GDPS and GDPS Virtual Appliance

Steven Cook
GDPS Software Designer
IBM
cooksd@us.ibm.com

VM Workshop
June 28, 2019

IBM

Agenda

- GDPS Concepts and Product Family Overview
- GDPS xDR Overview
- Synergy with z/VM (xDR GDPS Metro 4.2)
- GDPS Virtual Appliance
- GDPS xDR and Virtual Appliance Futures

When do we need Continuous Availability/Disaster Recovery solution?

Global disaster	Local disaster	Maintenance
- Hurricane - Earthquake - Power plants failure...	- Fire - Power supply problem - Unplanned IT Failure	- Hardware & software update. - Switch to a new datacenter - Test
Do we have safe backup? Do we have system ready to start outside of the region?	Could we avoid downtime and data loss? Is there a procedure to restart systems?	Can we do that transparently? How to reduce the risk of a rolling problem during a maintenance scenario?

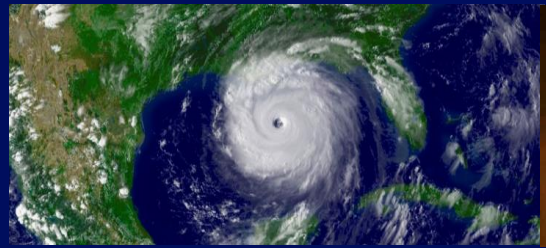

Resiliency is a growing concern

- Outages can cost a lot of money and reputation
- Public is constantly asking for more availability thru mobile access
- Government increase pressure to provide Disaster Recovery capabilities

"When a situation is within your control, take action.
When a situation is outside your control, make preparations."
-@jamesclear -Applied to CA/DR by Steven Cook

GDPS Concepts and Product Family Overview

GDPS™
Providing world class protection

IBM

Peer to Peer Remote Copy = PPRC = Metro Mirror

- **Synchronous copy implications :**
- Metropolitan distance
- Some impact on I/O (Write only) performance
- For CKD and FB disks
- At a given time, PRI volume and SEC volumes are exactly the same

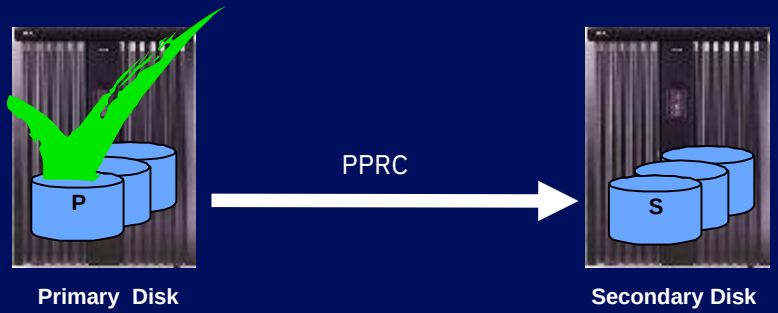

HyperSwap – the Technology

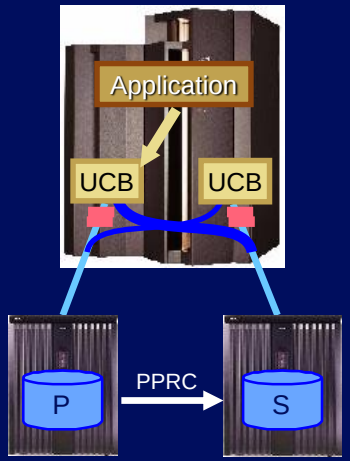

UCB = Unit Control Block

- Extends Availability to disk subsystems
- Substitutes Metro Mirror (PPRC) secondary for primary device
  - ✓ Automatic – No operator interaction
  - ✓ Fast – Can swap large number of devices
  - ✓ Non-disruptive – applications keep running
  - ✓ Includes volumes with Sysres, page DS, catalogs

- **Disk no longer a single point of failure**

Comprehensive application and data availability solution

Hyperswap Performance

Disk config
201 LSSs
20658Pairs

System config
2 Kp systems
22 zOS prod

HyperSwap UIT time?

UIT = 7,76sec

Disk config
251 LSSs
22434 Pairs

System config
2 Kp systems
2 zOS prod
14 zOS Proxy
46 Linux under VM

HyperSwap UIT time?

UIT = 8,24sec

UIT is the worst case I/O pause duration for any given device on any given system. Most devices on most systems will see a shorter pause.

Mirroring Failure vs. Primary I/O Failure

PPRC Mirroring Failure:

- Can not mirror to secondary
- GDPS will perform Freeze and follow on action based on PPRCFAILURE policy option
- **The GDPS solution for mirroring problems is Freeze automation**
- **Sec volumes are consistent after the Freeze**

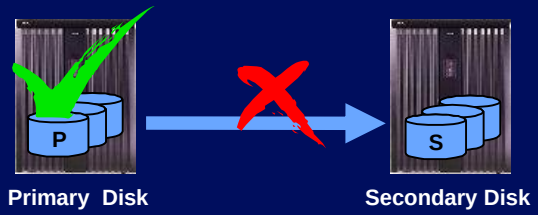

• Primary Failure (HW):

- Can not write to primary disk
- Captured by system doing write on 'HyperSwap-ready' systems
- GDPS notified via ENF
- If PRIMARYFAILURE=SWAP,xxx GDPS performs Swap else GDPS performs Freeze
- **The solution for primary failures is HyperSwap**

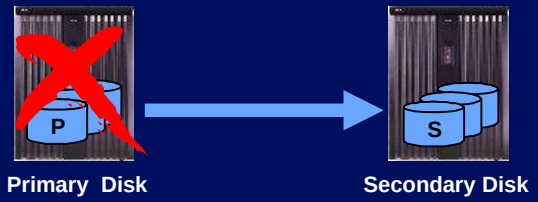

Standard CA/DR architecture : 2 sites

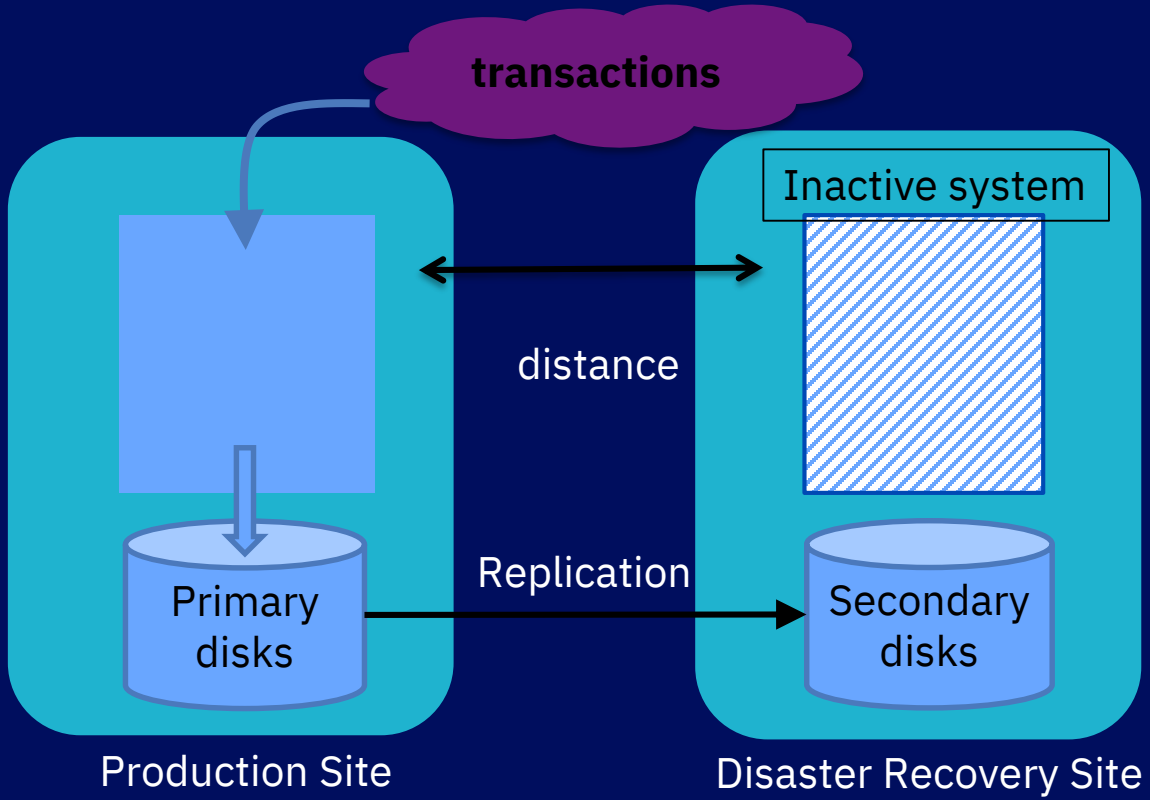

- System Clustering usually limit distance between sites
- How to use secondary data ?
- Complex commands when under stress

- Improved performances mandates Automation

- Is it possible to avoid impact on users ?

Business Continuity Metrics

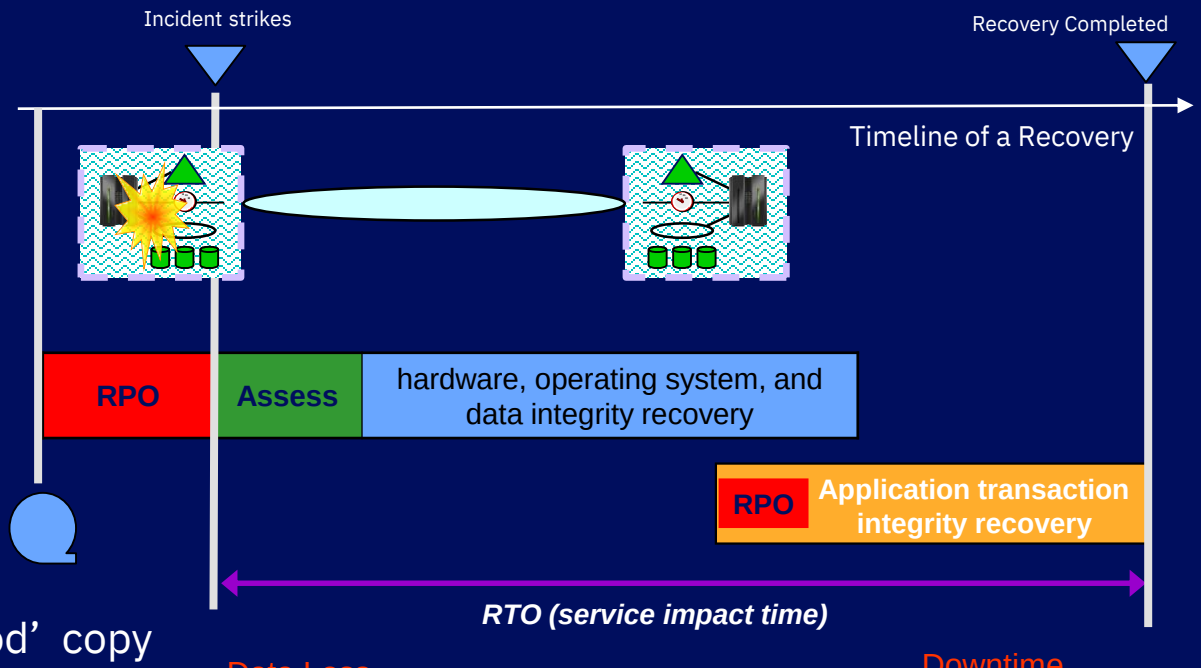

Data Loss

RPO = Recovery Point Objective

The time window of the amount of data loss in the event of a disaster.

Downtime

RTO = Recovery Time Objective

The amount of time required to resume operations in the event of a declared disaster.

Continuous Availability & Disaster Recovery Offering for IBM Z

First GDPS installation 1998, now more than 950 licenses in 51 countries

- **IBM has developed a solution since 1998**
- **Rely on Automation, Clustering and Data Replication protocols**
  - not limited to IBM Storage, but has differentiator when combined to IBM Storage
  - service offering : skill transfer to customer
- **Constantly Improved with customer**
  - Design Council : customers vote for new functionalities/priorities
  - RFE Process
  - IBM Montpellier GDPS Solution Test team

GDPS Solution Names and Configurations

GDPS Metro HM	HyperSwap Manager
GDPS Metro	Metro Mirror Single Leg
GDPS Metro	Metro Mirrror Dual Leg
GDPS Global GM	Global Mirror
GDPS Global XRC	XRC (z/OS Global Mirror)
GDPS Metro Global	Metro z Global Mirror 3-site and 4-site
GDPS Metro Global	Metro Global Mirror 3-site and 4-site
GDPS CA	Active-Active
GDPS VA	Metro Mirror

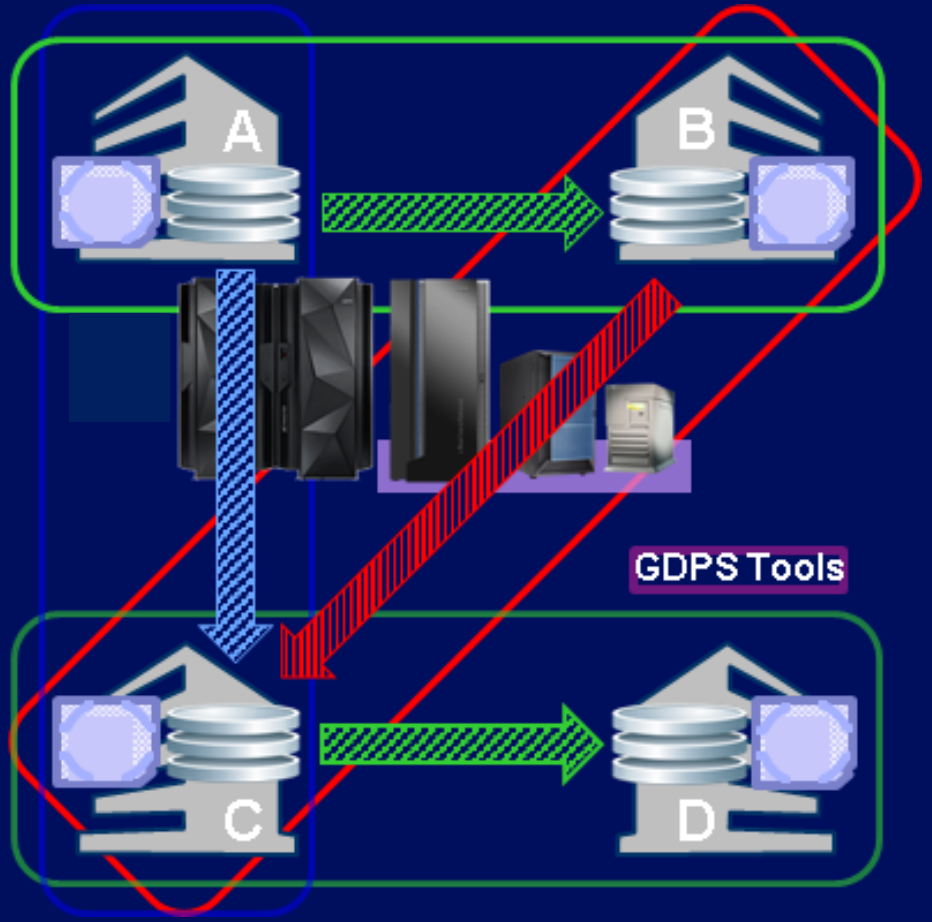

Near-continuous availability of data within a data center

Single data center

Applications can remain active
Near-CA in the event of a storage subsystem outage

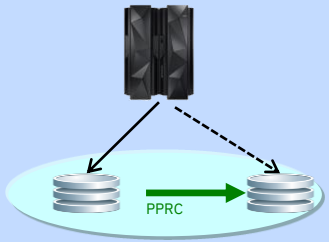

Near-continuous availability (CA) and disaster recovery (DR) within a metro. region

Two data centers

Systems can remain active
Multisite workloads can withstand site and storage failures

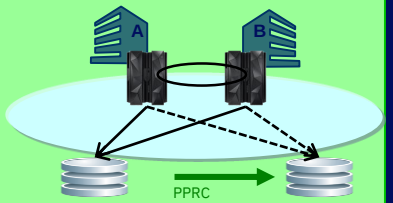

DR RPO = 0 and RTO<1h or
CA RPO = 0 and RTO minutes

Near-CA, DR and cross-site workload balancing at extended distance

Two or more data centers

Disaster recovery for out-of -region interruptions
All sites acti

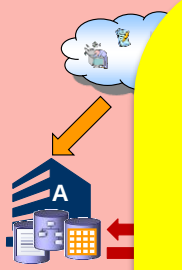

SW Re
RPO second

Disaster recovery at extended distance

Two data centers

More rapid systems disaster recovery with "seconds" of data loss
Disaster recovery for out-of-region interruptions

Near-CA and with metro DR for zVM – Linux on z only

Two data centers

Systems can remain active
Multisite workloads can withstand site and storage failures

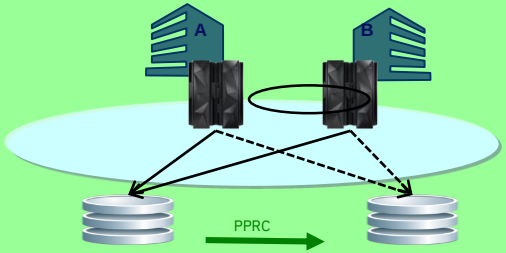

DR RPO = 0 and RTO<1h or
CA RPO = 0 and RTO minutes

Near-continuous availability (CA) and disaster recovery (DR) within a metro. region

Two data centers

Systems can remain active
Multisite workloads can withstand site and storage failures

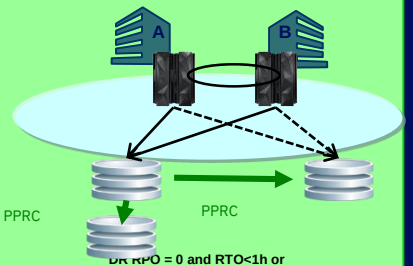

DR RPO = 0 and RTO<1h or
CA RPO = 0 and RTO minutes

Near- CA regionally and DR at extended distance

Three data centers

CA for site disasters and DR for regional disasters

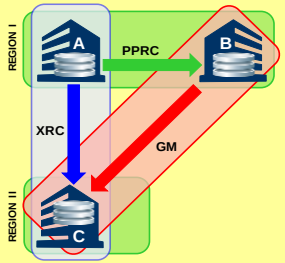

Within region = same as GDPS/PPRC
DR in other region = same as GDPS/GM - XRC

Near- CA regiona distance

Fo

CA for site dis
disas

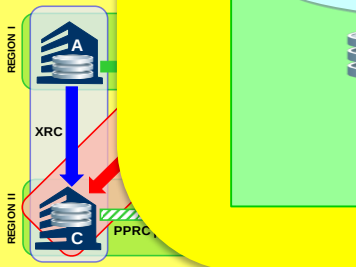

Within region = same as GDPS/PPRC
DR in other region = same as GDPS/GM - XRC

GDPS xDR Overview

xDR Platforms and Site configurations

GDPS™
Providing world class protection

IBM

GDPS xDR: Linux Guests and Native Linux on IBM Z

– Continuous Data Availability if you have z/OS

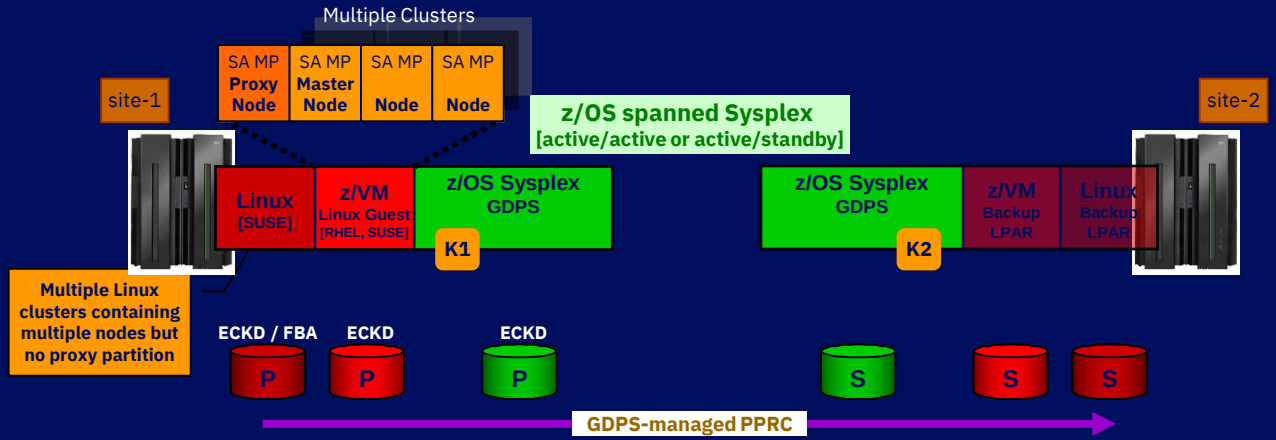

- Coordinated HyperSwap – z/OS, z/VM with its guests, and native Linux
- Graceful shutdown and startup (re-IPL in place) of Linux clusters or nodes
- Coordinated takeover – recovery from a Linux node or cluster failure
- Multiple Linux clusters are supported, as are multiple z/VM systems & Linux LPARs
- All members of a cluster must run under same z/VM system or SSI or in same Linux LPAR

Coordinated recovery for planned and unplanned events

xDR GDPS METRO v4.2 (unless noted)

Environments supported by GDPS and xDR

- z/VM
- SSC / IDAA
- Linux on Z (Native Linux), 3.14 only
- z/KVM
- z/OS with z/Proxy

What GDPS 2/3/4 site configs are supported by xDR

- z/VM are supported in GDPS Metro (single, dual(ss0 only) leg), Metro Global – GM 3 & 4 Site, Metro Global - XRC 3 & 4 site Site.
- z/OS Proxy are supported in GDPS Metro (single, dual leg), Metro Global - GM 3 & 4 site, Metro Global - XRC 3 & 4 site
- Linux on Z (Native Linux) GDPS PPRC, GDPS XRC, 3.14 only,
- z/KVM GDPS Metro (single leg) Note1
- SSC / IDAA GDPS Metro (single leg) Note1

Note1 : No HyperSwap support

Synergy with z/VM (xDR GDPS Metro 4.2)

xDR New Functions

GDPS™
Providing world class protection

IBM

- z/VM

Live guest relocation Keyword MAXTotal

GDPS now supports the MAXT keyword on the SSIRELOC script statement and panel operations, allowing more time to perform Live Guest Relocations (LGR) of z/VM guests.

Script verb
SSIRELOC '{MOVE|TEST} [DOMAIN(domain_name|ALL)]
SOURCE(sysname)
[TARGET(sysname|ANY)] [USER(userid|ALL)] [MAXQ(nn)] [MAXT(nn)]'

For GDPS 3.14 and 4.1 see New Function apar PH00538 and PH06828

xDR GDPS METRO 4.2

- SA Multi-platform (aka z/VM proxy) fixes and enhancements

IJ06476 + IJ10501 z/VM Live Guest Relocation fixes and enhancements

IJ08486 z/VM shutdown processing + xdrwatchdog

IJ15710 long running monitor interferes with e.g. planned HyperSwap
IJ15783 z/VM shutdown+reipl processing

All these are available on Fix Central.

Please always refer to the PSP bucket to determine the latest recommended level.

xDR GDPS METRO 4.2

- z/OS Proxy Performance Enhancement

- GDPS will now use up to 4K data buffers to send data over the CTC Access method to the target system.

Environments with the z/OS Proxy and > 10000 PPRC pairs in the entire consistency group gain from the enhancement during xDR INIT.

- GDPS now also handles systems in parallel during xDR INIT and HyperSwap PREPARE.

- Older z/OS proxy agent from GDPS 3.12, 3.13 , 3.14 and 4.1.0 is supported but will not gain from the enhancement.

- GDPSIVHP is updated to show the level of the z/OS proxy and post a warning when a older version of the z/OS Proxy agent is found

- New command MODIFY GEOPFSRV,QUERY LEVEL

- For GDPS 3.14 and 4.1 see New Function apar PH09213

GDPS Virtual Appliance (VA)

Overview and New Functions

IBM

GDPS Virtual Appliance (VA)

GDPS Virtual Appliance v.next

General availability soon*

GDPS Virtual Appliance v3.14

Includes z/OS 2.2 and GDPS 3.14

Support for z/VM only

Currently available

*IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.

GDPS Virtual Appliance brings High Availability capabilities of IBM System z to Linux-only enterprises

- No z/OS skills required
- Years of Intellectual Capital invested in GDPS logic available providing
  - Ability to avoid impact cause by primary disk outage (planned or not)
  - Scripted automation
- Monitoring and central point of control provided thru the GDPS interface

GDPS Virtual Appliance – What do we have in the box?

Step1. Download the package and install it from a partition on disks that will be dedicated to the appliance.

Step2. Load the new partition via the HMC

Step3. Access to the appliance via a dedicated command line interface (for updates, etc...)

Step4. Admin can now manage his GDPS control system via the new dedicated interface.

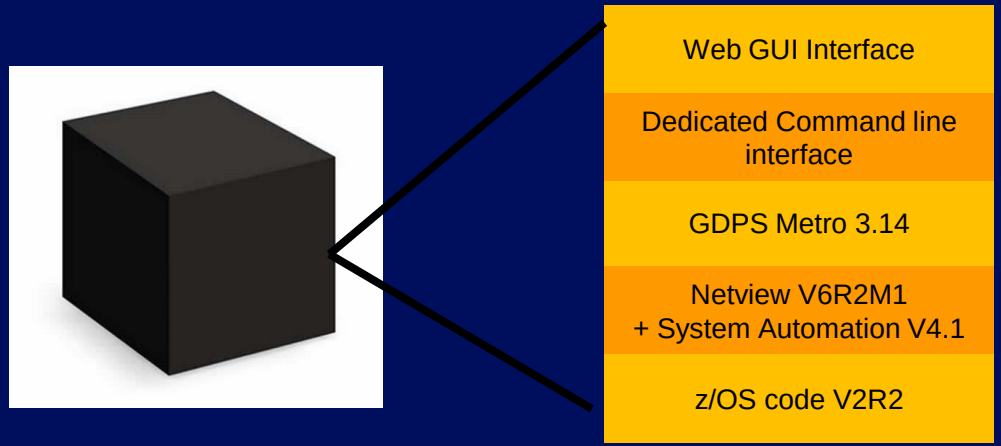

GDPS Virtual Appliance vs GDPS xDR feature

GDPS	xDR	Appliance
Enterprise Type	z/OS	No z/OS
Linux guest on z/VM	✓	✓
Native Linux on z	✓	⊘
Pre-packaged (No z/OS skill required)	⊘	✓
CKD disks	✓	✓
SCSI disks (native Linux on z)	✓	⊘
Multivendor (disks)	✓	✓
SSI cluster (zVM clustering)	✓	✓
MultiTarget Metro mirror, 3 sites solutions (MzGM)	✓	⊘
Multiple SubChannel set	✓	⊘

GDPS VA: Supported Linux guests

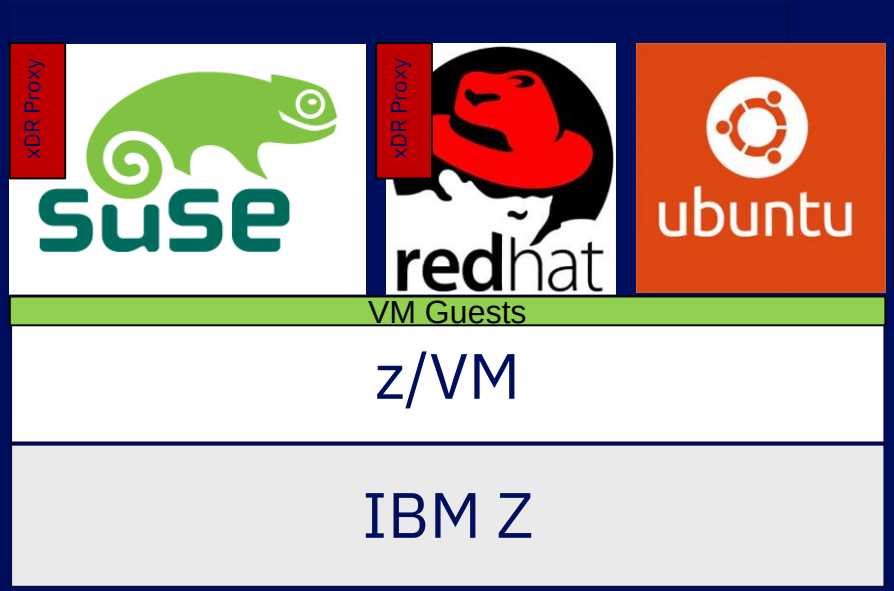

GDPS Virtual Appliance

Operate GDPS using the new GDPS web GUI
- View z/VM System and xDR Proxy
- HyperSwap planned/unplanned
- Site Switch planned/unplanned
- Freeze planned/unplanned
- Start/Stop z/VM image
- Start/Stop z/VM guest(s)
- Manage Linux clusters
- Live Guest Relocation

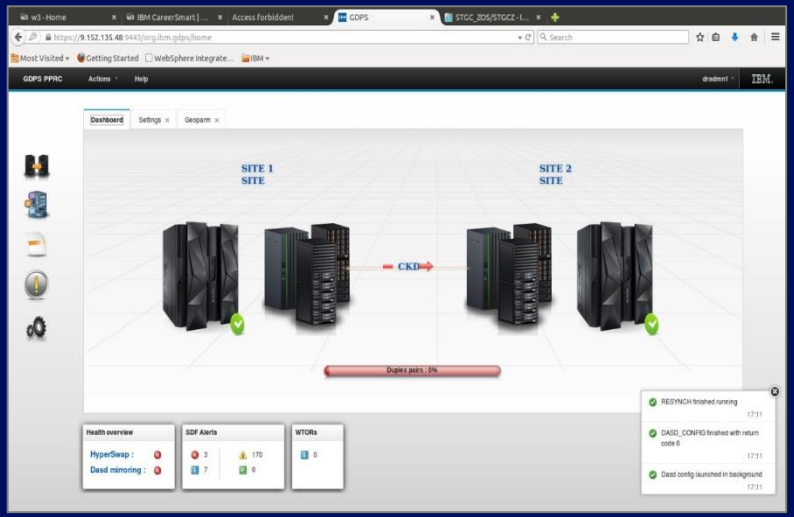

Provides non-z/OS customers the same benefits of high availability and D/R which were only available to z/OS customers

Provides additional benefit for moving workloads to Linux on IBM Z

GDPS Virtual Appliance: GDPS/PPRC xDR capabilities for z/VM & Linux on z Systems clients who do not use z/OS

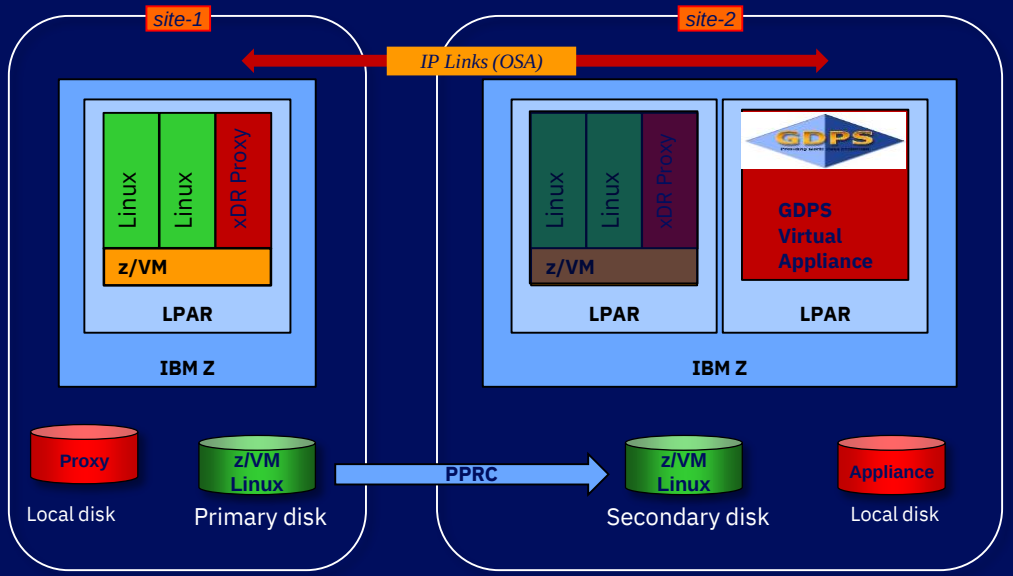

A comprehensive continuous availability and disaster recovery solution for z/VM clients

- **Monitors and manages**
  - PPRC replication management
  - Secondary data consistency
  - Planned and unplanned HyperSwap
- **System and Hardware management capabilities**
- Intuitive graphical user interface
- Simple scripting capability

z/VM High Availability Proxy configuration

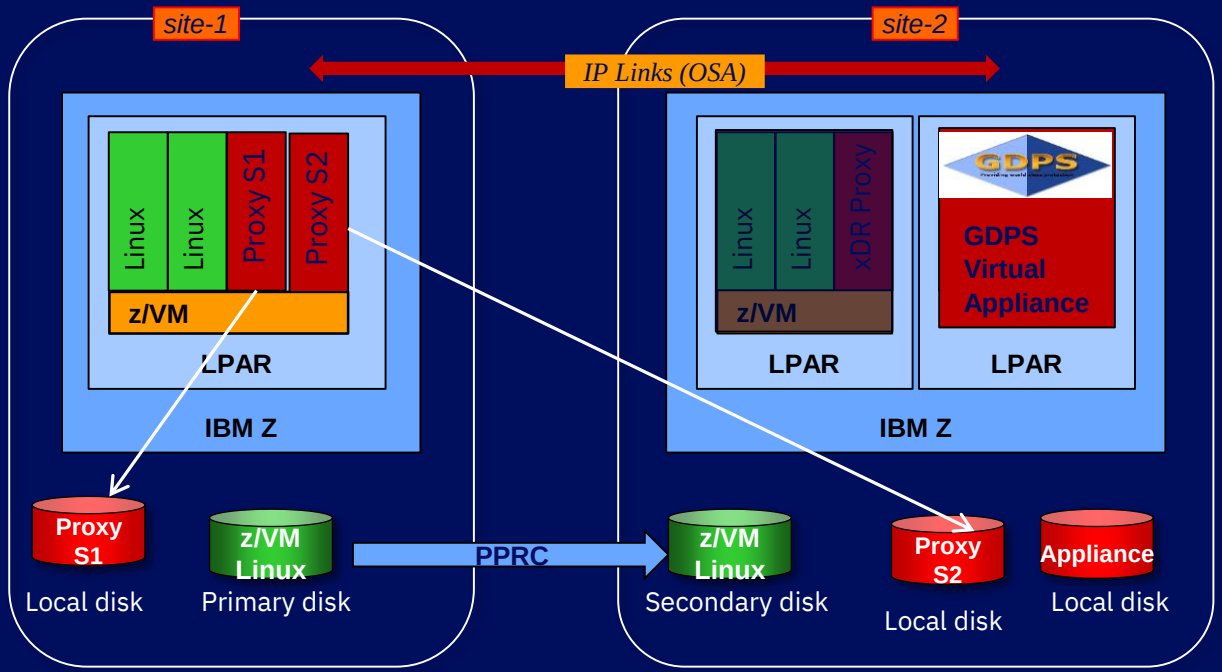

- Site2 proxy runs on local, non-mirrored disk in Site2
- Site2 proxy is “active” when PPRC primary disk is in Site1
- Site1 proxy runs on local, non-mirrored disk in Site1
- Site1 proxy is “active” when PPRC primary disk is in Site2
- Proxy switch from inactive to active automatic based on PPRC primary disk location
- Provides backup proxy protection in the event of planned or unplanned disk outage

GDPS xDR and Virtual Appliance Futures

You can influence our design

GDPS™
Providing world class protection

IBM

What is the IBM RFE Community?

A partnership to provide transparency in product support

- A community where IBM and IBM Customers interact and collaborate on Requests for Enhancements (RFEs) for IBM products

- Hosted on the IBM developerWorks website

Tivoli
PureSystems
Rational
Security Systems
Websphere
Information management
B2B & Commerce

NEW

Servers and Systems Software

GDPS xDR and Virtual Appliance Futures

Question 1

Do you use EDEVs or plan to use?

1 – Yes

2 – No

Question 2

Which storage do you use with z/VM?

1 – IBM DS8K

2 – Other IBM

3 – Other Brand

Question 3

Which storage do you use with z/VM?

1 – CKD

2 – FBA

3 – CKD & FBA

Question 4

How many pairs of replicated storage?

1 – <10,000 pairs

2 – >10,000 pairs

3 – >20,000 pairs

Question 5

How many z/VM instances do you have?

1->10 11->20 21->50 >50

Question 6

How many guests per z/VM hypervisor?

1->10 11->30 31->50 >100

Question 7

Which Linux distro do you run on z/VM?

1 – SUSE

2 – RHEL

3 – Ubuntu

GDPS xDR and Virtual Appliance Futures

Question 8

Do you mix lines of business per z/VM instance?

1 – Yes

2 – No

Question 9

Are you interested in joining the xDR Working Group (WG)? (NDA required)

1 – Yes

2 – No

Question 10

Are you interested in joining the GDPS Virtual Appliance Working Group (WG)? (NDA required)

1 – Yes

2 – No

Questions?


Grazie
Obrigado
Koszonom
Mahalo
arigato
THANK YOU
Спасибо
Takk
Рахмет
Prasom
Gracias
Danke
Dziekuje
Merci
Multumesc

Additional Information

Web sites:

- GDPS https://www.ibm.com/it-infrastructure/z/technologies/gdps
- IBM Z https://www.ibm.com/it-infrastructure/z
- IBM Z Resiliency https://www.ibm.com/it-infrastructure/z/capabilities/resiliency
- Storage https://www.ibm.com/it-infrastructure/storage
- Redbook – GDPS Family: An Introduction to Concepts and Capabilities
  http://www.redbooks.ibm.com/abstracts/sg246374.html?Open

GDPS Web site resources

- GDPS: The Enterprise Continuous Availability / Disaster Recovery Solution white paper
- GDPS pre-requisite information
- GDPS training schedule
- GDPS hardware qualification letters

- e-mail: gdps@us.ibm.com