



Improving IT Availability – New Concepts in Delivering Continuous Availability with Active / Active Sites

Dave Petersen
IBM Distinguished Engineer,
IBM Systems and Technology Group



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

| | | | |
|--|---|------------------------------------|-----------------------|
| AIX* | Geographically Dispersed Parallel Sysplex | OS/390* | Virtualization Engine |
| APPN* | HiperSockets | Parallel Sysplex* | VSE/ESA |
| CICS* | HyperSwap | PR/SM | VTAM* |
| DB2* | IBM* | Processor Resource/Systems Manager | WebSphere* |
| DB2 Connect | eServer | RACF* | z/Architecture |
| DirMaint | IBM logo* | Resource Link | zEnterprise |
| DRDA* | IMS | RMF | z/OS* |
| Distributed Relational Database Architecture | InfoPrint* | S/390* | z/VM* |
| e-business logo* | Language Environment* | Sysplex Timer* | z/VSE |
| ECKD | MQSeries* | System z | zSeries* |
| Enterprise Storage Server* | Multiprise* | System z9 | |
| ESCON* | NetView* | System z10 | |
| FICON* | On demand business logo | TotalStorage* | |
| GDPS* | | | |

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.


Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- **Level Set**
- **Active/Active Sites Overview**
- **Preliminary Test Results**
- **Components**
- **Summary**



System z
environments





Improving IT Availability - New Concepts in Delivering Continuous Availability with Active / Active Sites




Level Set

Active/Active Sites Overview

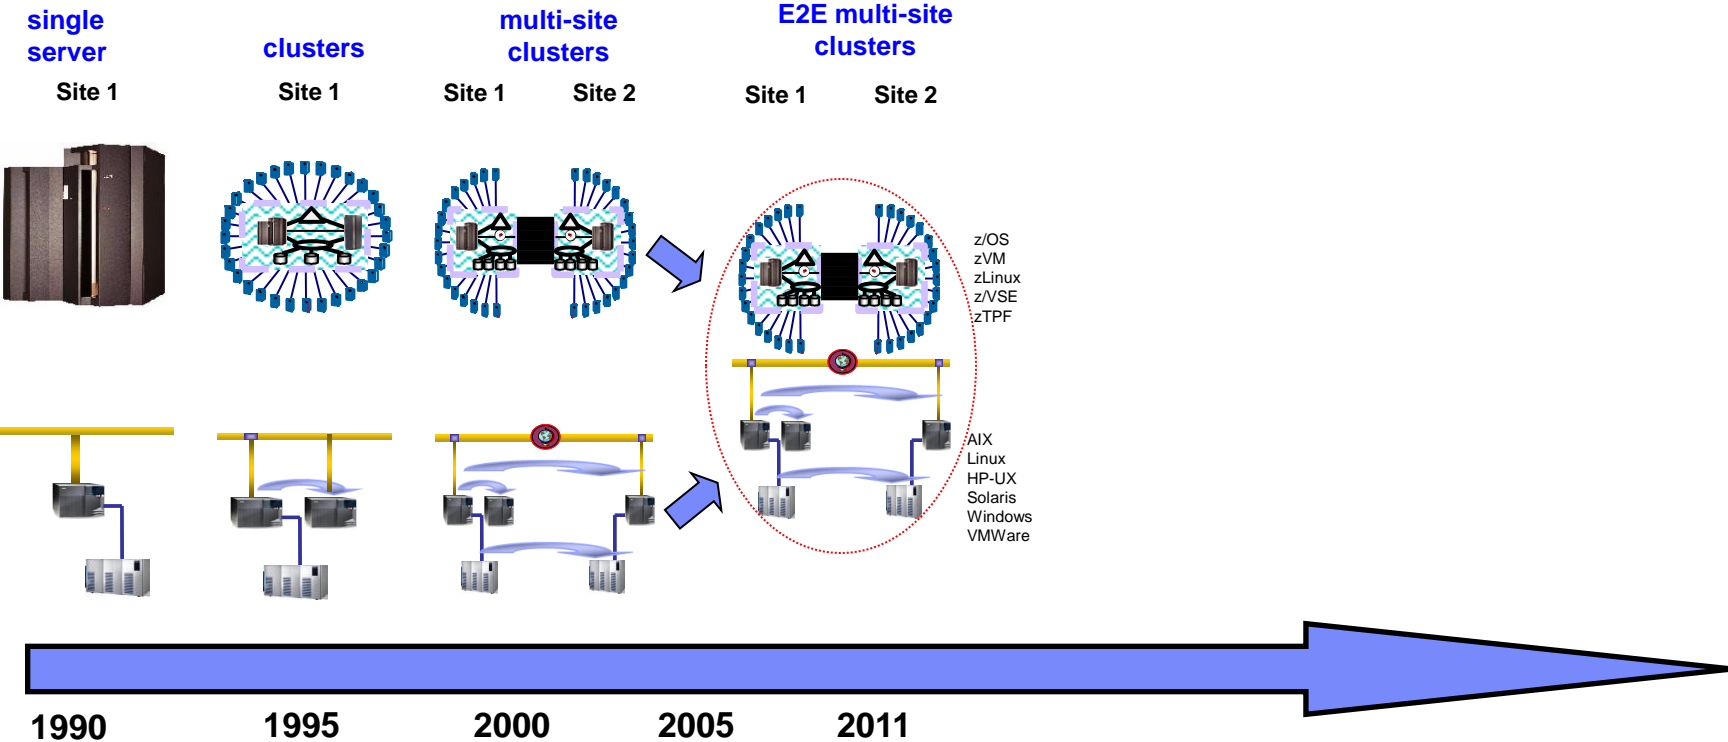
Preliminary Test Results

Components

Summary



Enterprise Wide BC Solution Direction



z/OS
zVM
zLinux
zVSE
zTPF

AIX
Linux
HP-UX
Solaris
Windows
VMWare

- Built In Redundancy
- Capacity Backup
- Hot Pluggable I/O
- Sys z RPO 48-72H / RTO 48-72H
- Provides CA / DR
- Sync/Asynch Data Mirroring
- Eliminates Tape/Disk SPoF
- Application Independent
- Sys z RPO 0-3 sec / RTO < 1H
- Planned/unplanned HW/SW outages
- Flexible, non-disruptive growth
- Dynamic Workload and Resource Management
- Sys z RPO 48-72H / RTO 48-72H
- Provides E2E CA / DR
- Automated failover/switchover
- Browser-based Administration
- E2E RPO 0-3 sec / RTO < 1H

Business Continuity

▪ The top Causes of Business Interruption:

- Planned Maintenance
 - System and Software Upgrades or Reconfiguration
 - Database Administration
- Component Failure
 - Caused by Operator Errors, Software defects, Disk Failure, Subsystems, Hardware, Power Grid
 - Data is recoverable.
 - But, changes might be stranded until component is restored
- Disaster
 - Flood, Earthquake, Fire, ..., Loss of a site
 - Data is not recoverable

Establishing the Objectives:

- *Recovery Time Objective (RTO)* = How much time is needed to restore business operations?
- *Recovery Point Objective (RPO)* = How much data could we afford to *lose*?

How Much Interruption can your Business Tolerate?

Ensuring Business Continuity:

Standby

- **Disaster Recovery**

- Restore business after an unplanned outage

- **High-Availability**

- Meet Service Availability objectives e.g., 99.9% availability or 8.8 hours of down-time a year

- **Continuous Availability**

- No downtime (planned or not)

Active/Active

Global Enterprises that operate across time-zones no longer have any 'off-hours' window. Continuous Availability is required.

What is the cost of 1 hour of downtime during core business hours?

| Cost of Downtime by Industry | |
|------------------------------|----------------------|
| <u>Industry Sector</u> | <u>Loss per Hour</u> |
| 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."

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

September 14, 2010

Outage Curtails Chase's Online Banking

The New York Times

... with 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
- Some industries fine for downtime and inability to meet regulatory compliance
- 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

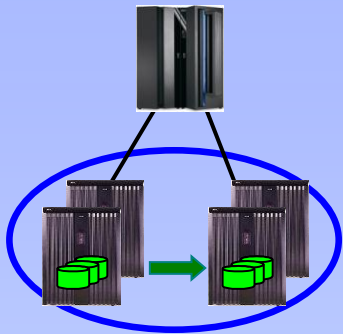
² Continuity Central, "Business Continuity Unwrapped," 2006, <http://www.continuitycentral.com/feature0358.htm>

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

What are customers doing today ?

Continuous Availability of Data within a Data Center


Single Data Center
Applications remain active
Continuous access to data in the event of a storage subsystem outage



GDPS/HyperSwap Mgr
RPO=0 & RTO=0

Continuous Availability / Disaster Recovery within a Metropolitan Region


Two Data Centers
Systems remain active
Multi-site workloads can withstand site and/or storage failures



GDPS/PPRC
RPO=0 & RTO<1 hr

Disaster Recovery at Extended Distance


Two Data Centers
Rapid Systems Disaster Recovery with “seconds” of Data Loss
Disaster recovery for out of region interruptions



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

Continuous Availability Regionally and Disaster Recovery Extended Distance

Three Data Centers
High availability for site disasters
Disaster recovery for regional disasters



GDPS/MGM & GDPS/MzGM

Customer Requirements

- **Want to shift focus from a failover model to a nearly-continuous availability model (RTO near zero)**
- **Access data from any site (unlimited distance between sites)**
- **No application changes**
- **Multi-sysplex, multi-platform solution**
 - “Recover my business rather than my platform technology”
- **Ensure successful recovery via automated processes (similar to GDPS technology today).**
 - Can be handled by less-skilled operators
- **Provide workload distribution between sites (route around failed sites, dynamically select sites based on ability of site to handle additional workload).**
- **Provide application level granularity**
 - Some workloads may require immediate access from every site, other workloads may only need to update other sites every 24 hours (less critical data).
 - Current solutions employ an all-or-nothing approach (complete disk mirroring, requiring extra network capacity).
- **Replace RYO solutions**

What are GDPS/PPRC customers doing today?



- **GDPS/PPRC, based upon a multi-site Parallel Sysplex and synchronous disk replication, is a metro area Continuous Availability (CA), Disaster Recovery solution (DR)**
- **GDPS/PPRC supports two configurations:**
 - Active/standby
 - Active/active
- **Some customers have deployed GDPS/PPRC active/active configurations**
 - All critical data must be PPRCed and HyperSwap enabled
 - All critical CF structures must be duplexed
 - Applications must be parallel sysplex enabled
 - Signal latency will impact OLTP thru-put and batch duration resulting in the sites being separated by no more than a couple tens of KM (fiber)
- ***Issue:* the GDPS/PPRC active/active configuration does not provide enough site separation for some enterprises**

What are GDPS/XRC & GDPS/GM customers doing today ?



- **GDPS/XRC and GDPS/GM, based upon asynchronous disk replication, are unlimited distance DR solutions**
- **The current GDPS async replication products require the failed site's workload to be restarted in the recovery site and this typically will take 30-60 min**
 - Power fail consistency
 - Transaction consistency
- **There are no identified extensions to the existing GDPS async replication products that will allow the RTO to be substantially reduced.**
- ***Issue:* GDPS/XRC and GDPS/GM will not achieve an RTO of seconds being requested by some enterprises**



Improving IT Availability - New Concepts in Delivering Continuous Availability with Active / Active Sites

Level Set

⇒ ***Active/Active Sites Overview***

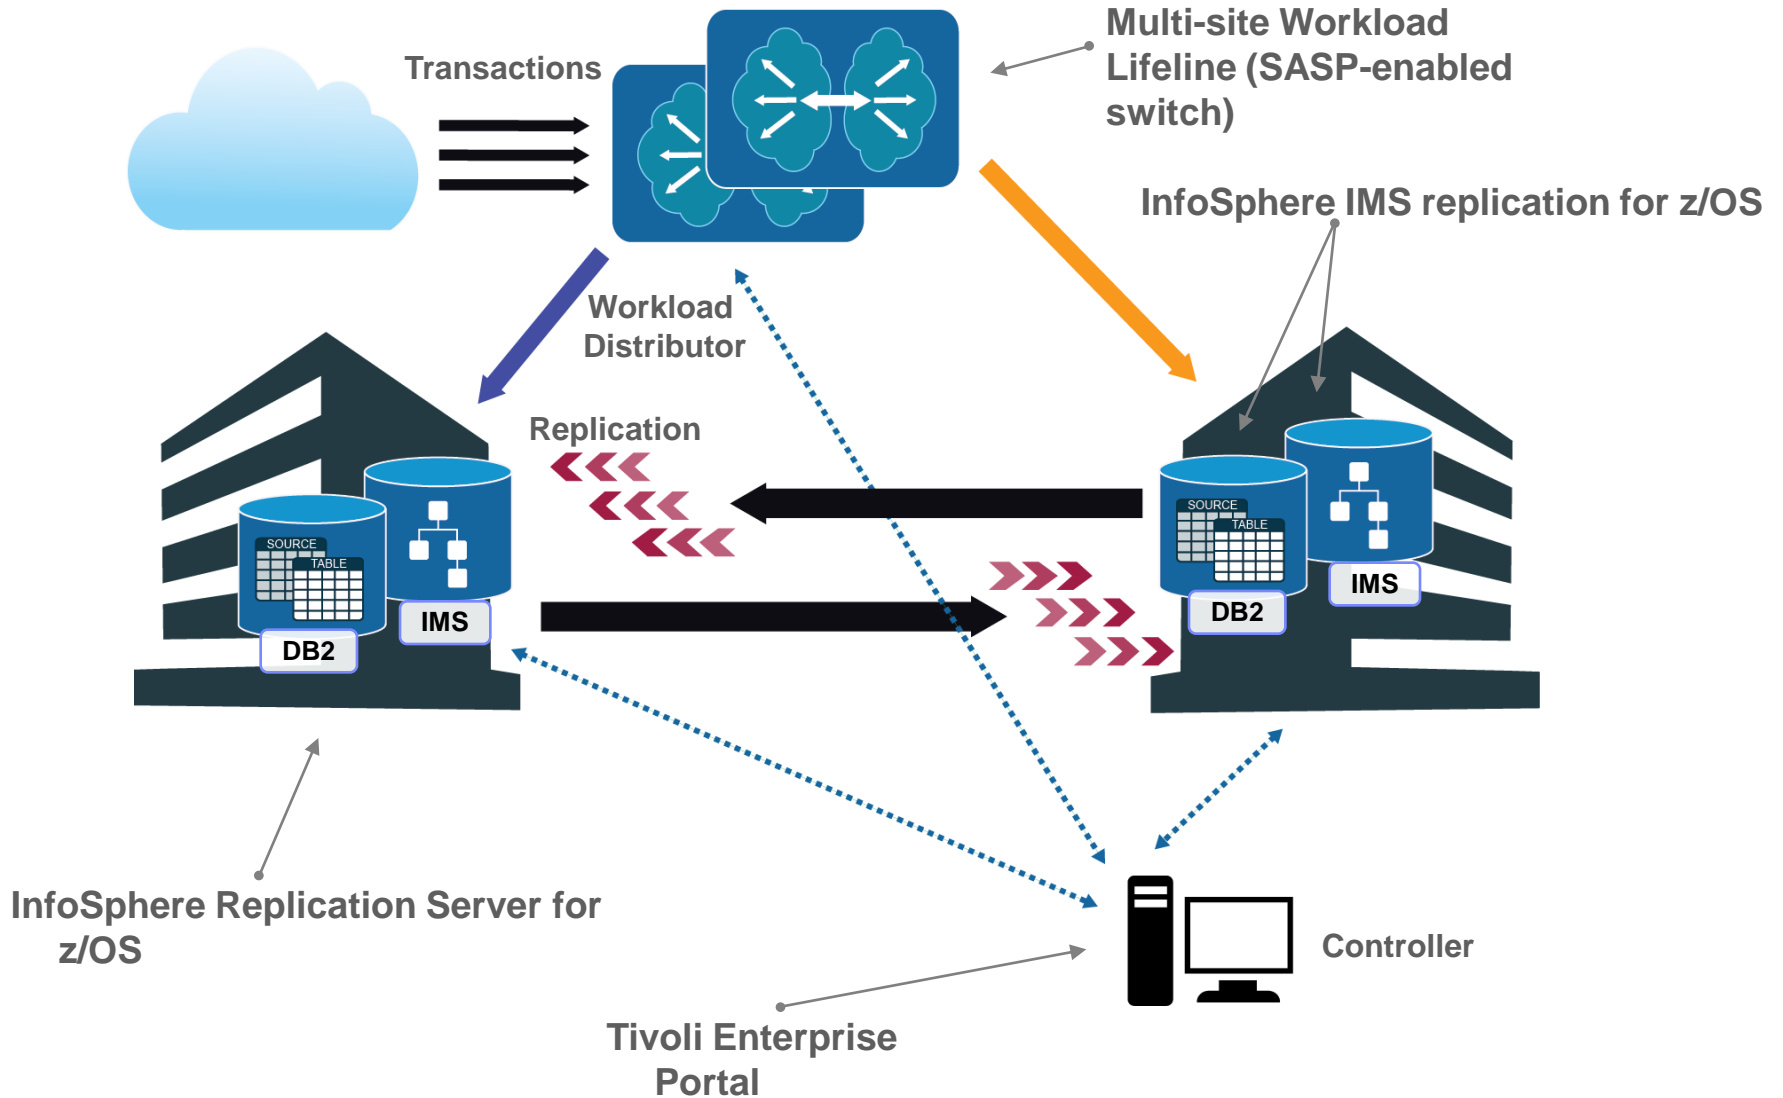
Preliminary Test Results

Components

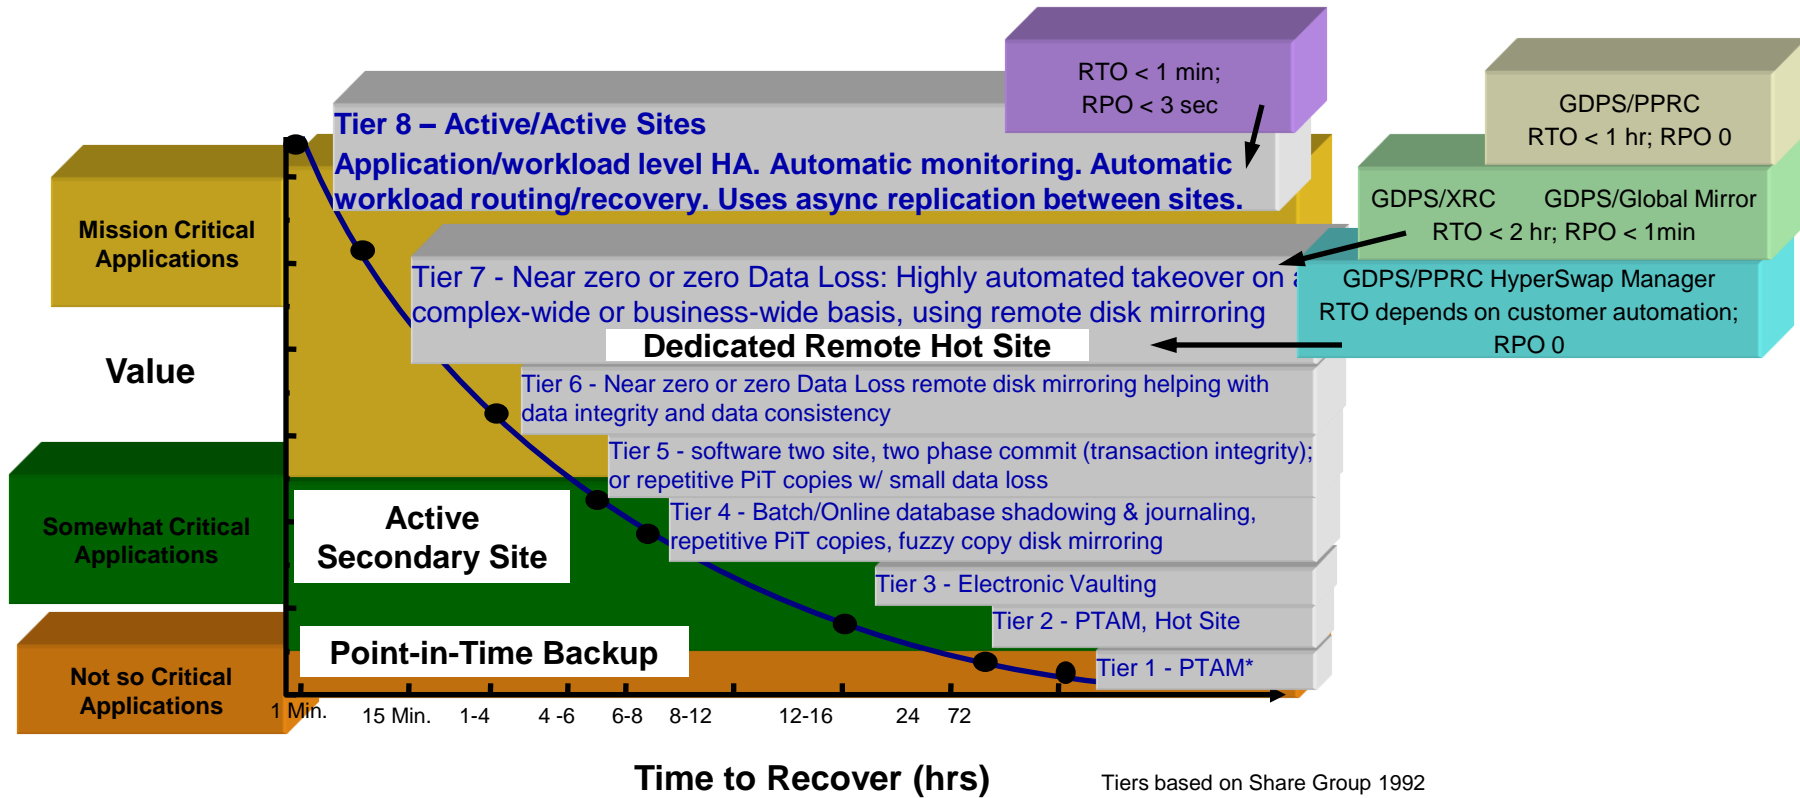
Summary



Active/Active Sites concept



Tiers of disaster recovery: Level-setting Active/Active Sites



Tiers based on Share Group 1992
*PTAM = Pickup Truck Access Method

Best D/R practice is blend tiers of solutions in order to maximize application coverage at lowest possible cost . One size, one technology, or one methodology does not fit all applications

IBM United States Services Announcement 611-023, dated May 24, 2011

- IBM GDPS active/active continuous availability

At a glance

- IBM® GDPS® active/active continuous availability is the next generation of GDPS and represents a fundamental paradigm shift for near continuous availability solutions.

Overview

- **IBM 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 combines the best attributes of the existing suite of GDPS services and expands them to allow you to achieve unlimited distances between your data center sites with recovery time objectives measured in seconds. IBM GDPS active/active continuous availability is a solution for an environment consisting of two sites, separated by unlimited distances, running the same applications and having the same data with cross-site workload monitoring, data replication, and balancing. IBM GDPS active/active continuous availability, as with previous GDPS solutions, provides a complete set of services to help achieve near continuous availability. This solution, which is an integration of IBM products and GDPS control software, is delivered through an IBM service engagement which includes project management throughout the implementation cycle.

Statement of direction

- **IBM intends to deliver, over time, additional configurations that comprise GDPS active/active continuous availability.** In addition to the Active/Standby configuration, IBM plans to make available the Active/Query configuration, which will provide the ability to selectively query data in either site.*

* This statement represents the current intention of IBM. IBM development plans are subject to change or withdrawal without further notice.

Any reliance on this statement of direction is at the relying party's sole risk and does not create any liability or obligation for IBM.

Active/Active Sites configurations

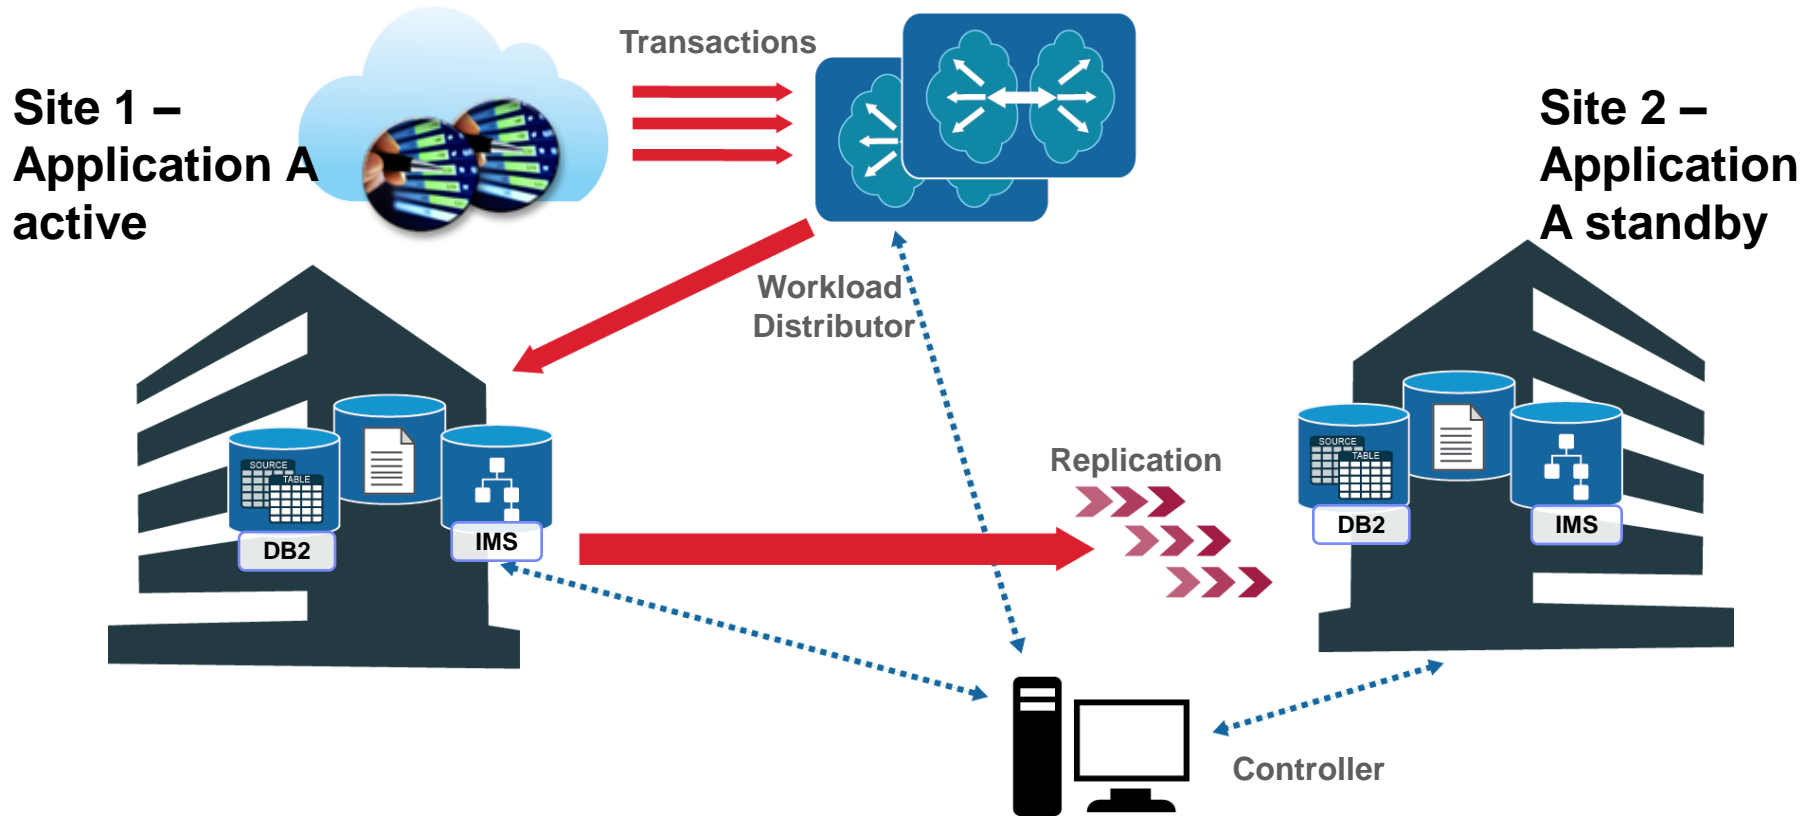
- **Configurations**
 1. Active/Standby – general availability on June 30
 2. Active/Query – statement of direction
 3. ...

- **A configuration is specified on an application basis**

- **An application is the aggregation of these components**
 - **Software:** applications (e.g., COBOL program) and the middleware run-time environment (e.g., CICS region & DB2 subsystem)
 - **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)

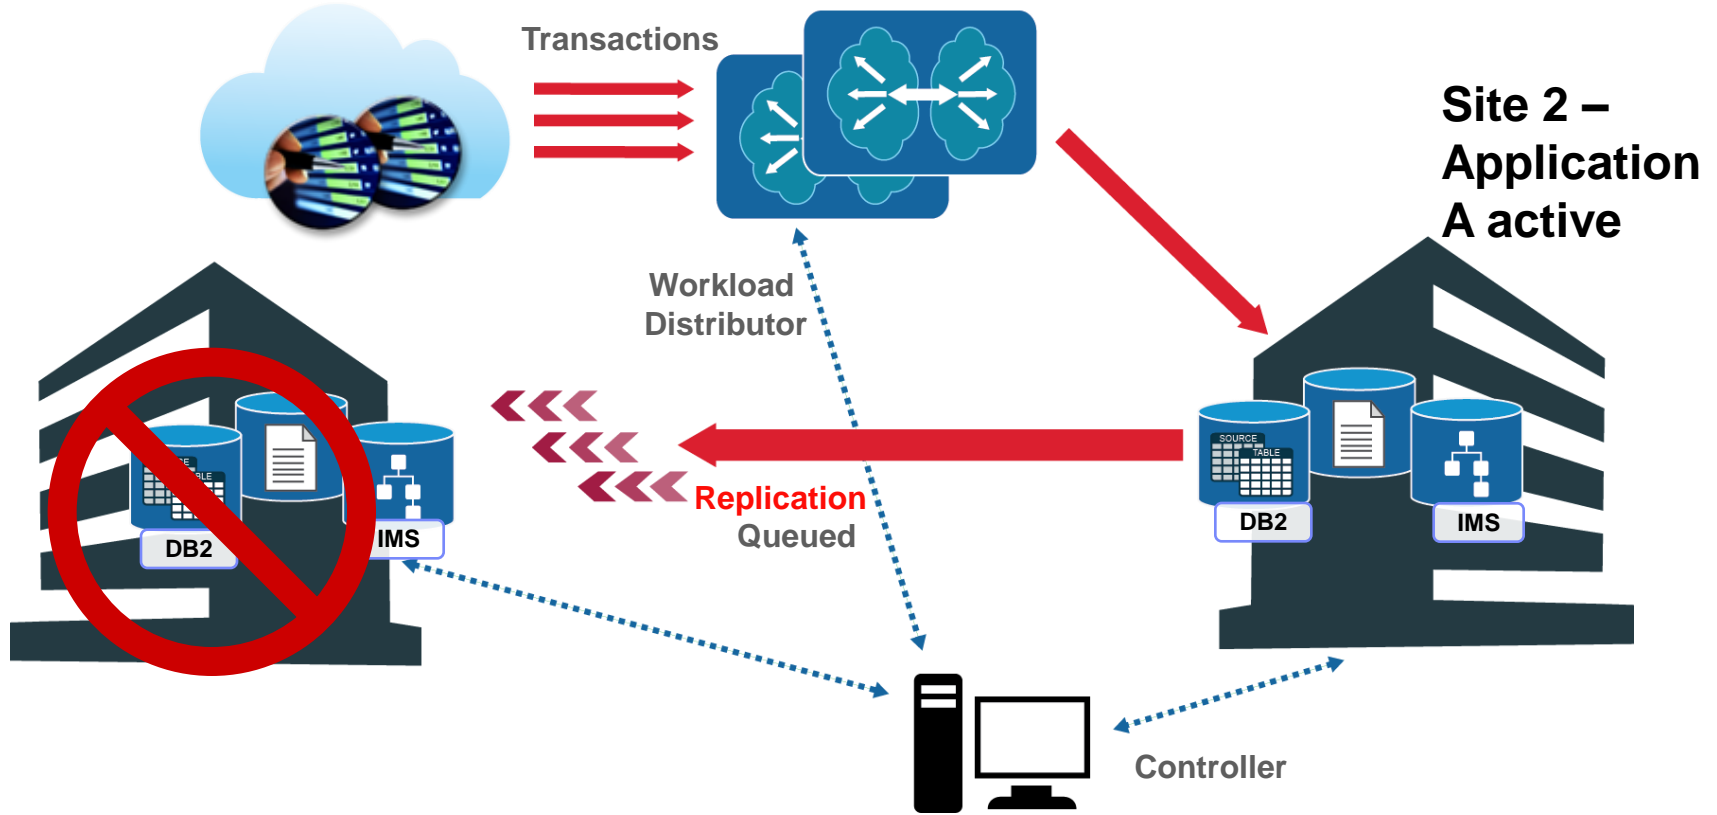
Active/Standby configuration

- Static routing
- Automatic failover



Active/Standby configuration (continued)

- Static routing
- Automatic failover



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 OK, the workload will be switched automatically by changing the routing of transactions 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.**

Active/Active Sites functions

- **Start/stop a controller** – start and stop an A/A Sites controller
- **Start/stop a site** – start and stop individual sysplexes (each sysplex maps to a site)
- **Stop/start a workload** – start and stop individual workloads
- **Monitoring** – monitor the A/A Sites configuration and, if any conditions that will potentially impact a workload and/or site switch, generate an alert
- **Planned workload switch** – switch the workload site to the other site initiated by operator action
- **Unplanned workload switch** – switch failed workload to the other site, either automatically or based upon operator prompt, after the workload failure detection interval
- **Planned site switch** – switch all workload executing to the other site initiated by operator action
- **Unplanned site switch** – switch the failed site workloads to the other site, either automatically or based upon operator prompt, after the site failure detection interval

Improving IT Availability - New Concepts in Delivering Continuous Availability with Active / Active Sites



Level set


Active/Active Sites overview



Preliminary testing results

Components

Summary



Preliminary testing results*

- **Planned workload switch**
 - Operations initiated switch from the active instance of a workload to the standby instance took 20 seconds
 - Not possible with disk replication
- **Unplanned workload switch**
 - Automatic switch from the active instance of a workload to the standby instance took 120 seconds (workload failure detection interval is 60 seconds)
 - Not possible with disk replication
- **Planned site switch (9 * CICS-DB2 and 1 * IMS workloads)**
 - Operations initiated switch of the workloads in a site to the other site took 20 seconds
 - Current GDPS and disk replication will take 1-2 hours
- **Unplanned site switch**
 - Automatic switch of failed site workloads to the surviving site took 107 seconds (site failure detection interval is 60 seconds)
 - Current GDPS and disk replication will take about one hour

* IBM laboratory results; actual results may vary.



Improving IT Availability - New Concepts in Delivering Continuous Availability with Active / Active Sites

Level set


Active/Active Sites overview

Preliminary testing results



Components

Summary



Minimum releases of required products installed on z/OS production and controller images for Active/Standby configuration

▪ Operating system

- z/OS V1R11

▪ Applications/Middleware

- DB2 for z/OS V9
- IMS V10
- WS MQ V7.0

▪ Replication

- InfoSphere Replication Server for z/OS (DB2) V10
- InfoSphere IMS Replication for z/OS V10.1 (new product)

▪ Management and monitoring

- GDPS/Active-Active V1.1 (new product)
- NetView for z/OS V6.1
- System Automation for z/OS V3.3
- IBM Multi-site Workload Lifeline V1.1 (new product)
- IBM Tivoli Monitoring V6.2.2
- Optional OMEGAMON products (required only if the customer wants to monitor the behavior of the respective products/resources that they deal with (DB2, CICS, storage, etc.)
 - OMEGAMON XE on z/OS V4.2.0
 - OMEGAMON XE for Mainframe Networks V4.2.0
 - OMEGAMON XE for Storage V4.2.0
 - OMEGAMON XE for DB2 Performance Expert (or Performance Monitor) on z/OS V4.2.0 (if DB2 is running)
 - OMEGAMON XE on CICS for z/OS V4.2.0 (if CICS is running)
 - OMEGAMON XE on IMS V4.2.0 (if IMS is running)
 - OMEGAMON XE for Messaging V7.0 (if MQ is running)

Improving IT Availability - New Concepts in Delivering Continuous Availability with Active / Active Sites



Level set


Active/Active Sites overview

Preliminary testing results

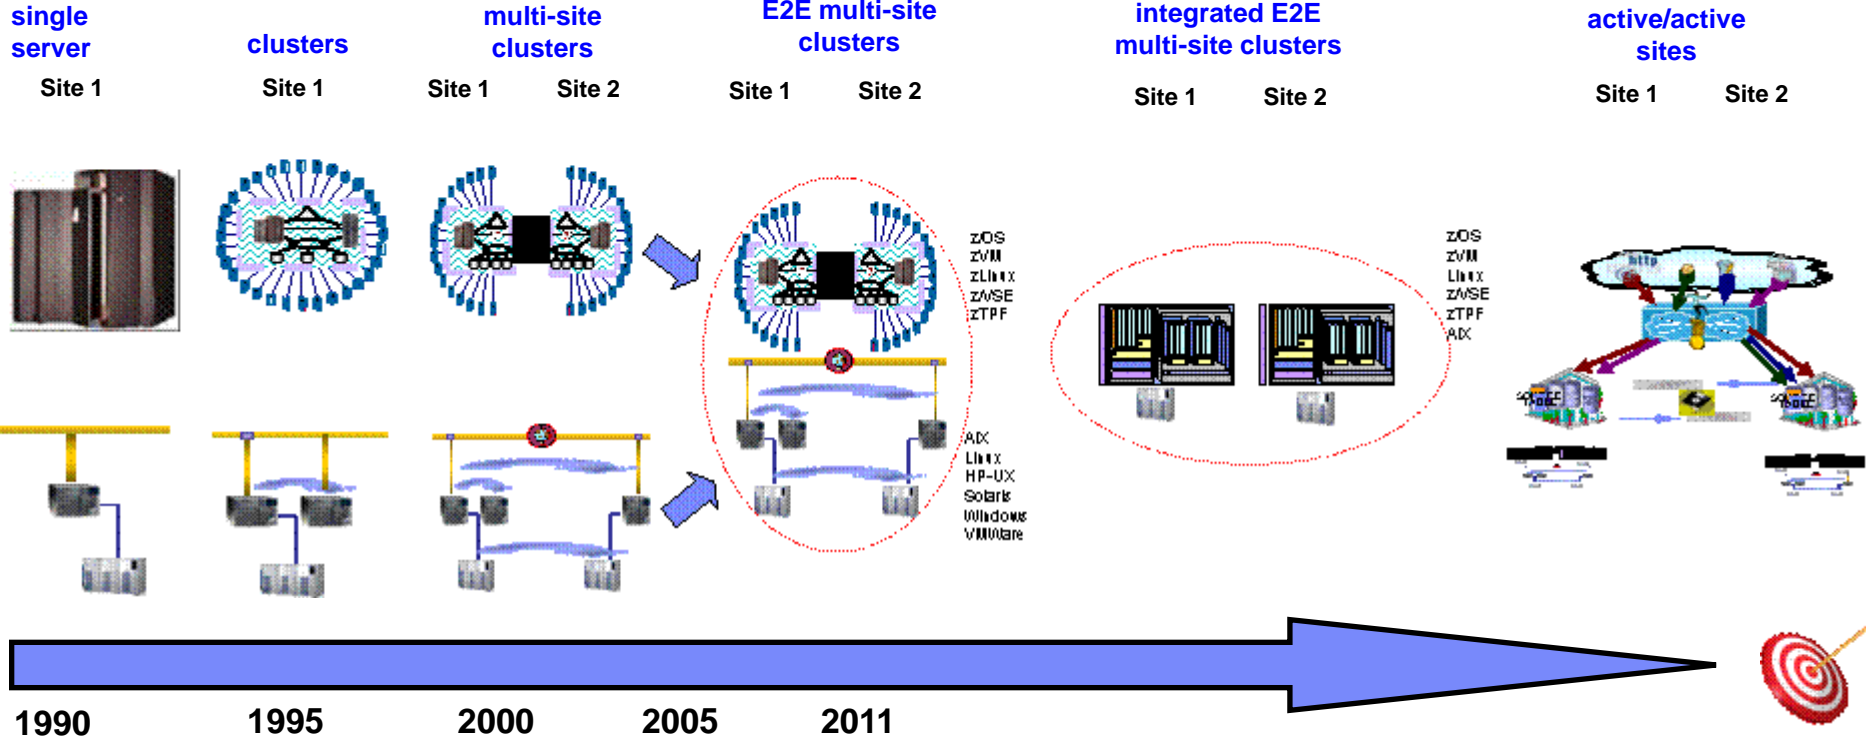
Components



Summary



Enterprise Wide BC Solution Direction



- Built In Redundancy
- Capacity Backup
- Hot Pluggable I/O
- Sys z RPO 48-72H / RTO 48-72H

- Planned/unplanned HW/SW outages
- Flexible, non-disruptive growth
- Dynamic Workload and Resource Management
- Sys z RPO 48-72H / RTO 48-72H

- Provides CA / DR
- Sync/Asynch Data Mirroring
- Eliminates Tape/Disk SPoF
- Application Independent
- Sys z RPO 0-3 sec / RTO < 1H

- Provides E2E CA / DR
- Automated failover/switchover
- Browser-based Administration
- E2E RPO 0-3 sec / RTO < 1H

- E2E integration
- Optimizes floor space & utilities
- Improved systems management
- Single point of control
- Extensive virtualization
- Application level granularity
- E2E RPO 0 sec / RTO < 1H

- All sites active
- Automated workload management
- Capacity on demand
- HA tunable / site
- Unlimited distance between sites
- Continuous data protection
- E2E RPO 0 sec / RTO 30 sec

Summary

- **Current GDPS family of offers**

- Over 12-year history of disaster recovery and continuous availability for System z customers
- A proven track record of success, with almost 600 clients worldwide and growing
- Ongoing investment and updated – up to GDPS V3.8

- **NEW GDPS/ Active-Active family of offers**

- The next generation of GDPS
- Concept: Active applications, transactional integrity, shared data, replication, and automation over global distances for true continuous availability worldwide
 - First configuration is Active/Standby
 - Statement of Direction on Active Query configuration
 - Additional configurations planned for the future



Thank
YOU