#SHAREorg

SHARE
Technology · Connections · Results

GDPS® End to End Support (xDR and DCM)

Steven Cook
IBM Corporation

August 7, 2012
Session Number 11662

SHARE in Anaheim 2012

Trademarks

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

IBM*	ESCON*	Redbooks*
IBM (logo)*	FlashCopy*	Sysplex Timer*
Ibm.com*	GDPS*	System p*
AIX*	HyperSwap	System z*
DB2*	IBM*	Tivoli*
DS6000	IBM logo*	z/OS*
DS8000	Parallel Sysplex*	z/VM*
Dynamic Infrastructure*	POWER5	

* 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.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.
InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
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.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
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.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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.

Disclaimer

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- 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.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here

Speaker Bio

Steven Cook
Level 2 Certified IT Specialist
Managing Consultant / GDPS Developer
Cooksd@us.ibm.com

From Kaypro to Microsoft Windows to UNIX to Linux to Mainframe

Clusters, clusters, and more clusters

Follow Me: www.twitter.com/cooksd

Session Contents

- Distributed Cluster Management
  - Concept / Requirements
  - DCM Agent Architecture
  - What's New
  - DCM Future Releases
- Distributed Systems Hardware Management Tool
  - Overview
- xDR
  - What's New
  - Future Enhancements

DCM CONCEPT & REQUIREMENTS

Suite of GDPS service products to meet various business requirements for availability and disaster recovery

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
Single Data Center Application 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
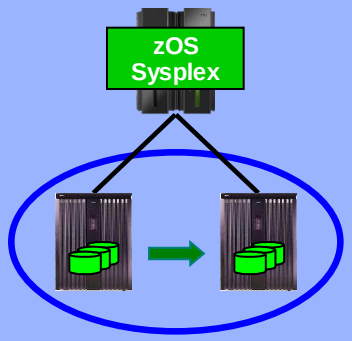	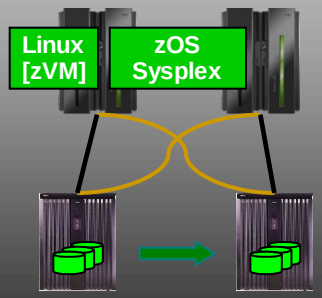	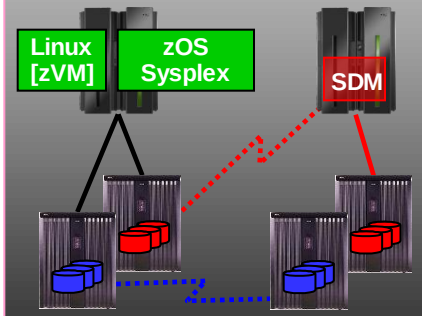	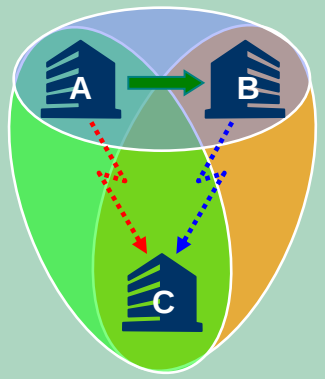
GDPS/PPRC HM	GDPS/PPRC active/active, active/standby configs	GDPS/GM & GDPS/XRC	GDPS/MGM & GDPS/MzGM
RPO 0 sec & RTO 0 sec	RPO 0 sec & RTO 1-2 min / <1 hr	RPO few sec & RTO 1hr	RPO 0 sec & RTO 1-2 min / <1 hr RPO few sec & RTO 1 hr

RPO – recovery point objective (data loss)
RTO – recovery time objective (downtime)

Synch replication
Asynch replication

DCM Concept

- Manage and coordinate availability and disaster recovery across distributed systems using high availability clustering solutions and the System z workload(s) that GDPS is responsible for
- Provide a single point of control for heterogenous site, cluster, and single-node failover
- Assist business process failover (BIA studys showed that apps spanned multiple server platforms and OSes)
- Preserve's data consistency across Mainframe and Open System data
- Monitors Tivoli SA AppMan clients or VCS health
- Simplify site to site DR tests End-to-End Enterprise(audit,monitor,control)
- One business process, but many islands (mainframe, blades, Vmware, System x, Power, apps, ...)

Not a replacement for Distributed Systems Admins

Integrated, Automated, Industry-unique

Requirements

- GDPS v3.8 or later
  - PPRC
  - GM
  - XRC
- DCM Agent
  - Tivoli System Automation Application Manager v3.1 or later
  - Veritas Cluster Server 5.0 MP3 or later
- TCP/IP Inter-site Connectivity
- Distributed Systems Disk Replication (HW or SW based replication)

GDPS DCM Tivoli SA Application Manager Supported Endpoint OS platforms (GDPS 3.9 & SA AppMan 3.2.2 running on zLinux)

- Windows Server
  - V2008 (32 bit) & (64 bit)
- AIX
  - V5.3 ML 4
  - V6.1
- SUSE SLES
  - V10 (32 bit) & (64 bit)
  - V11 (32 bit) & (64 bit)
- Red Hat RHEL
  - V5.0 (32 bit) & (64 bit)
- Windows & Linux for System X running under VMware ESX
- Clusters
  - MSCS (Windows)
  - PowerHa (AIX)
  - Veritas Cluster Server (VCS)
  - SA Multi-Platform (Linux, AIX, Windows, and Sun Solaris)

GDPS DCM VCS supported Server / OS platforms (GDPS 3.8+ & VCS 5.x MP3)[1]

Server / OS	AIX	HP-UX	Linux	Solaris	Windows
IBM System x (Intel / AMD x86_64)			▪Suse SLES 9/10/11sp2 ▪ RH 4/5/6	▪Solaris 9 ▪Solaris 10	▪Windows 2003 ▪Windows 2008
IBM System p (Power 5)	▪AIX 5.3+		▪Suse SLES 9/10/11sp2 ▪ RH 4/5/6		
HP (Itanium / PA RISC)		▪HP-UX 11.23+			
Sun (Sparc)				▪Solaris 9 ▪Solaris 10	
IBM System p / pHype (Power 5)	▪AIX 5.3+				

Notes: (1) VCS 5.0 supports Unix/Linux & VCS 5.1 supports Windows (32-bit and 64-bit) different Agents

DCM AGENT ARCHITECTURE AND INTERFACES

GDPS DCM Integration with Tivoli SA AppMan

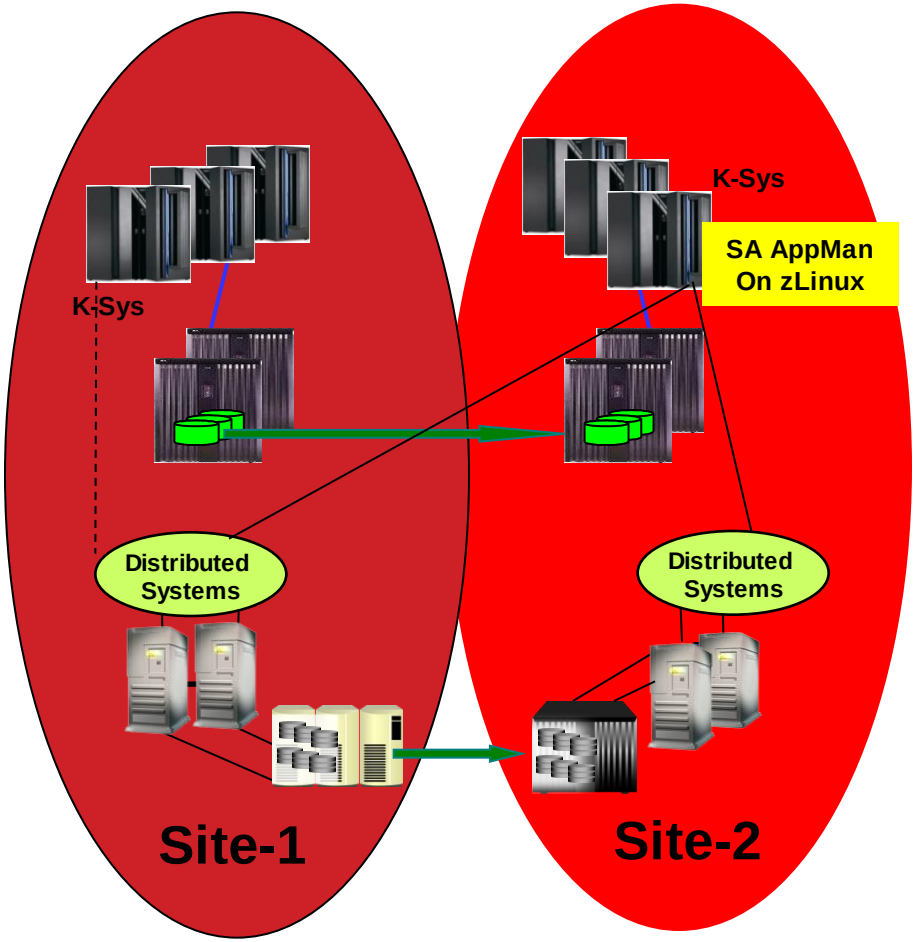

GDPS DCM Integration with Veritas Cluster Server

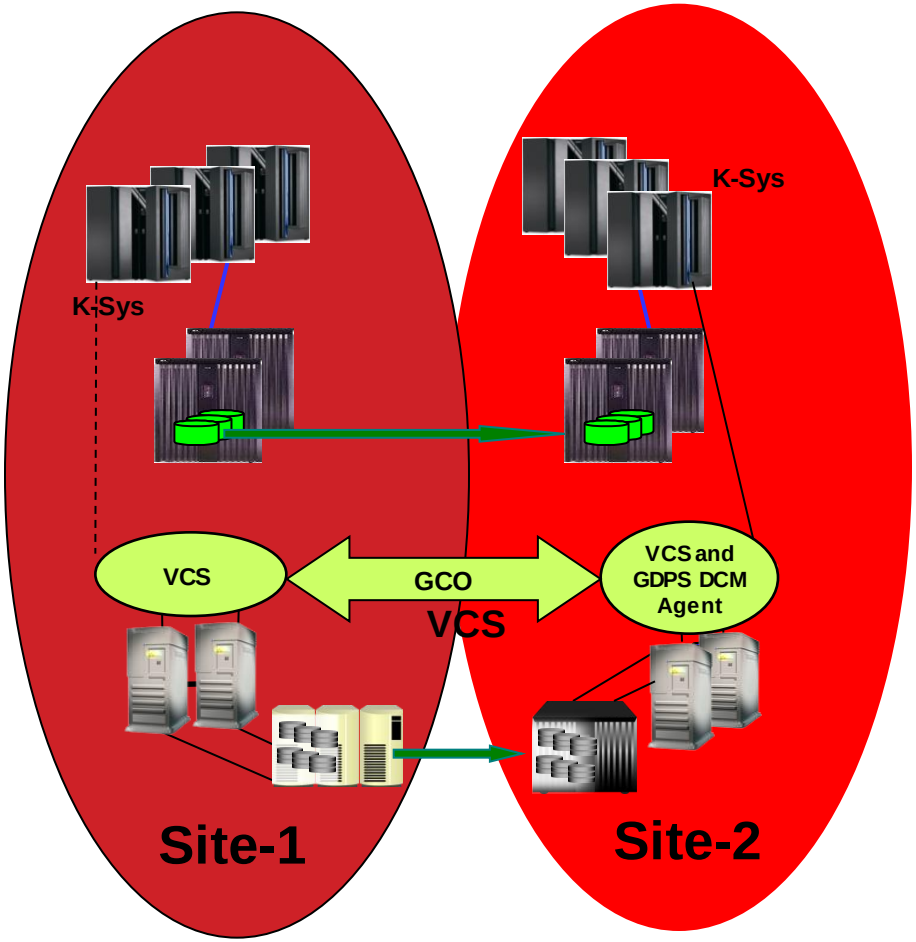

DCM functions

- Monitoring and alerting
  - Monitoring of DCM resources
  - Heart beating (DCM agent ←-→ K-sys)
  - GDPS generates alerts for DCM resources in abnormal state
- Manual operations through GDPS 3270 panels or Web GUI
  - Status of DCM resources can be queried
    - Single application “Service Group”, all applications on a single cluster, all clusters in a site
  - Perform planned actions on DCM resources
    - Such as START, STOP, SWITCH, Power On, Power Off
- Takeovers
  - Failures associated with DCM resources can be detected
  - Automation scripts can be executed
- Automation scripting for planned actions and takeovers
  - New CLUSTER script statements to manage VCS resources
    - Used to start, stop, switch, failover distributed applications
- Site-to-Site Failover Testing (VCS FireDrill)

GDPS DCM 3270 Panel

```
VPCPDCMP                    DCM - Distributed Cluster Management          SAM7
Actions:   V View   T Start   S Stop   W Switch Site
           PF Power OFF      PO Power On                  Sender : OK   Receiver: OK
Site/CL/AG/Group     Relation Info      State    HBI   HIM    Last Act. Time
_ Site1                                 OK
_ Site2                                 OK
_   FriendlyE2E        SA-AppMan        ACTIVE   30    0      HB         0121_15:04:42
 Agent View selected
Cmd ===>
  F1=Help   F3=Return    F4=Agent View F5=Refresh F6=Roll F7=Up F8=Down
  F9=Cluster View  F10=Agents with exceptions     F12=Group View
```

GDPS/PPRC ALLSITE1 CONTROL script executed on active master k-sys

- COMM='MOVE WORKLOAD FROM APPLICATION SITE TO RECOVERY SITE'
- SYSPLEX='STOP ALL' – stop all sys z systems except K1
- CLUSTER='STOP DISCRETIONARY CLUSTER(ALL) SITE(2)' - stop all discretionary workload
- CLUSTER='STOP CLUSTER(ALL) SITE(1)' - stop all applications in all clusters in Site1. Non-stretched clusters will be unavailable
- DASD='SWITCH DELPAIR' - switch PPRC mirror so that Site2 disks will be primary and Site1 will be secondary (applicable to sys z server CKD disk & distributed server FB disk)
- DASD='STOP SECONDARY' - suspend PPRC mirroring to ensure no I/O activity occurs to former PPRC primary volumes (applicable to sys z server CKD disk & distributed server FB disk)
- SYSPLEX='CDS SITE2' - use only the Site2 CDSs.
- CBU='ACTIVATE CPC=CPC2A' - activate CBU capacity on the CPC in Site2 in preparation for moving the Site1 systems to this CPC. The default is to wait until the CPC is activated.
- SYSPLEX='ACTIVATE CF2B LPAR' - activate the backup CF LPAR in Site2 (LP2A5). Activation of the LPAR will cause CFCC code to be loaded automatically.
- SYSPLEX='CFRECOVER UNCOND' - stop using Site1 CFs; clean up and use Site2 CFs only. This will force failed-persistent structures and switch to your Site2 CFRM policy.
- IPLTYPE='*ALLGDPS MODE=SITE2' - point all systems to load off site 2 disk (defined in GEOPLEX DOMAINS) for subsequent IPLs.
- IPLTYPE='*ALLGDPS ABNORMAL - point all systems to load in their alternate LPARs in site 2
- SYSPLEX='ACTIVATE P1 LPAR' - activate the backup LPAR for P1 in Site2.
- SYSPLEX='ACTIVATE P2 LPAR' - activate the backup LPAR for P2 in Site2.
- SYSPLEX='ACTIVATE P3 LPAR' - activate the backup LPAR for P3 in Site2.
- SYSPLEX='LOAD P1' - IPL P1 in its alternate location using Site2 DASD.
- SYSPLEX='LOAD P2' - IPL P2 in its alternate location using Site2 DASD.
- SYSPLEX='LOAD P3' – IPL P3 in its alternate location using Site2 DASD
- CLUSTER='START BUSINESSCRITICAL CLUSTER(ALL) SITE(2)' - start business-critical workload in all clusters in Site2. For the non-stretched cluster in Linux, workload is also restarted in Site2. The AIX active-active stretched cluster most likely has been running in Site2 all along since LVM facilitates data availability and the active-active concept. If the cluster was not already running in Site2 at the time of the failure, this statement will see to it that the workload is started..

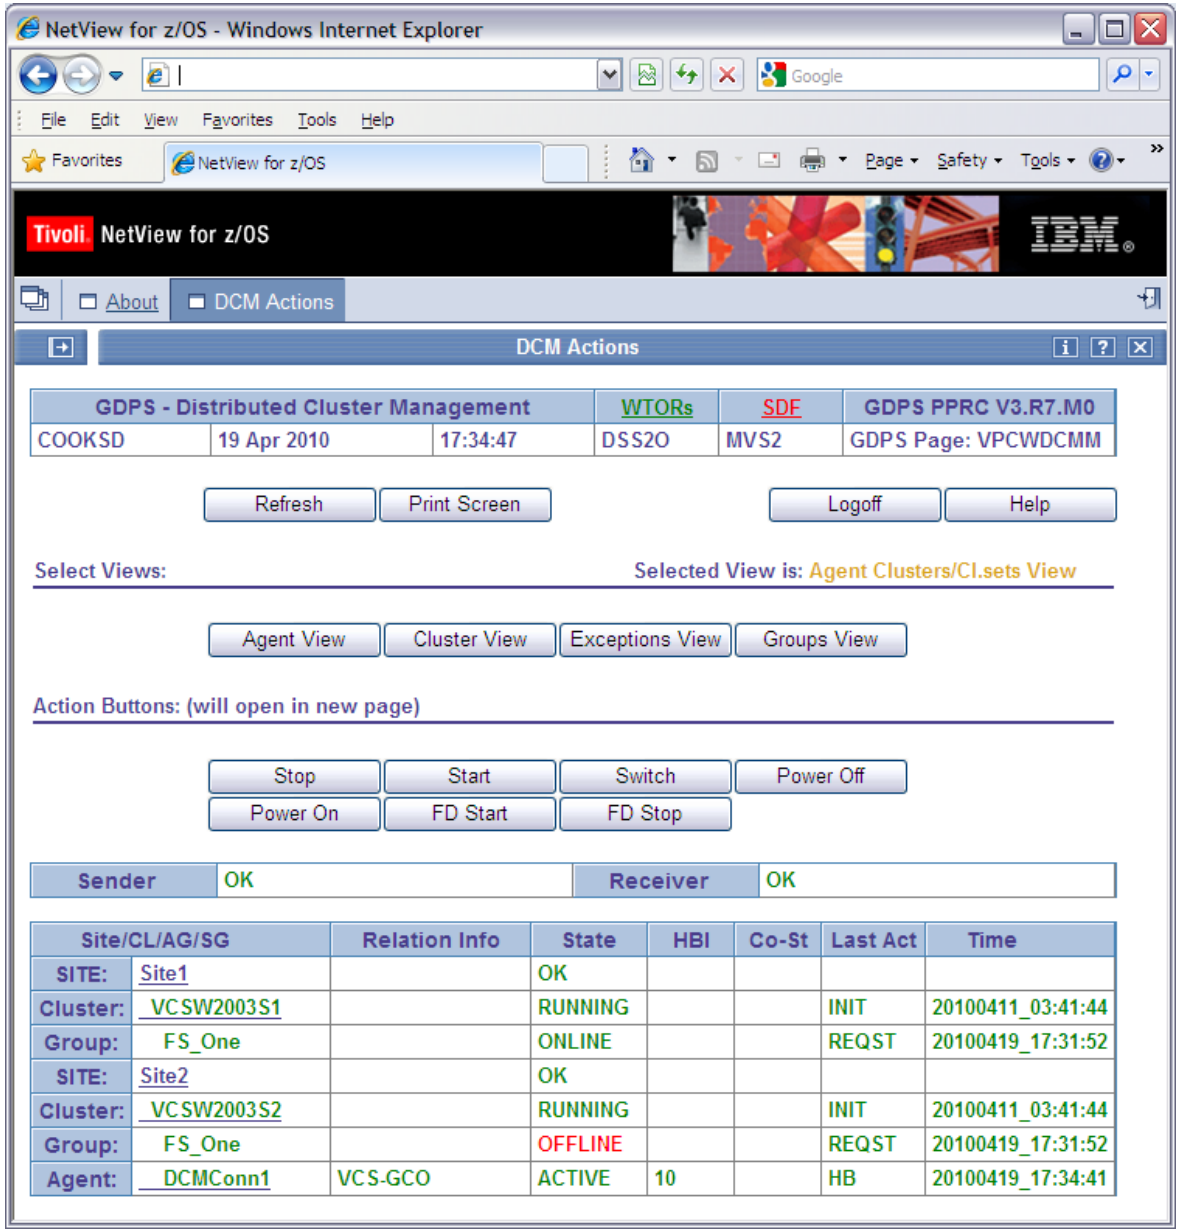

WHAT'S NEW

New Features in GDPS DCM

- IBM System z – zEnterprise BladeCenter Extension (zBX) support
  - IBM AIX on Power blades
  - Linux on System x blades
  - Microsoft Windows on System x blades

- GDPS/PPRC DCM supports distance up to 300 KM

- GDPS/GM support for Tivoli SA AppMan at unlimited distance
  - Tested at 25,000 KM between sites

- New GDPS DCM Tool

GDPS zEnterprise Support

Existing function extended to zEnterprise environments

- Application CA/DR at up to 300 KM using GDPS/PPRC DCM and AppMan
  - 200 KM max for non-Business Recovery Service (BRS) configurations
- Application CA/DR at up to 300 KM using GDPS/PPRC DCM and VCS
  - 200 KM max for non-Business Recovery Service (BRS) configurations
- Application CA/DR at unlimited distance using GDPS/GM DCM or GDPS/ XRC DCM and VCS
- GDPS V3R8 SPE GA November 18th, 2011

New function

- Application CA/DR at unlimited distance using GDPS/GM DCM and AppMan
- GA March 2012

GDPS/GM DCM Support of SA AppMan

GDPS DCM

SA AppMan DDR

Managed resources and data

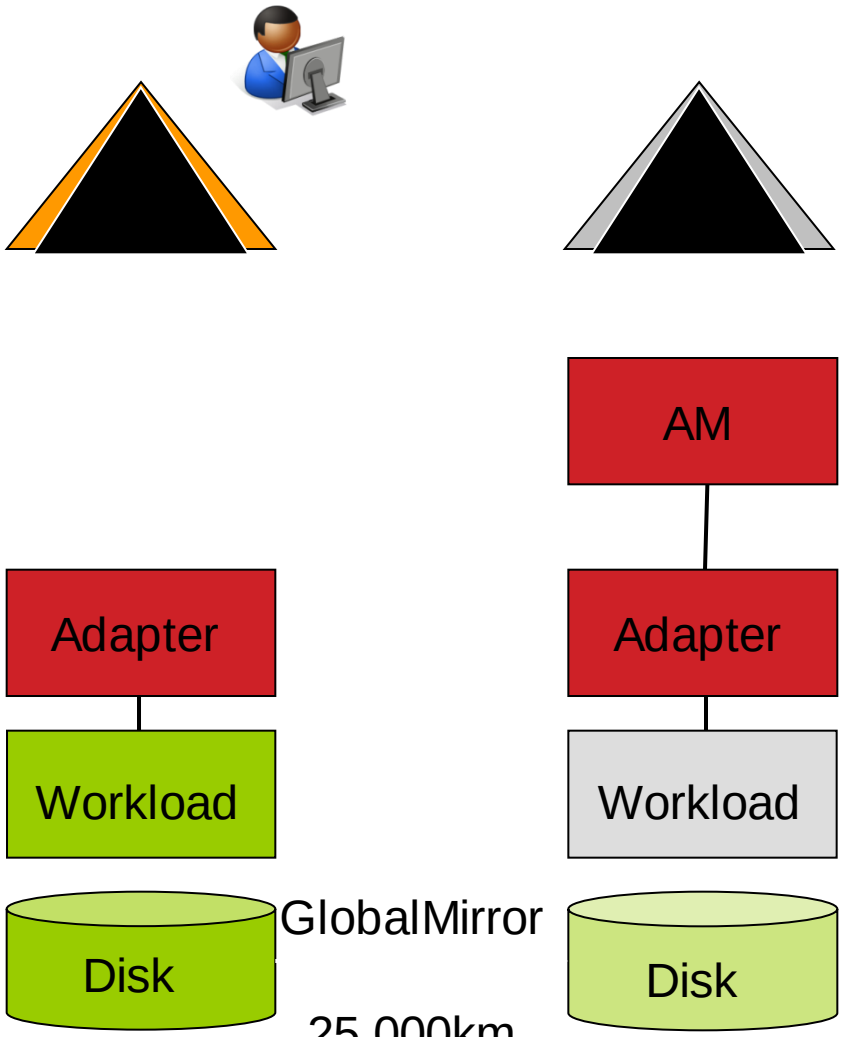

New GDPS Tool

Distributed Server Hardware Management Tool

- Allows a GDPS operator to control distributed systems hardware from a REXX script
- Hardware control actions such as boot, power down, start VM,…
- Works in concert with SA AppMan
  - AppMan controls the workload
- Supports many different vendor's server hardware
  - Supported for GDPS/PPRC, GM, and XRC environments
  - Support for zEnterprise zBX
- Shipped May 2012

FUTURE RELEASES

New Feature – AppMan Agentless Adapter

Tivoli SA AppMan Agentless Adapter Support

- Extend application CA/DR support to single-node (unclustered) servers integrated with AppMan's Agentless Adapter technology
  - Most distributed systems customers are not running clustered configurations
  - Removes requirement to install heavy-weight software prereqs for AppMan on single-node servers
  - Customer Benefit:
    - Simple integration of single node-servers into configurations managed by GDPS DCM and AppMan
    - Supported for Metro Mirror, Global Mirror, and XRC environments
  - Support for zEnterprise zBX
  - Support for wide array of distributed systems hardware
- **This is only a Statement of Direction**

New Feature – AppMan Toggle

Tivoli SA AppMan Toggle Support

- Support multiple AppMan instances
  - Will provide HA for AppMan
  - Automated site switching for AppMan
  - Supported for Metro Mirror, Global Mirror, and XRC environments
  - Support for zEnterprise zBX
  - AppMan must run on zLinux

- **This is only a Statement of Direction**

DISTRIBUTED SERVER HW MANAGEMENT TOOL

GDPS Distributed Cluster Management (DCM) Distributed Server Hardware Management Tool

- Monitor and Control Distributed Systems from GDPS
- End-2-End Enterprise wide synchronized failover
- A toolset of recovery automation templates
- GDPS REXX script driven from the mainframe
- Works with synchronous and asynchronous replication technologies
- Tivoli System Automation Application Manager provides the application workload management
- Veritas Cluster Server can provide application workload management
- Application agnostic

Distributed Server Hardware Management Tool Monitoring and Control Interfaces

- pSeries HMC[1]
- PHYP/PowerVM
- RSA[2] Cards
- BladeCenter

- VMware
- Native Windows / Linux / AIX
- VCS[3] / MSCS[4] / LinuxHA / PowerHA / Tivoli SA MP[5]

Extensible and Flexible

Server HW

Operating System

Cluster SW

Hypervisor

1Hardware Management Console
2Remote Supervisor Adapter
3Veritas Cluster Server
4Microsoft cluster Server
5System Automation Multi-Platform

How it works - GDPS/PPRC Unplanned Site Takeover

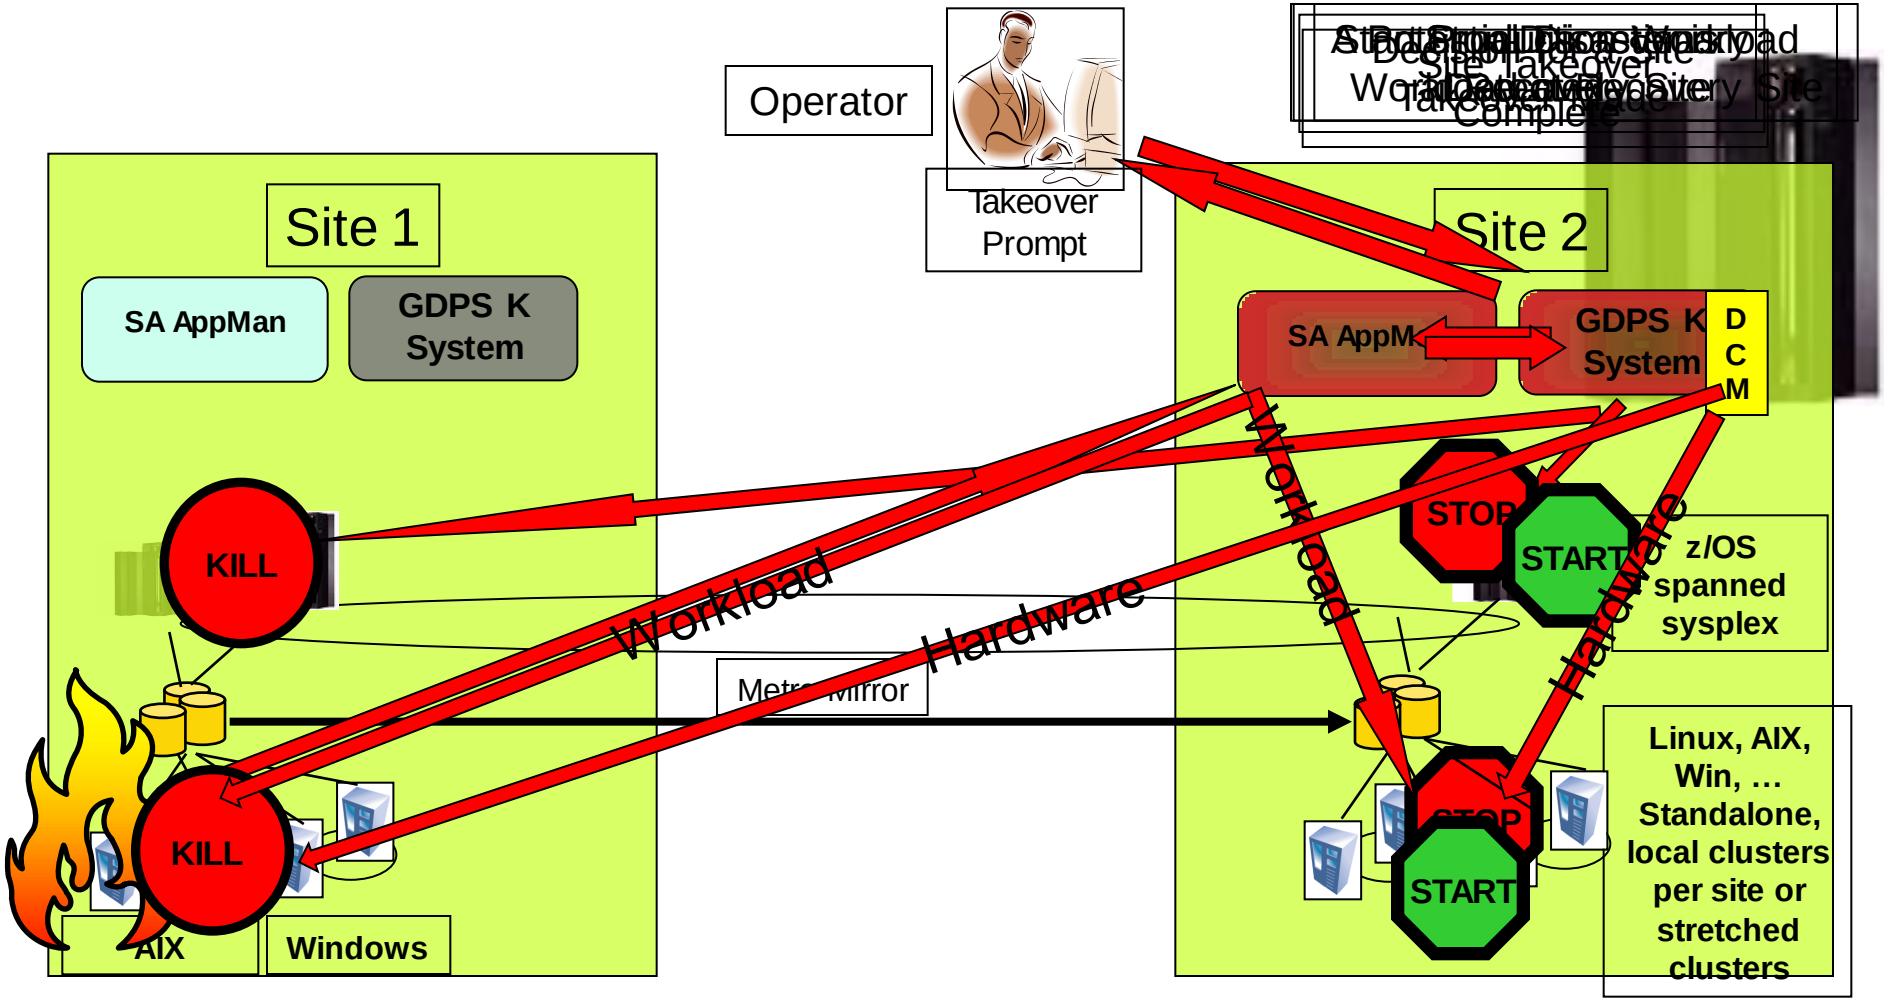

Prioritize this list of possible features

Monitors and Controls Distributed Systems in sync with MF

- Stop / Start VMs and LPARs
- Power Down/Up systems
- Boot / Reboot / IPL
- Select boot disks
- Monitor and Control Cluster Software
- Monitor System power states
- Monitor System boot state
- Stop / Start workloads on Distributed Systems w/Tivoli SA Application Manager or Veritas Cluster Server

System

IBM BladeCenter®

GDPS® XDR

AGENDA

- What is xDR?
- What's new or changed in GDPS 3.8 and GDPS 3.9
- New support planned for GDPS 3.9
  - xDR z/VM 6.2 SSI toleration
- xDR future enhancement of GDPS xDR
  - SSI Exploatition
  - User Requirement

GDPS/PPRC Multiplatform Resiliency for System z (xDR guest Linux on System z)

- GDPS/PPRC Multiplatform Resiliency for System z addresses high availability and disaster recovery for Linux for System z by using System Automation for Multiplatforms (Tivoli System Automation for Multiplatforms also referred to as SA MP) and by exploiting HyperSwap to expand GDPS/PPRC functionality previously only offered for the z/OS cluster to Linux for System z clusters.
- GDPS offers two "flavors" of the GDPS/PPRC Multiplatform Resiliency for System z:
  - GDPS/PPRC Multiplatform Resiliency for System z, when used for coordinated management of Linux on System z guests of z/VM, is referred to as *xDR guest Linux on z*.
  - GDPS/PPRC Multiplatform Resiliency for System z, when used for Linux on System z running native in a System z LPAR, is referred to as *xDR native Linux on z*.

What's new or changed in GDPS 3.8 and GDPS 3.9

- **SPE's in GDPS 3.8**

  GDPS PM38466 Multiple subchannel support.

  GDPS PM48808 z/VSE support running as guest

- **NEW in GDPS 3.9**

- **GEO–396 Shutdown of an xDR-managed z/VM system**

  - Timeout value for stopping of z/VM guests during stopvm are now based on historical statistics kept by GDPS rather than using an arbitrary, fixed timeout value.

  - Control the sequence in which you stop guest systems during a shutdown of z/VM.

What's new or changed in GDPS 3.8 and GDPS 3.9 CONT...

- **GEO-395 2 Proxy nodes for each z/VM host (Master/Backup)**
- **GEO-414 Re-IPL Action against xDR-managed z/VM or Native Linux systems.**

- **CHANGED**
  - SOCKET COMMUNICATION replace REXEC in GDPS 3.9 .

What's new or changed in GDPS 3.8 and GDPS 3.9 CONT...

- IPLing z/VM V6.2 systems

APAR VM65080 provide support for IPL z/VM 6.2 from GDPS using a IPLPARM selector.
FN:SITE1(PDVOL=1001) FN:SITE2(PDVOL=2001)

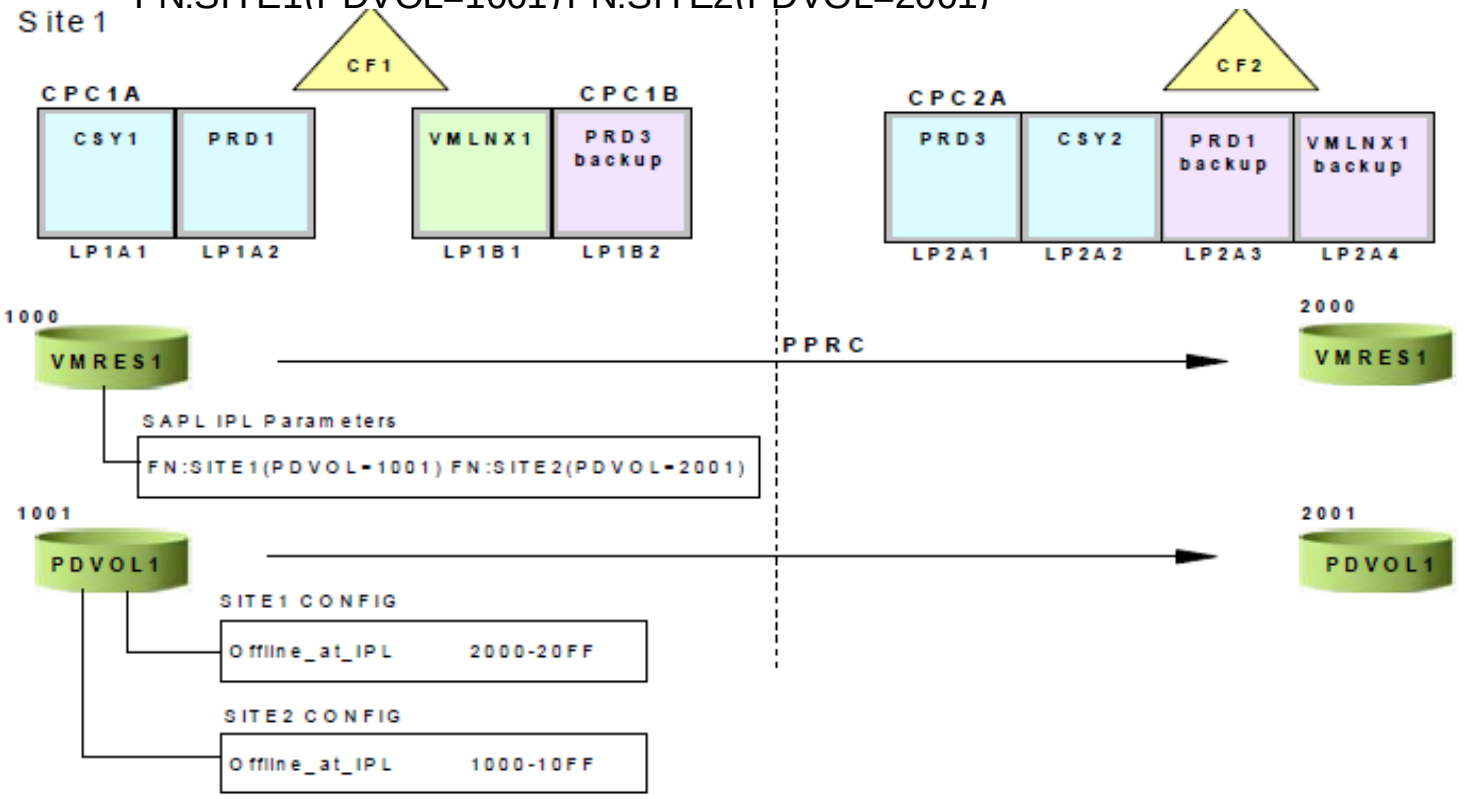

New support planned for GDPS 3.9

xDR SSI toleration

During a planned or unplanned Hyperswap , it may be necessary for the z/OS Ksys LPAR to reset an LPAR that is not responding to Hyperswap commands within the set time limits.

If one or more LPAR(s) is a joined member of a z/VM 6.2 Single System Image (SSI), then special action maybe required to reIPL a VM member LPAR and allow it to rejoin the SSI. Today, the action is a manual process to validate that a member LPAR is indeed down and not sharing resources shared with other member LPARs.

New support planned for GDPS 3.9 CONT...

- PLANNED SOLUTION

GDPS APAR PM64211 and z/VM VM65176 and SA MP 3.2.2 sp3

- Support is added to GDPS and xDR to automate the manual recovery of an LPAR(s) that has been reset by GDPS.
- New option to z/VM SET SSI "**FORCE DOWN sysname**"

GDPS will send
SET SSI FORCE DOWN sysname1-4 to each proxy in the SSI cluster following a GDPS RESET.

- New z/VM IPL Dialog prompt when joined members are not responsive.

GDPS will answer
GO if all members are found shutdown Else a WTOR will be raised for OPERATOR action.

New support planned for GDPS 3.9 CONT...

- Applicable to GDPS 3.9 only

- With PM64211 (ssi teleration) LOAD, RESET, ACTIVATE, DEACTIVATE, MODIFY, QryxDR, RE-IPL, DUMP.. will not be allowed from a NON-Controlling Master.

xDR future enhancement of GDPS xDR

- SSI Exploitation
  - Live Guest Relocation (LGR) of critical VM virtual servers from GDPS
  - GDPS management, including heart-beat monitoring, of critical VM virtual servers after a LGR executed from VM (Stretched Cluster support)

Dank u
Dutch

Merci
French

Спасибо
Russian

Gracias
Spanish

شكراً
Arabic

감사합니다
Korean

Tack så mycket
Swedish

धन्यवाद
Hindi

תודה רבה
Hebrew

Obrigado
Brazilian Portuguese

谢谢
Chinese

Dankon
Esperanto

Thank You

ありがとうございます
Japanese

Trugarez
Breton

Tak
Danish

Danke
German

Grazie
Italian

நன்றி
Tamil

děkuji
Czech

ขอบคุณ
Thai

go raibh maith agat
Gaelic

Related Sessions

Tuesday, 3pm in Salon D – Session 11662

- **GDPS End to End Support (xDR and DCM)**

Wednesday, 8am in Salon H – Session 11663

- **GDPS Active/Active Sites Update**

Friday, 11am in Salon H – Session 11661

- **GDPS 3.9 Update**

Session Survey

- Session 11662
- QR code or Paper

Related Sessions:

Tuesday, 3pm in Salon D – Session 11662

- **GDPS End to End Support (xDR and DCM)**

Wednesday, 8am in Salon H – Session 11663

- **GDPS Active/Active Sites Update**

Friday, 11am in Salon H – Session 11661

- **GDPS 3.9 Update**