same data to provide:


- Cross-site Workload Balancing
- **Continuous** Availability
- Disaster Recovery

Replication:


Data at geographically dispersed sites are kept in sync via software-based data replication

Active/Active concepts



San Jose



Replication



London



Two or more sites, separated by unlimited distances, running the same applications & having the same data to provide:

- Cross-site Workload Balancing
- **Continuous** Availability
- Disaster Recovery



Tivoli Enterprise Portal

Tivoli Enterprise Portal:

Monitoring spans the sites and now becomes an essential element of the solution for site health checks, performance tuning, etc.

GDPS/Active-Active

- ***Configurations***

- Active/Standby – announce 24th May 2011, GA 30th June 2011
- Active/Query – statement of direction
- Active/Active – intended direction

- ***Specify configuration on a “workload” consisting of:***

- Software:
user written applications (*e.g., COBOL program*) and the middleware run time environment (*e.g., CICS region*)
- Data:
related set of objects that must preserve transactional consistency and optionally referential integrity constraints
(*e.g., DB2 Tables*)
- Network connectivity:
one or more TCP/IP addresses & ports
(*e.g., 10.10.10.1:80*)


What is a GDPS/Active-Active environment?

Active-Standby Configuration Today



- ***Two Production Sysplex environments (“sites”)***
 - One active, one standby – for each defined workload
 - Either site can be active or standby for specific workloads
 - Sites can be widely geographically dispersed
 - Software-based replication between the two sites
 - IMS and DB2 databases are supported now
- ***Two Controller Systems***
 - Primary and Backup
 - Typically one in each of the production locations, but there is no requirement that they are co-located in this way
- ***Workload balancing/routing switches***
 - Must be Server/Application State Protocol compliant (SASP)
 - RFC4678 describes SASP

Conceptual view




What Software comprises a GDPS/Active-Active environment?

Integration of multiple software products

- z/OS 1.11 or higher
- IBM Multi-site Workload Lifeline v1.1
- IBM Tivoli NetView for z/OS v6.1
- IBM Tivoli Monitoring v6.2.2 FP3
- IBM InfoSphere Replication Server for z/OS v10.1
- IBM InfoSphere IMS Replication for z/OS v10.1
- System Automation for z/OS v3.3
- GDPS/Active-Active v1.1
- Optionally the OMEGAMON suite of monitoring tools to provide additional insight

High level architecture



Functional breakdown

- ***IBM Multi-site Workload Lifeline v1.1***

- **Advisor**.....runs on the Controllers & provides information to the external load balancers on where to send transactions and information to GDPS on the health of the environment. There is one primary and one secondary advisor
- **Agent**.....runs on all production images with active/active workloads and provides information to Lifeline Advisor on health of that system

- ***IBM Tivoli NetView for z/OS v6.1***

- Runs on all systems, providing automation and monitoring functions.
- NetView Enterprise Master normally runs on the Primary Controller

- ***IBM Tivoli Monitoring v6.2.2 FP3***

- Run on the Controllers, **or** Linux on System z, **or** distributed servers
- Monitoring infrastructure & portal plus alerting/situation management via:
 - Tivoli Enterprise Portal
 - Tivoli Enterprise Portal Server
 - Tivoli Enterprise Monitoring Server

Functional breakdown

- ***IBM InfoSphere Replication Server for z/OS v10.1***
 - Runs on production images where required to capture (active) and apply (standby) data updates for DB2 data.
 - Relies on MQ as the data transport mechanism (QREP).
- ***IBM InfoSphere IMS Replication for z/OS v10.1.1***
 - Runs on production images where required to capture (active) and apply (standby) data updates for IMS data.
 - Relies on TCPIP as the data transport mechanism.
- ***System Automation for z/OS v3.3***
 - Runs on all images, providing many critical functions, e.g.:
 - BCPii (communicate with consoles)
 - System Automation infrastructure for workload and server management
 - Remote communications capability to enable GDPS to manage sysplexes from outside the sysplex

Functional breakdown

GDPS/Active-Active v1.1:

Runs only on the Controllers. Provides the following functions:



- Workload management (e.g.: start/stop of workloads)
- Replication management (e.g.: start/stop replication between sites)
- Routing management (e.g.: start/stop routing to a site)
- System & Server management (e.g.: STOP, LOAD, RESET, CBU/OOCOD)
- Monitoring of the environment
- Alerting based on any unexpected deviations
- Planned/Unplanned situation management & control
- Complex/compound scenario automation via scripting capability

Optional:


OMEGAMON suite of monitoring tools for additional insight

Functional Breakdown


Active Production




Standby Production



Primary Controller




Backup Controller



Sample environment –

All workloads active in Site 1, Site 2 is only for stand-by




Sample scenario –

Unplanned workload outage

- Failure is detected and an alert is issued
- GDPS receives the alert and runs a monitor to check the status of both sites
- If GDPS monitor finds an issue that could prevent a switch to the standby site, a prompt is generated on the console
- If all is OK, transactions are automatically routed to the stand-by site, making it active, as the workload is already available to process work
 - There is a policy option to prompt the operator for a switch decision rather than automatic switching.
- Once the switch takes place, GDPS schedules an unplanned workload switch script to carry out any post switch actions that may be required such as adding capacity via CBU/OOCOD.

Sample environment –

Some workloads active in Site1, others active in Site 2



Sample scenario – *Site1 planned outage*

- Stop routing transactions to workloads active in Site1
- Stop replication from Site1 to Site2
- Stop replication for Site2 active workloads sending updates to Site1 (if any)
- Start routing transactions for workloads previously active in Site1 to Site2

The workloads are now processing transactions in Site2 for all workloads, with no replication to Site1






- Then, as required by the scenario:
 - Stop workloads in Site1
 - Close down all systems in Site1

Positioning

- GDPS/Active-Active is for mission critical workloads that have stringent recovery objectives that can not be achieved using existing GDPS solutions.
 - RTO measured in seconds for unplanned outages
 - RPO measured in seconds for unplanned outages
 - Non-disruptive site switch of workloads for planned outages
 - At any distance
 - NOT intended to substitute for local availability solutions (eg: Parallel Sysplex enabled applications)

What are customers doing today ?



Continuous Availability of Data within a Data Center	Continuous Availability w/ Disaster Recovery within a Metropolitan Region	Disaster Recovery at Extended Distance	Regional Continuous Availability w/ Disaster Recovery @ Extended Distance	Continuous Availability, Disaster Recovery, and Cross-site Workload Balancing at Extended Distance
GDPS/HyperSwap Mgr	GDPS/PPRC	GDPS/GM & GDPS/XRC	GDPS/MGM & GDPS/MzGM	GDPS Active-Active Sites
RPO = 0 / RTO = 0	RPO = 0 / RTO <1hr (>20 km) RPO = 0 / RTO = 0 (<20 km)	RPO secs / RTO <1 hr		RPO seconds / RTO seconds
Single Data Center Applications remain active	Two Data Centers Systems remain active	Two Data Centers Rapid Systems Disaster Recovery with "seconds" of Data Loss	Three Data Centers High availability for site disasters	Two or More Data Centers All sites active
Continuous access to data in the event of a storage subsystem outage	Multi-site workloads can withstand site and/or storage failures	Disaster recovery for out of region interruptions	Disaster recovery for regional disasters	
				
Components				
Tivoli – NetView, SAz System z, DS8K, PPRC GDPS control code, Services	Tivoli – NV, SAz, SA-MP, AppMan System z, DS8K, VTS, PPRC GDPS control code, Services	Tivoli – NV, SAz System z, DS8K, Global Mirror, XRC GDPS control code, Services	Tivoli – NV, SAz System z, DS8K, MGM, MzGM GDPS control code, Services	Tivoli – SA, NetView Multi-site Workload Lifeline DB2 & IMS replication System z, DS8K, Global Copy GDPS control code, Services

RPO – recovery point objective
RTO – recovery time objective

Thank You