IBM Resiliency Services
Always there, in an always-on world

IBM GDPS Virtual Appliance Introduction

Disaster Recovery for Linux on z Systems

Trademarks

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

CICS*	Geographically Dispersed Parallel Sysplex	IMS	System Storage*	zSeries*
DB2*	HyperSwap	MQSeries*	System z*	z Systems
DS8000*	IBM*	NetView*	Tivoli*	
FlashCopy*	IBM eServer	OMEGAMON*	WebSohere*	
GDPS*	InfoSphere	Parallel Sysplex*	z/OS*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Windows Server and the Windows logo are trademarks of the Microsoft group of countries.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

* Other product and service names might be trademarks of IBM or other 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.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Agenda

- **The CA/DR Requirement and Challenges**
- **IBM® GDPS® Virtual Appliance Overview and Benefit**
- **Basic Operation Introduction**

IT Infrastructure Service Availability Building Blocks

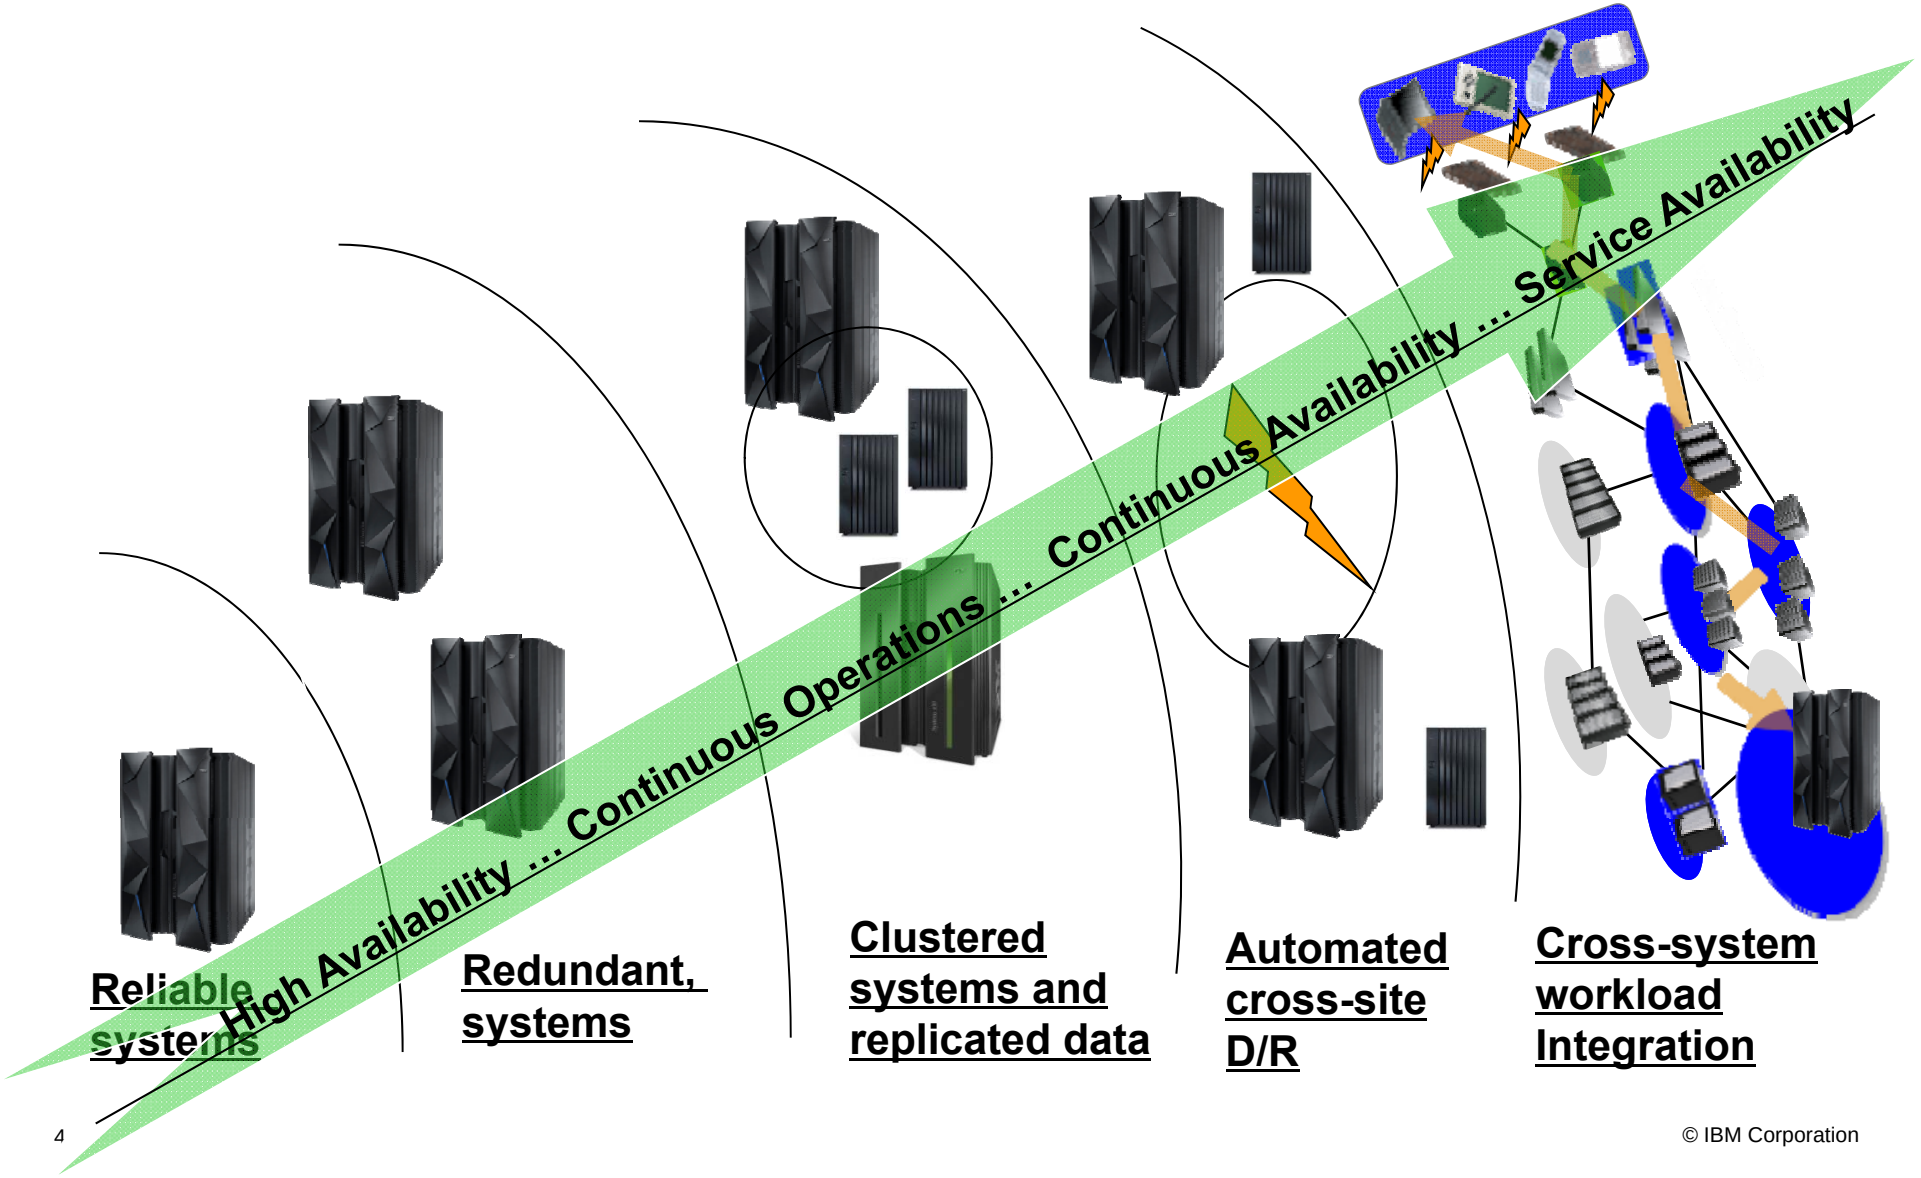

Strive to provide Continuous Availability

- **High Availability (HA)–** The attribute of a system to provide service during defined periods, at acceptable or agreed upon levels and masks *unplanned* outages from end-users. It employs Fault Tolerance; Automated Failure Detection, Retry, Bypass, Reconfiguration.
- **Continuous Operations (CO)-** Attribute of a system to continuously operate and mask *planned* outages from end-users. It employs Non-disruptive hardware and software changes, non-disruptive configuration, software coexistence.
- **Continuous Availability (CA)-** Attribute of a system to deliver non disruptive service to the end user 7 days a week, 24 hours a day (there are no planned or unplanned outages).

What are the reasons for system outage

Non-Disaster Events

- **Typical planned outages:**
  - Backups
  - OS installs and maintenance
  - Application maintenance
  - Hardware/Software upgrades

- **Typical unplanned, non-disaster outages:**
  - Application failure
  - Operator error
  - Local power outages
  - Network failure
  - Hardware failure

Disaster Events

- **Typical disaster outages:**
  - Outages that are caused by natural disasters or other catastrophes that damage the production facilities beyond usability (e.g – fire, flood, earthquake, bombing)
  - Failure of regional power grid
  - Outages not caused by computer hardware or software defects
  - Outages that require a recovery procedure at an off-site location

- **A disaster is a rare event that most customers will never experience but must plan for it.**

- **In fair weather prepare for foul**

Disaster Recovery

- **Disaster Recovery (DR) -** Given a major disruption caused by a disaster, fully recover essential data and operations. DR encompasses the ability of the total design to transfer data and workload offsite and to restart the workload at a new site. The switch often involves a system and service outage but in the extreme the switch can be made completely seamless to important users.

- Two objectives (RTO and RPO) are common measurements for DR, and can often be balanced against each other to optimize the cost/benefit ratio.
  - **Recovery Point Objective (RPO)** defines the amount of data that you can afford to recreate during a recovery, by determining the most recent point in time for data recovery.
  - **Recovery Time Objective (RTO)** is the time needed to recover from a disaster or how long the business can survive without the systems.

Drivers for improvements in CA/DR

- **Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System [Docket No. R-1128] (April 7, 2003)**
  - Focus on mission critical workloads, their recovery and resumption of normal processing
- **Cost of an outage**
  - Financial
  - Reputation
- **Global Business Model**
  - 24x365 processing
  - Planned outage avoidance

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."

CA/DR challenges for Data Centers

Complexity of datacenters through:

1
- Heterogeneous, multi-tiered applications
- Hundreds of heterogeneous servers (platform,OS, …)
- Heterogeneous storage landscape

4 Separation of application/server management and storage management (data replication)

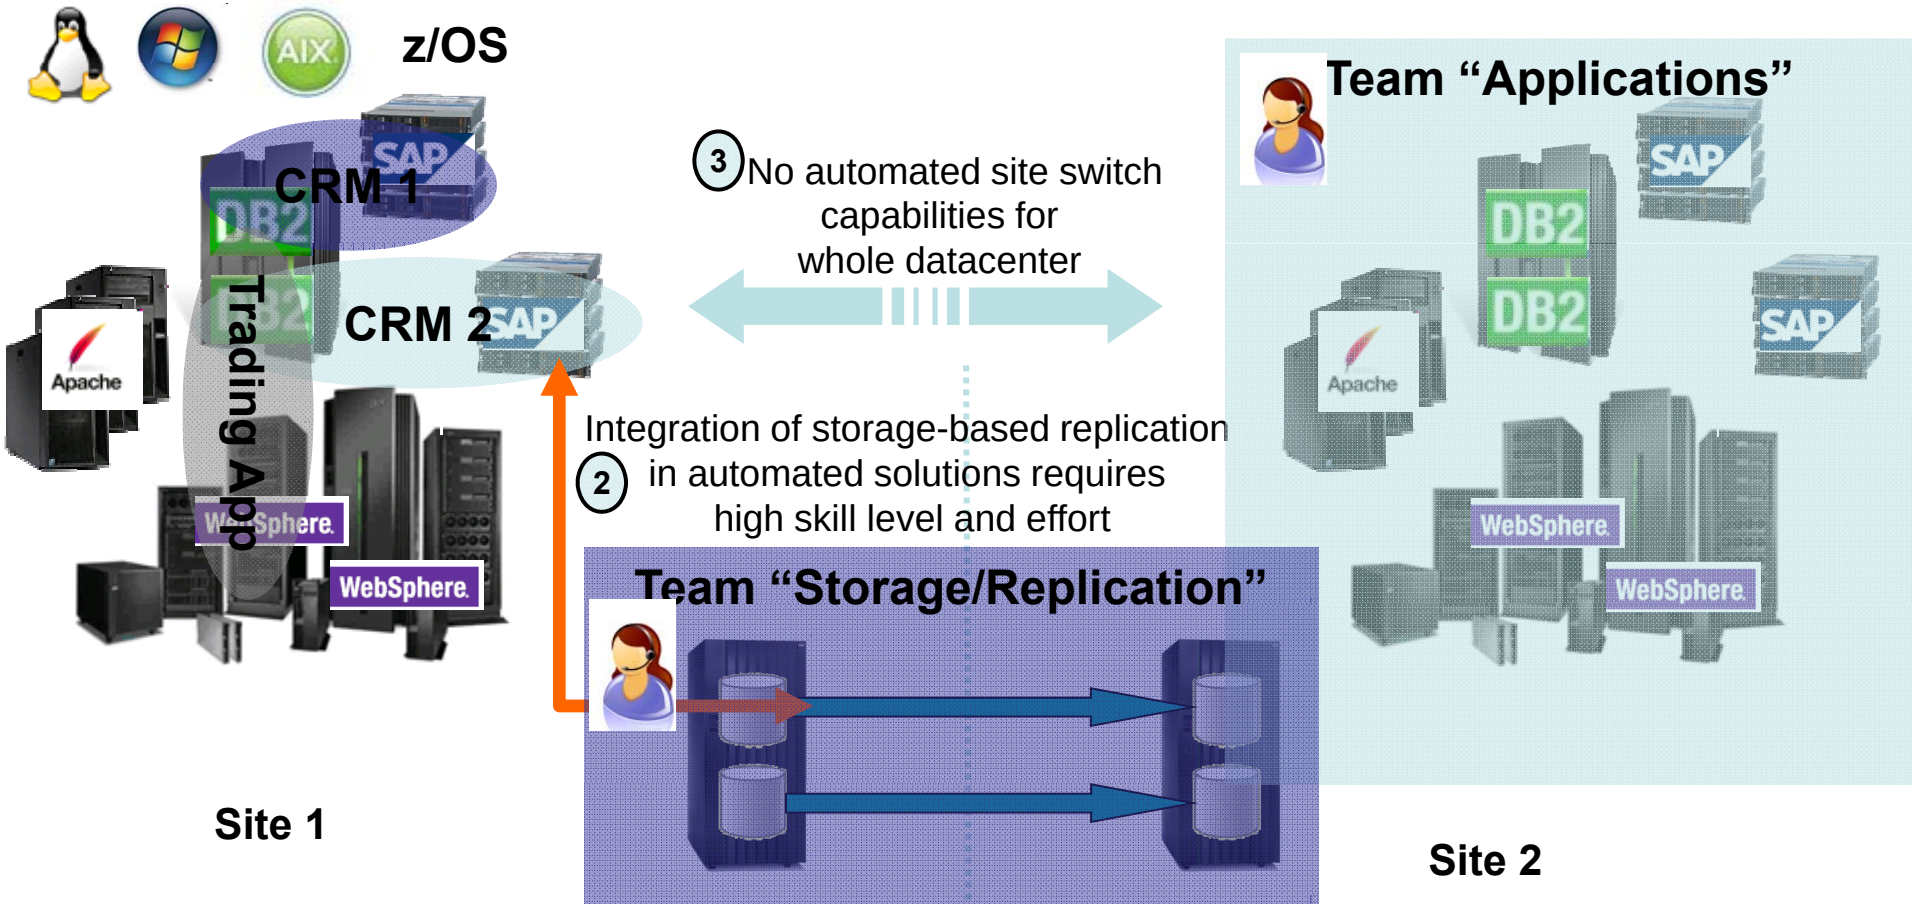

Automation: Critical for successful CA and DR

- **The benefits of automation:**
  - Reliable, consistent recovery time
  - Consistent recovery times as environment scales
  - Reduces infrastructure management cost and staffing skills
  - Reduces or eliminates human intervention and human error
  - Facilitates regular testing for repeatable, reliable, scalable business continuity
  - Helps maintain recovery readiness by managing/monitoring server, data replication, workload and network with the notification of events that occur within the environment

Automate – Automate – Automate

Agenda

- The CA/DR Requirement and Challenges
- IBM GDPS Virtual Appliance Overview and Benefit
- Basic Operation Introduction

GDPS solutions aimed at CA/DR

- GDPS (Geographically Dispersed Parallel Sysplex) is a very successful GTS Service offering that provides multiple solutions to meet various customer requirements for Continuous Availability and Disaster Recovery

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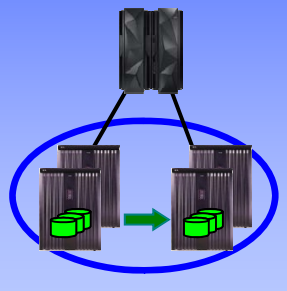 RPO=0 & RTO=0	Multi-site workloads can withstand site and/or storage failures 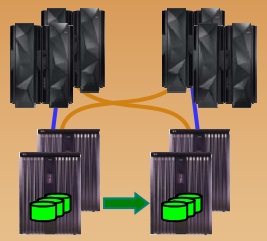 RPO=0 RTO mins / RTO <1 hr (<20km) (>20km)	Two Data Centers Rapid Systems Disaster Recovery with "seconds" of Data Loss Disaster recovery for out of region interruptions 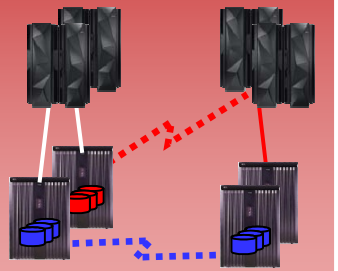 RPO secs, RTO <1 hr	Three Data Centers High availability for site disasters Disaster recovery for regional disasters 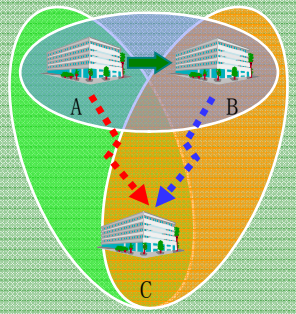 RPO=0, RTO mins/<1 hr & RPO secs, RTO <1 hr	Two or More Data Centers All sites active 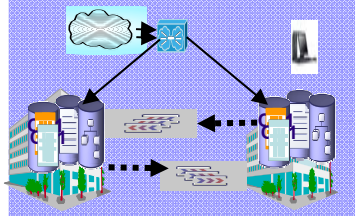 RPO secs, RTO secs

GDPS/PPRC Experience in Linux on z

Stockholm, Sweden
[Finance]

6-way z/OS Parallel Sysplex & 2 z/VM Clusters (30 Linux Guests)

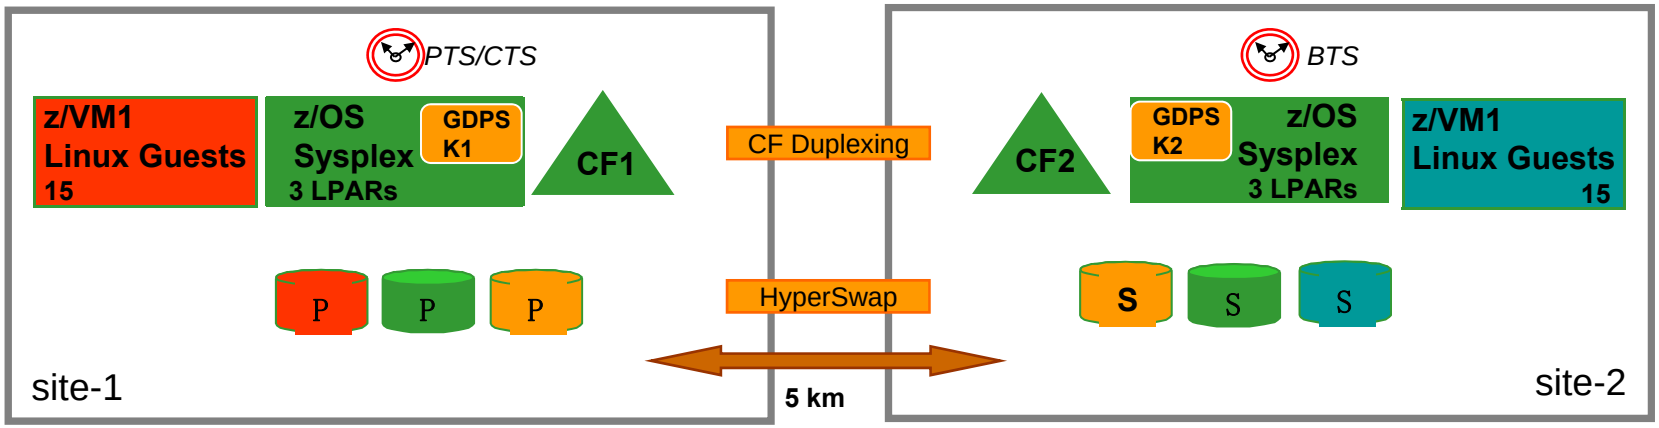

PRIMARYFAILURE = SWAP,GO
PPRCFAILURE = COND

z/OS 1.13, NV 6.1, SA 3.3, GDPS 3.8, z/VM 6.1
Linux RHEL 5.2, SA MP 3.2.1.2

Business Requirements:

- No data loss (RPO 0 seconds)
- Continuous data availability for z/OS® and Linux guests hosted by z/VM® Supporting site maintenance without application outage
- Coordinated D/R for heterogeneous z Systems™ applications (RTO < 2hr)

Web site: www.handelsbanken.se

Perceived benefits

- DR site failure simulation every 6 months
- Avoiding outages due to single component failures or compound failure

PPRC Volume Pairs	Logical Subsystems (LSS)	Planned HS Suspend UIT	Unplanned HyperSwap UIT
1,587 (71TB)	28	15 sec	10 sec

UIT = User Impact Time (seconds)

Solution Description: 'IBM GDPS Virtual Appliance'

- **GDPS Virtual Appliance is a fully integrated software solution that provides Continuous Availability & Disaster Recovery functions for Linux on z Systems customers in conjunction with the GDPS/PPRC xDR facility running on those systems**
  - **It is an image comprising of an operating system, the application components, an appliance management layer which makes the image self-containing, and APIs / UIs for customization, administration, and operation tailored to the appliance function.**
  - **It improves both usability and time-to-value for customers.**

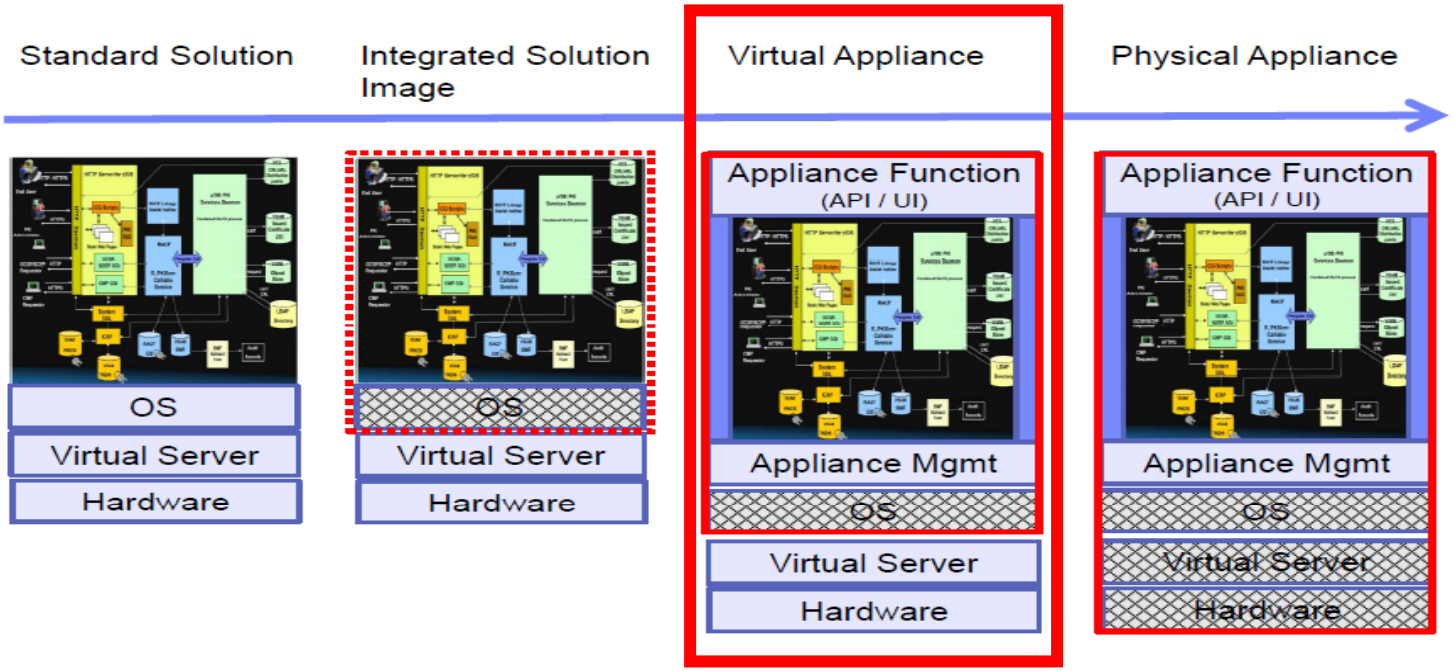

GDPS Virtual Appliance extends GDPS capabilities into z/VM, Linux and z/VSE on z Systems environments that do not have z/OS.

GDPS Virtual Appliance features:

- Single point of control and automation reduces the need for highly specialized skills to handle recovery and planned site switches
- Manages remote copy environment and keeps data available and consistent for operating systems and applications.
- HyperSwap® function protects against failures to disk subsystems.
- Monitoring and automation to ensure reliable and rapid recovery via automated processes
- Virtual Appliance requires:
  - General purpose engine
  - z/VM and Linux on z Systems
  - ECKD Disk

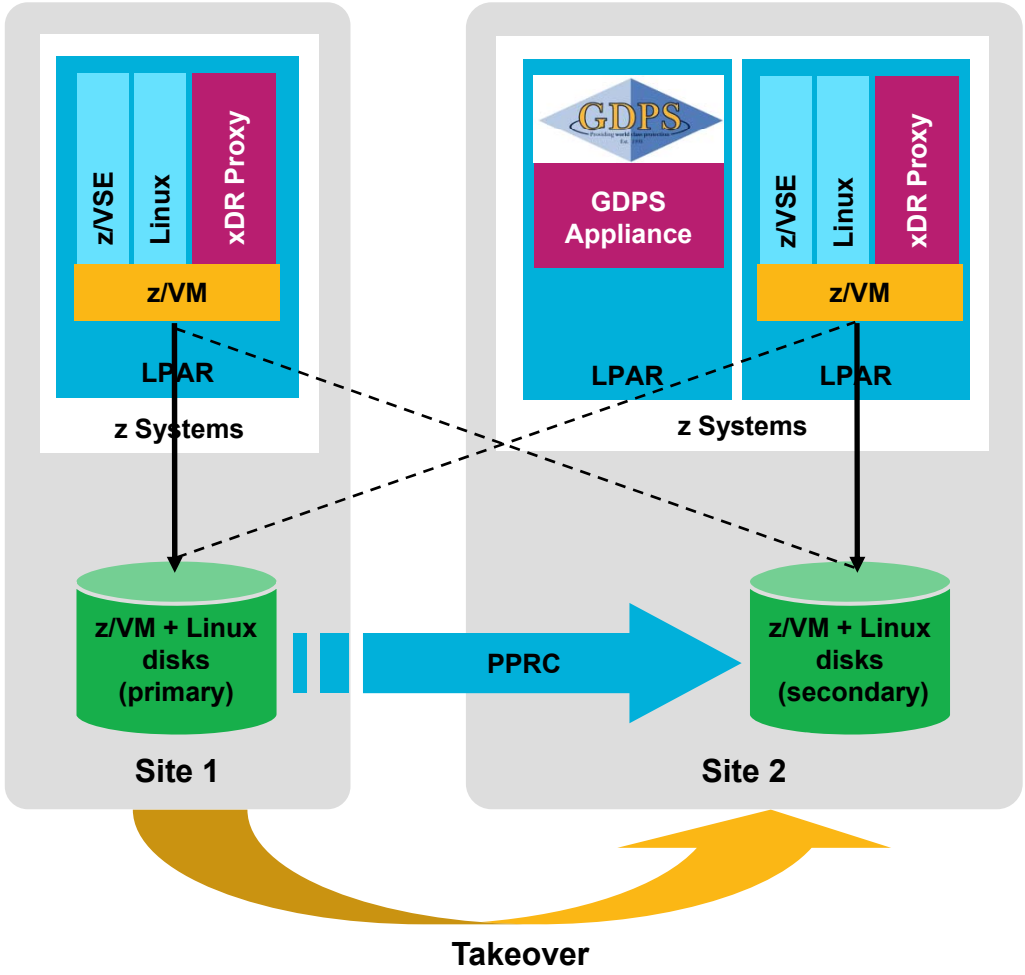

Planned Disk Outage with GDPS Virtual Appliance

- HyperSwap provides the ability to non-disruptively swap from using the primary volume of a mirrored pair to using what had been the secondary volume.

- A planned HyperSwap is invoked manually by operator action using GDPS facilities. One example of a planned HyperSwap is where a HyperSwap is initiated in advance of planned disruptive maintenance to a disk subsystem.

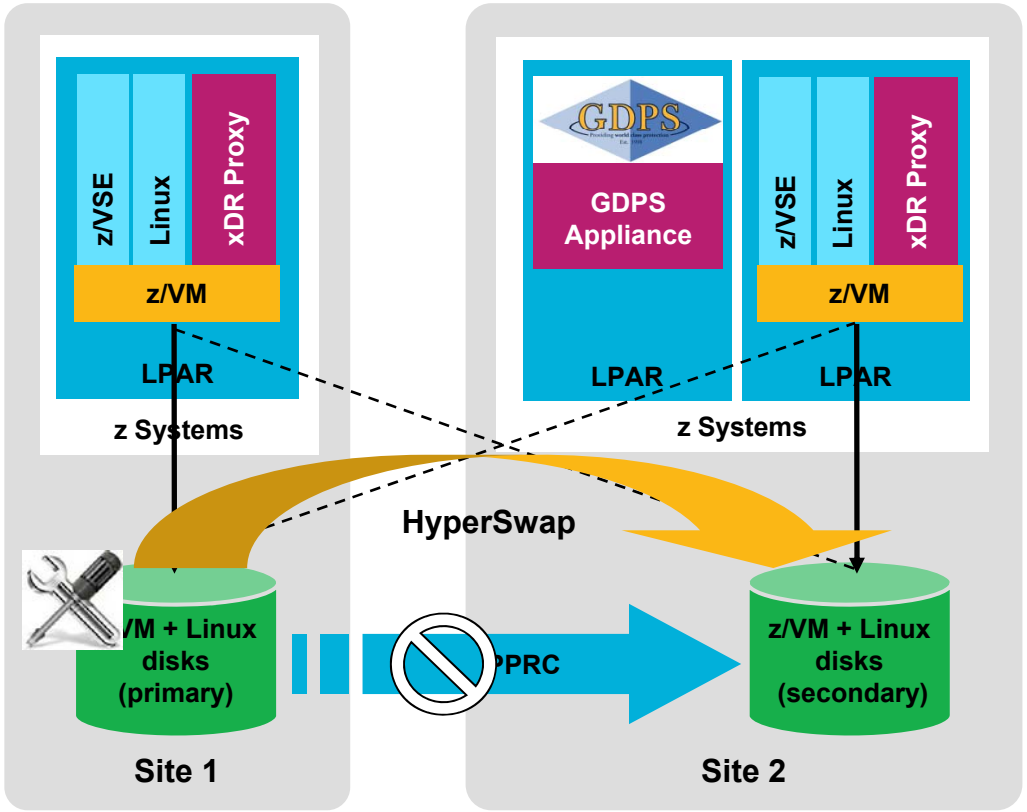

Without HyperSwap

Planed Action Initiated

Shutdown systems, remove systems from cluster, reverse PPRC (suspend PPRC), restart systems

1-2 hrs (approx)

With HyperSwap

14 systems - 10K volume pairs = 6 secs UIT

PPRC Failover, swap the primary & secondary PPRC UCBs, systems remain active

Unplanned Disk Outage with GDPS Virtual Appliance

- An unplanned HyperSwap is invoked automatically by GDPS, triggered by events that indicate the failure of a primary disk device.

- HyperSwap Triggers
  - Hard Failure Triggers
    - I/O errors
    - Boxed devices
    - Control Unit failures
    - Loss of all channel paths
  - "Soft Failure"–I/O response time triggers

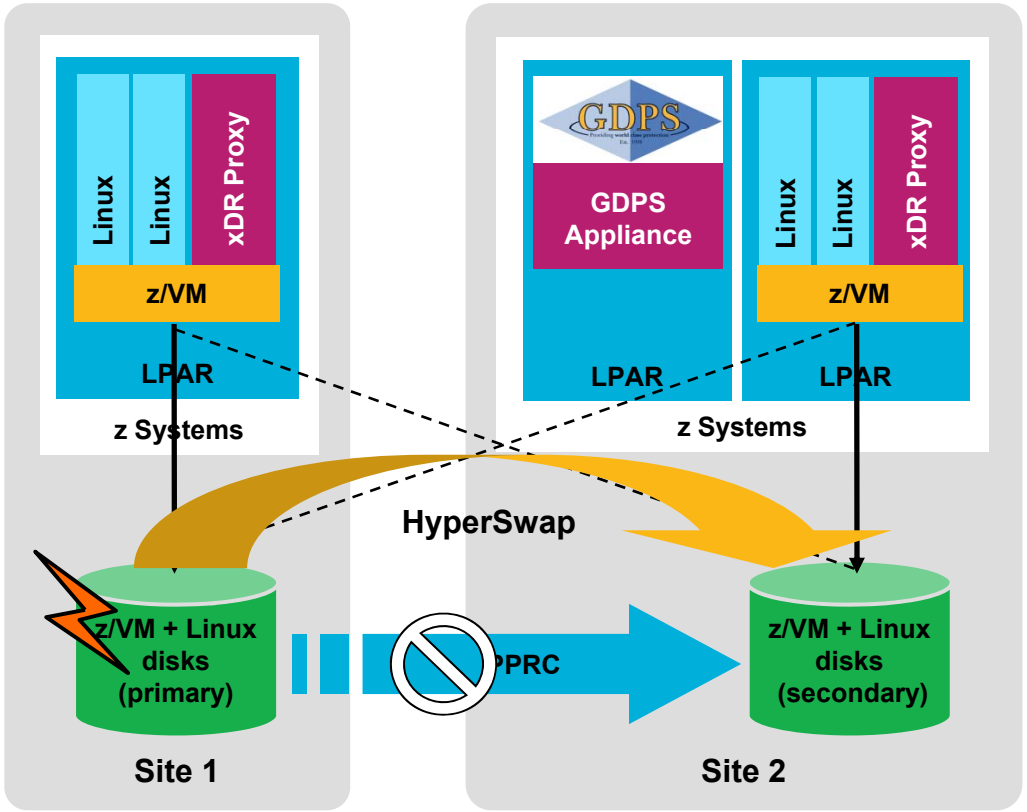

Without HyperSwap

Takeover Prompt → Systems quiesced

Decision Time //

Takeover Started → Remove systems from cluster, restart systems → *>2 hours or more*

With HyperSwap

14 systems - 10K volume pairs = 6 secs UIT

PPRC Failover, swap the primary & secondary PPRC UCBs, systems remain active

Disaster Recovery with GDPS Virtual Appliance

- In case of a site disaster, GDPS Virtual Appliance will immediately issue a freeze for all applicable primary devices. This is done to protect the integrity of the secondary data.
- It will Reset Site1 and Site2 systems and update all the IPL information to point to what were the secondary devices, and re-IPL all the production systems in LPARs in Site 2.
- The GDPS Virtual Appliance scripting capability is key to recovering the systems in the shortest possible time following a disaster. All of this is carried out with a single operator instruction.

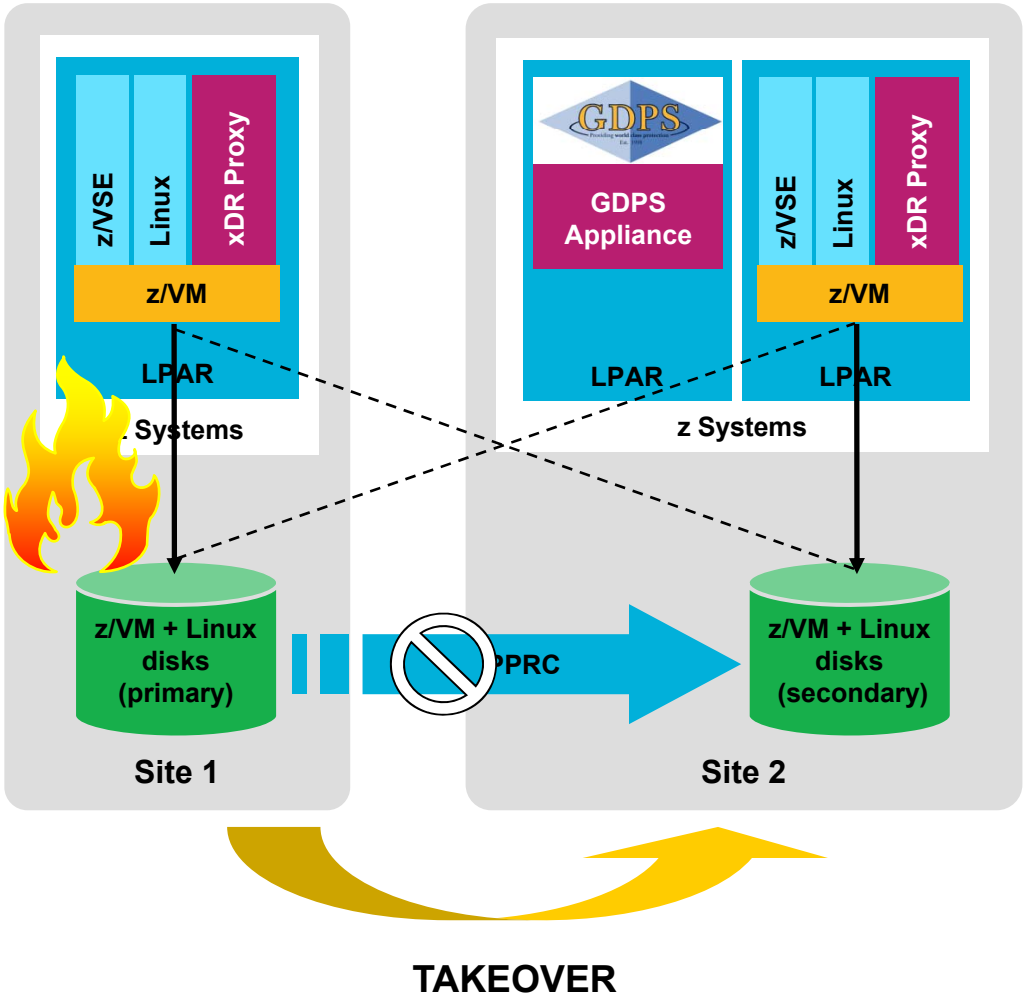

Tiers of Disaster Recovery

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.

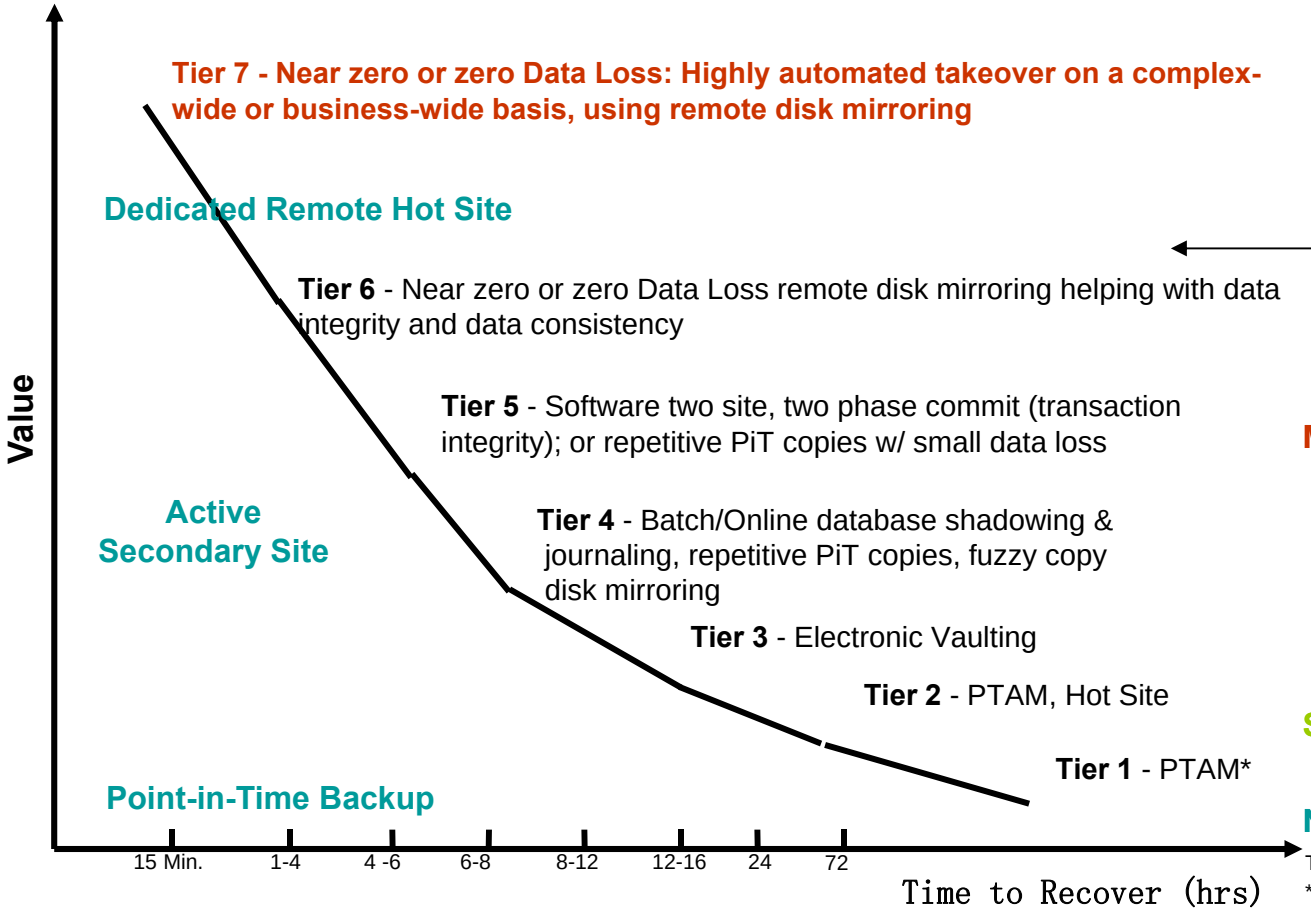

GDPS Virtual Appliance
RTO < 1 hr; RPO 0

- RTO does not include decision time.
- Variable Subsystem specific recovery time

Mission critical application

Somewhat critical application

Not so critical application

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

Summary of Requirements:

- **IBM GDPS Virtual Appliance hardware requirements:**
  - 1 LPAR with one single logical CP
  - 2 GB Memory
  - 4 3390 volumes
  - 1 OSA attachment

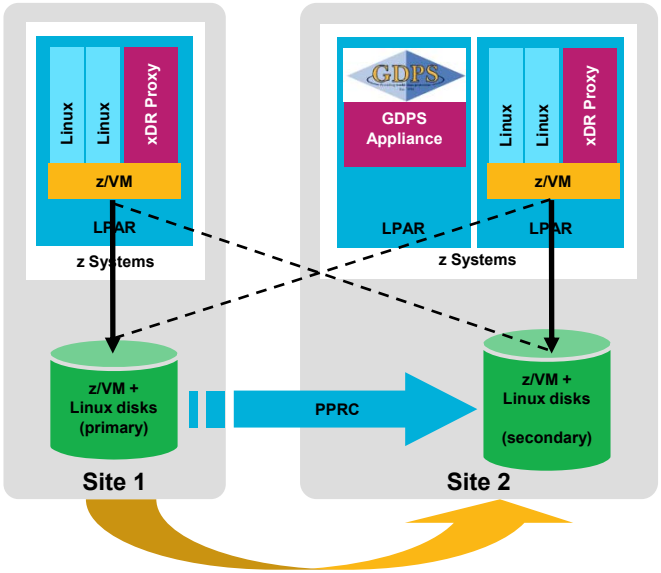

- **Managed system prerequisites:**
  - The disks being used by z/VM and Linux on z Systems which are to be mirrored must be ECKD disks on disk subsystems that support Metro Mirror.
  - z/VM Version 5.4 or z/VM Version 6.2 or higher.
  - A supported distribution of Linux on z Systems with the latest recommended fixpack.
  - IBM Tivoli System Automation for Multiplatforms with the latest recommended fixpack. The separately priced xDR for Linux feature is required.
    - SA MP with xDR feature required on two Linux on z Systems proxy guests per managed z/VM system.
    - SA MP with xDR feature optional on other Linux on z Systems application guests for GDPS management and monitoing of application guest availability.

Agenda

- The CA/DR Requirement and Challenges
- IBM GDPS Virtual Appliance Overview and Benefit
- Basic Operation Introduction

GDPS Virtual Appliance Operation

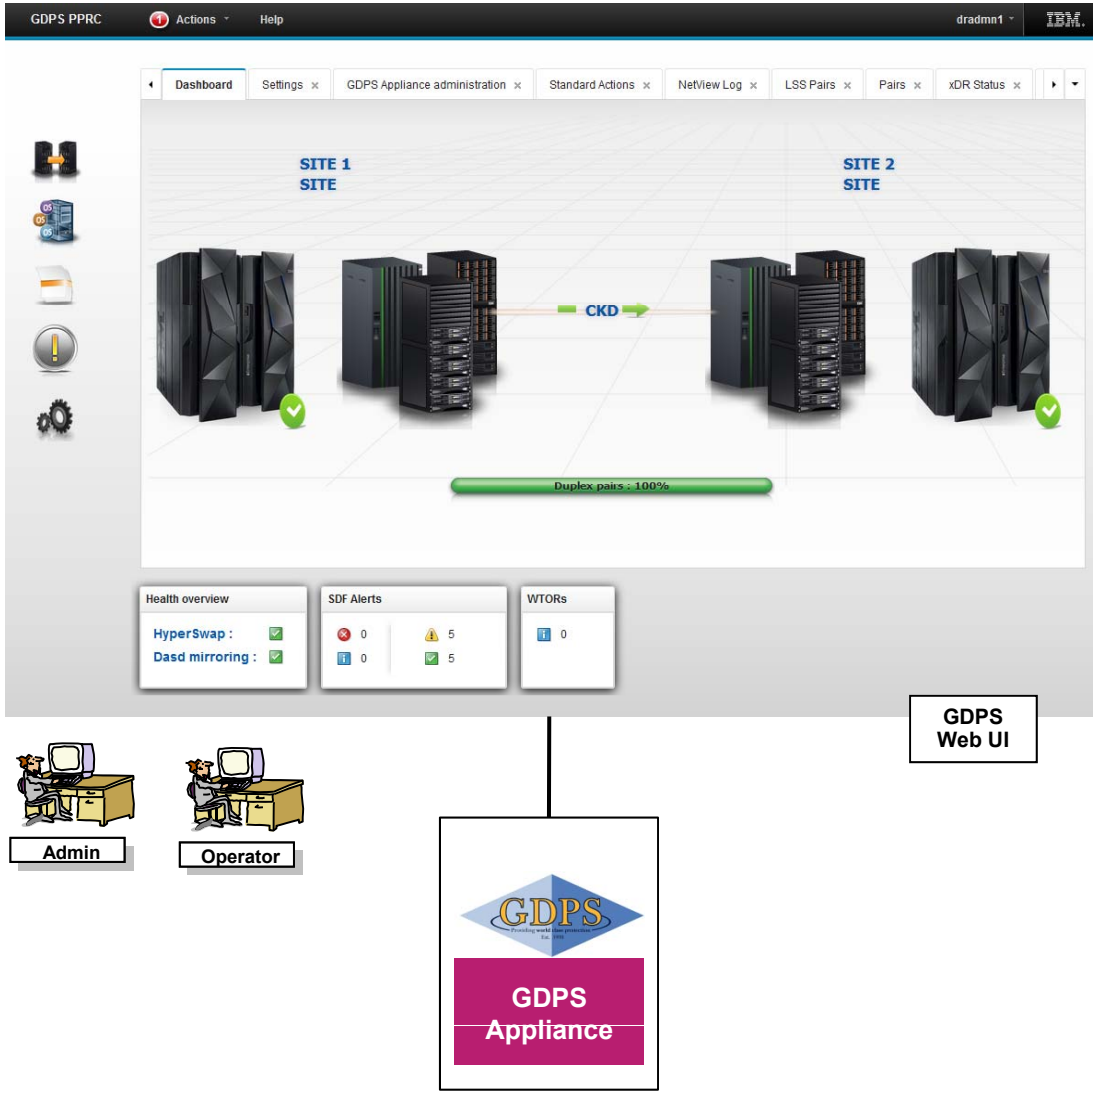

- Pre-configured self-containing image with GDPS K-System
- Customization
  - GDPS customization by IBM Service
- Operation and Administration
  - GDPS Web UI used for operation, administration and problem determination

Note: Only ECKD disks are supported

Initial Installation

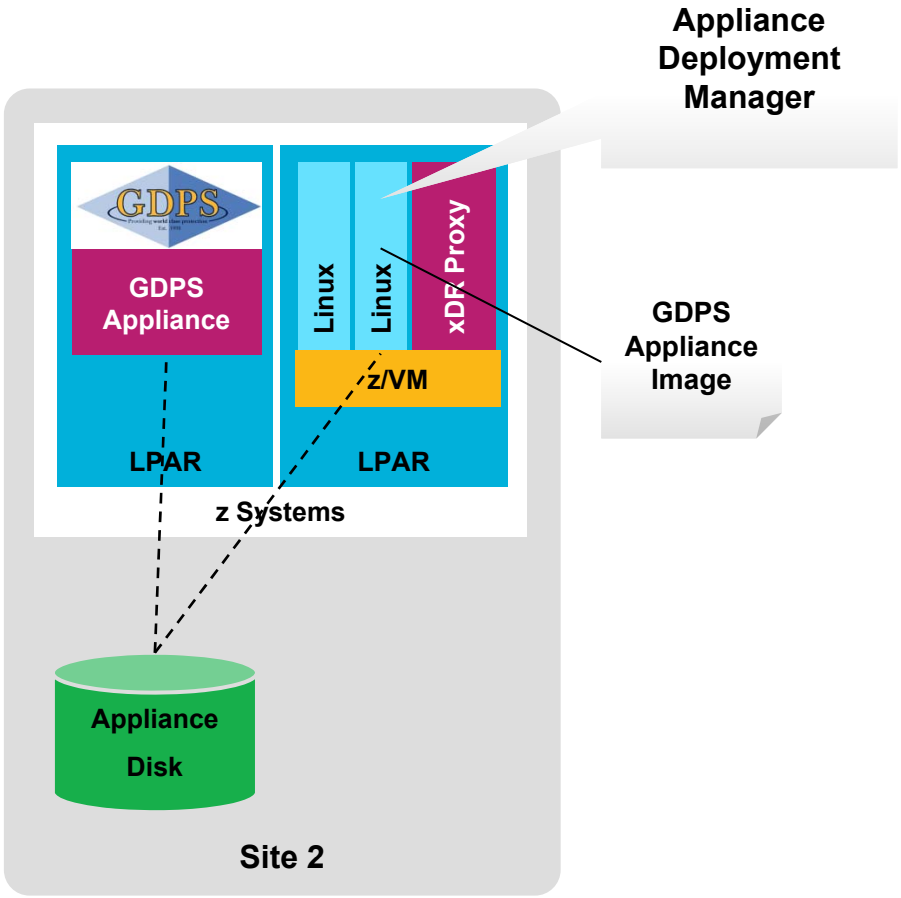

1. Setup GDPS Virtual Appliance LPAR meeting appliance requirements
   – Access to DASD (ECKD) device from which appliance should be IPLed
2. Install Deployment Manager on Linux z/VM guest or native in LPAR with access to DASD device
3. Copy appliance image on disk accessible by Linux
4. Run Deployment Manager
   -to customize TCPIP and timezone in appliance image
   -to copy customized image to DASD device (Appliance Disk)
5. IPL LPAR from Appliance Disk
6. GDPS Customization by IBM Service

Time to deploy: minutes vs days

GDPS Virtual Appliance Upgrade (version n → n+1)

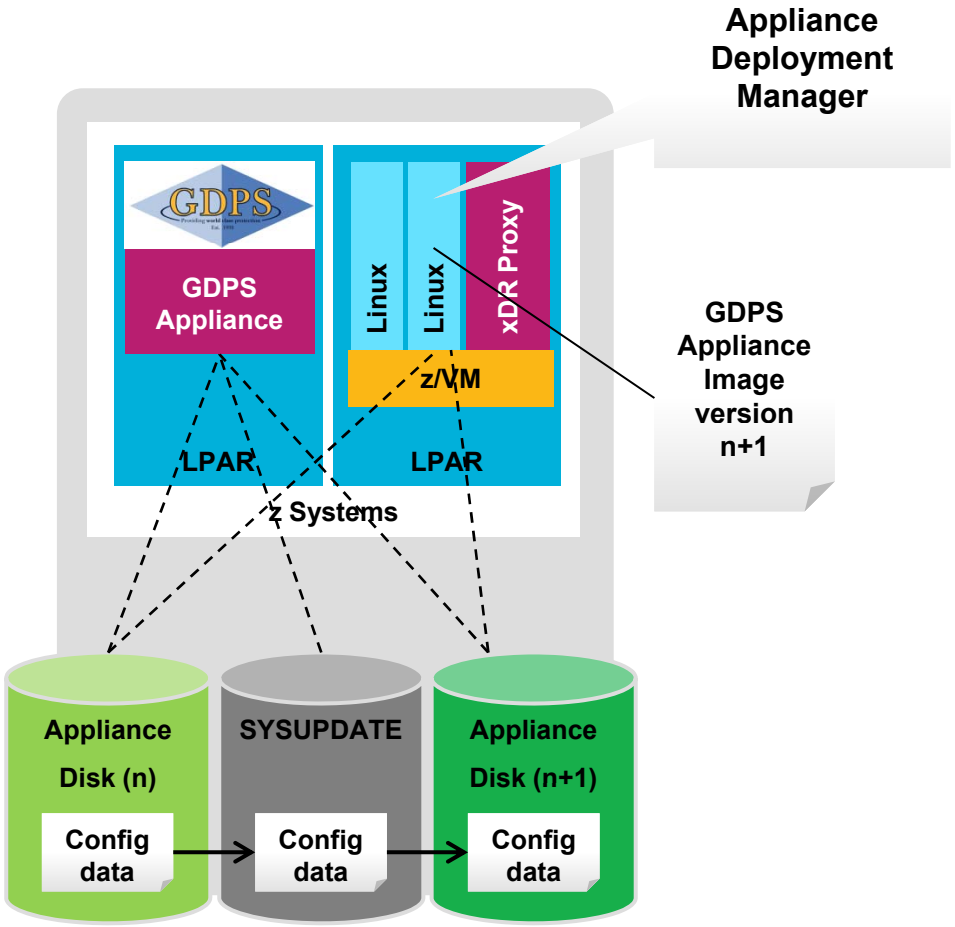

1. Download version n+1 from IBM
2. Use Deployment Manager to copy n+1 image to DASD device Appliance Disk (n+1)
3. Issue CLI command to save configuration and user data on volume SYSUPDATE
4. Shut Down GDPS Virtual Appliance
5. IPL Appliance using Disk (n+1). Migration procedure automatically copies, migrates and distributes saved configuration and user data from SYSUPDATE to the new Appliance Disk.

Done!

GDPS Disk Configuration

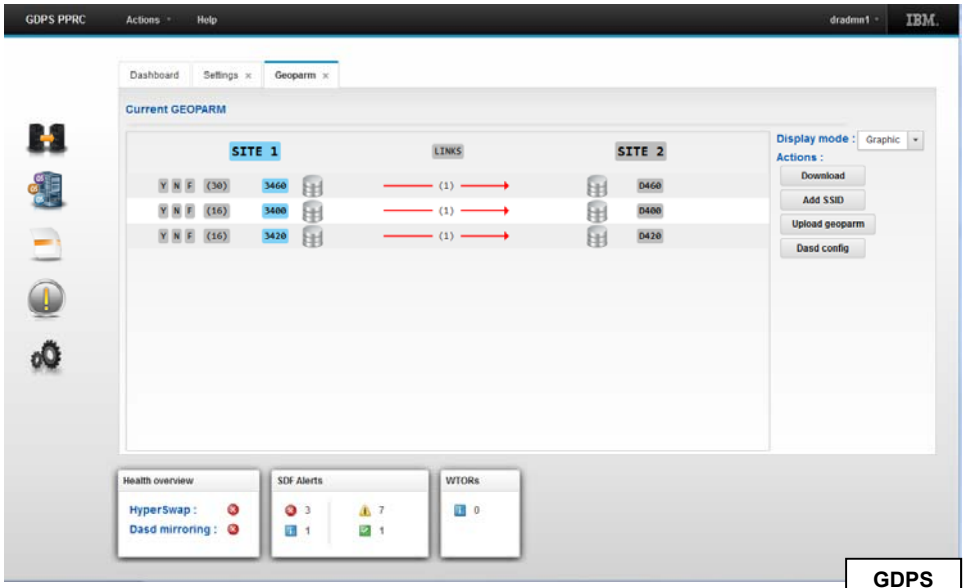

1. Edit GEOPARM file in WebUI editor

2. Upload and handover GEOPARM file to GDPS via WebUI

3. Activate GEOPARM file via GDPS Web UI

Done!

Common Operation

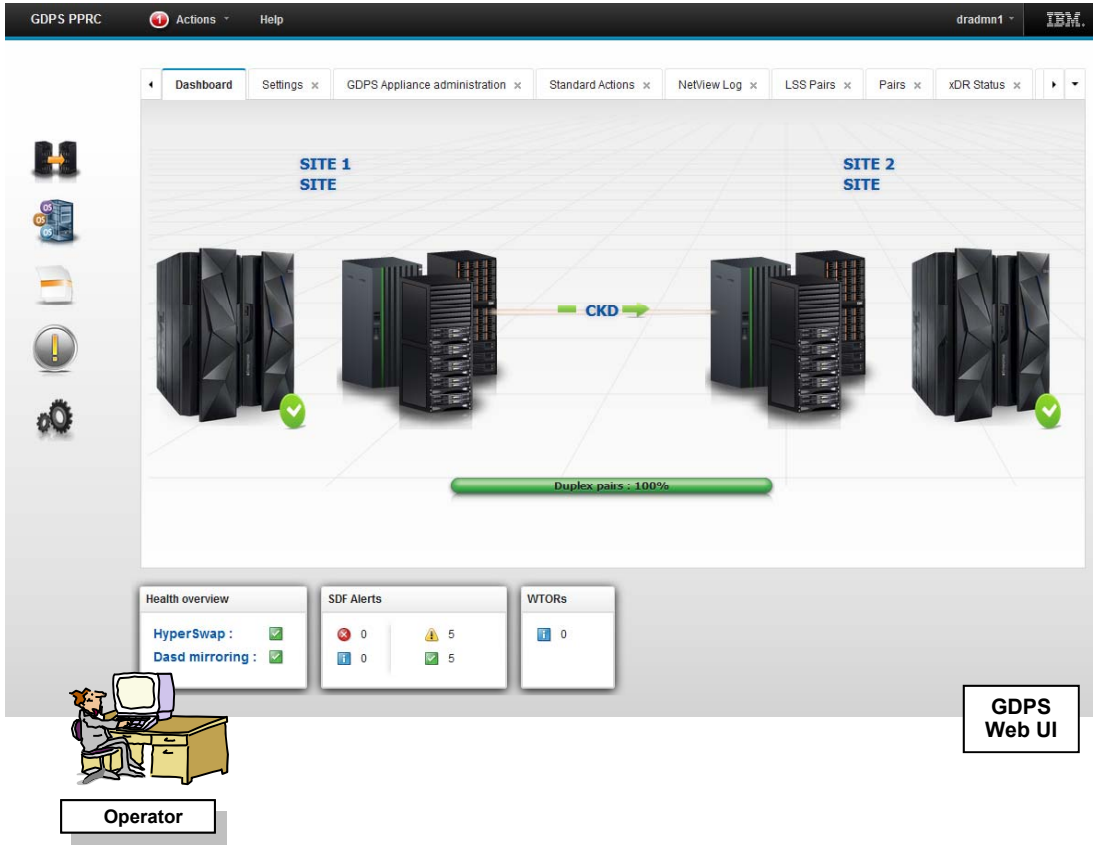

- Operate GDPS using the GDPS Web UI
  - View z/VM System and xDR Proxy
  - HyperSwap planned/unplanned
  - Site Switch planned/unplanned
  - Freeze planned/unplanned
  - Start/Stop z/VM image

Problem Determination

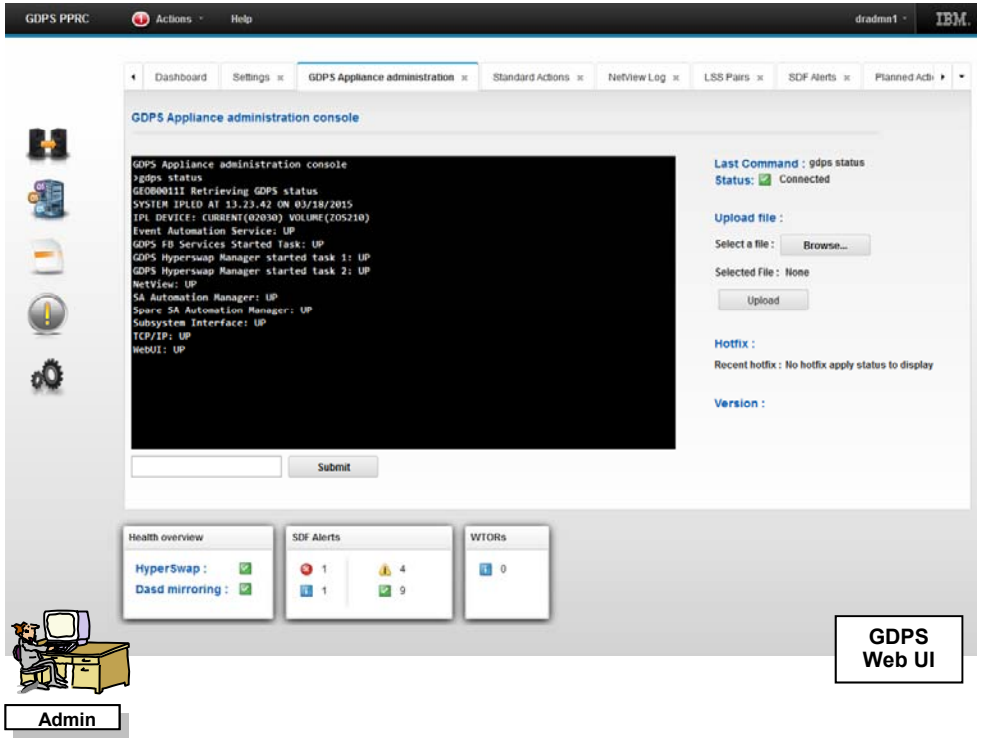

1. Issue CLI command to collect log and trace data via Web UI
2. Download log data via SFTP
3. Send log data to IBM for analysis