Université IBM i 2018

16 et 17 mai

IBM Client Center Paris

thirty years

Session S55 – Implémenter Hyperswap en environnement GDR

Ludovic Ménard – Jean-Luc Bonhommet

IBM France

ludovic_menard@fr.ibm.com – jeanluc_bonhommet@fr.ibm.com

Overview of IBM Spectrum virtualize High Availability Technologies

Spectrum Virtualize Stretched Cluster

- Base function for many years supported on SVC systems
- Enhanced Stretched Cluster since V7.2
- Provides automatic failover in most failure scenarios
- No specific host software support needed
- Configuration of two sites from a single point
- Wide host and storage support

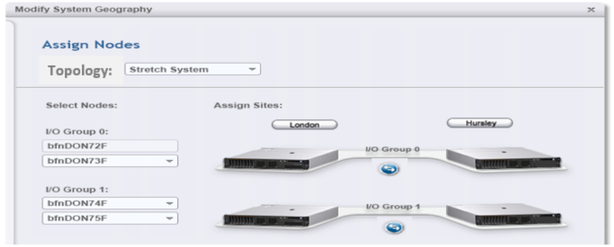

Spectrum Virtualize HyperSwap

- Licensed feature in SVC and Storwize products
- Available since V7.6
- Uses Metro Mirror, non-disruptive volume move technologies
- Provides automatic failover in most failure scenarios
- No specific host software support needed
- Configuration of two sites from a single point
- Wide host and storage support, including internal storage

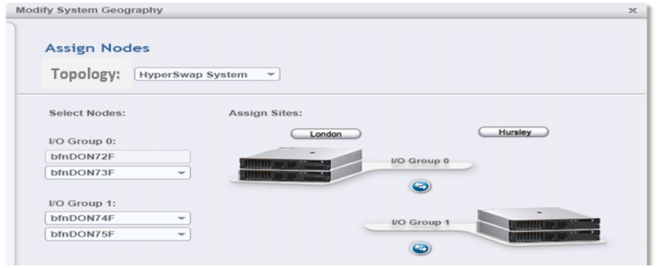

Hyperswap Operation

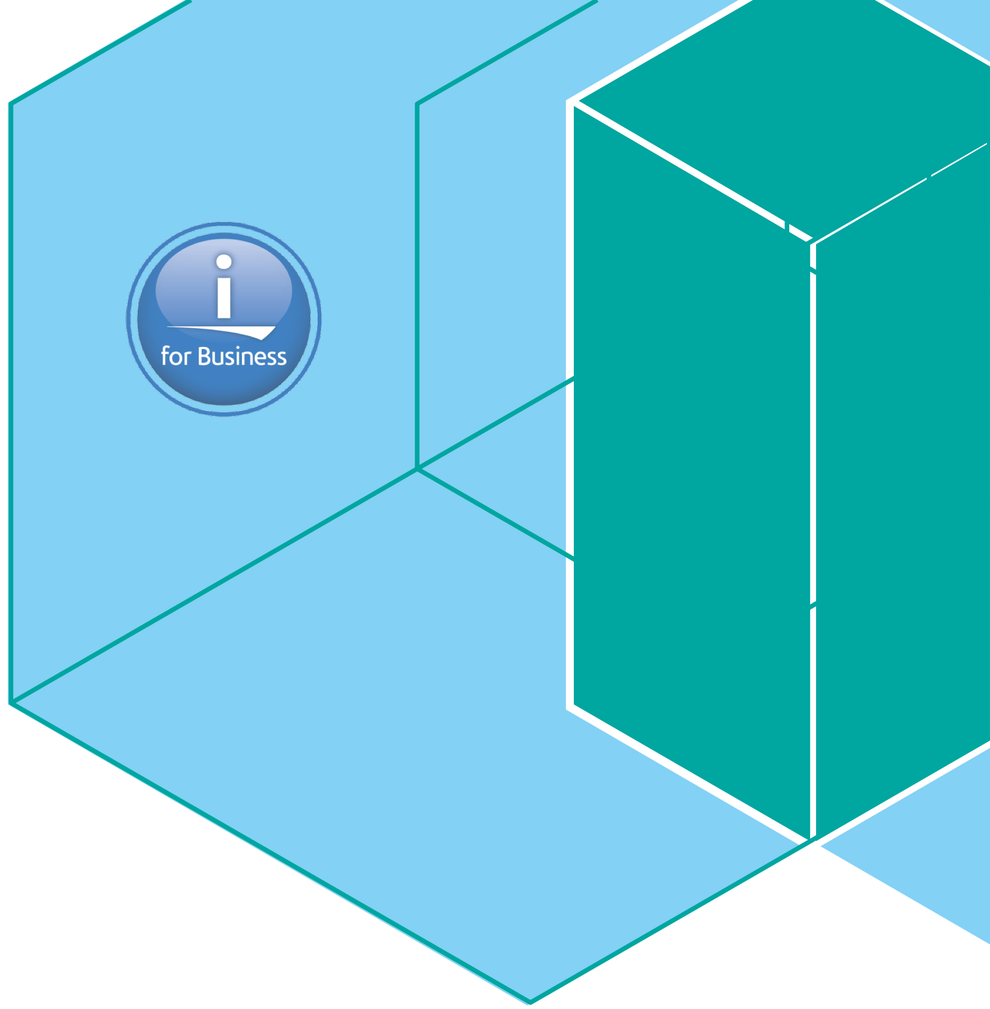

IBM Systems

How HyperSwap Works : Omnidirectional Metro Mirroring

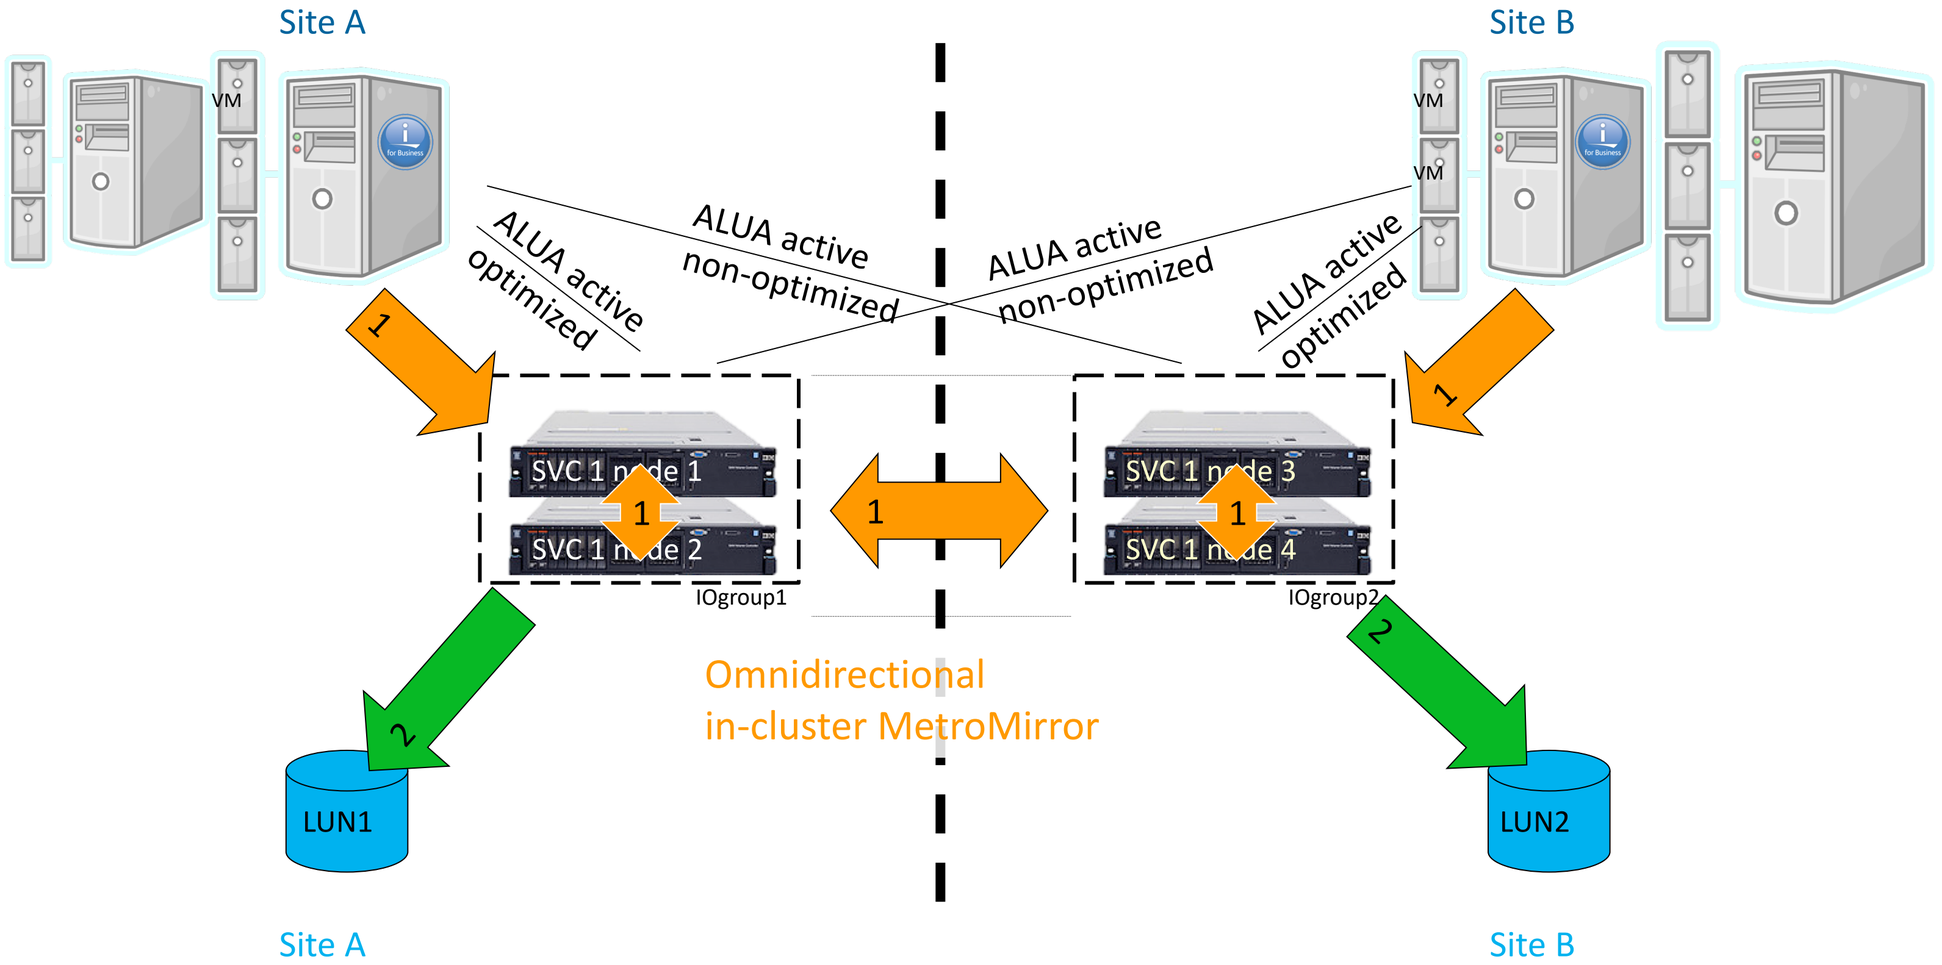

How Omnidirectional Metro Mirror Works

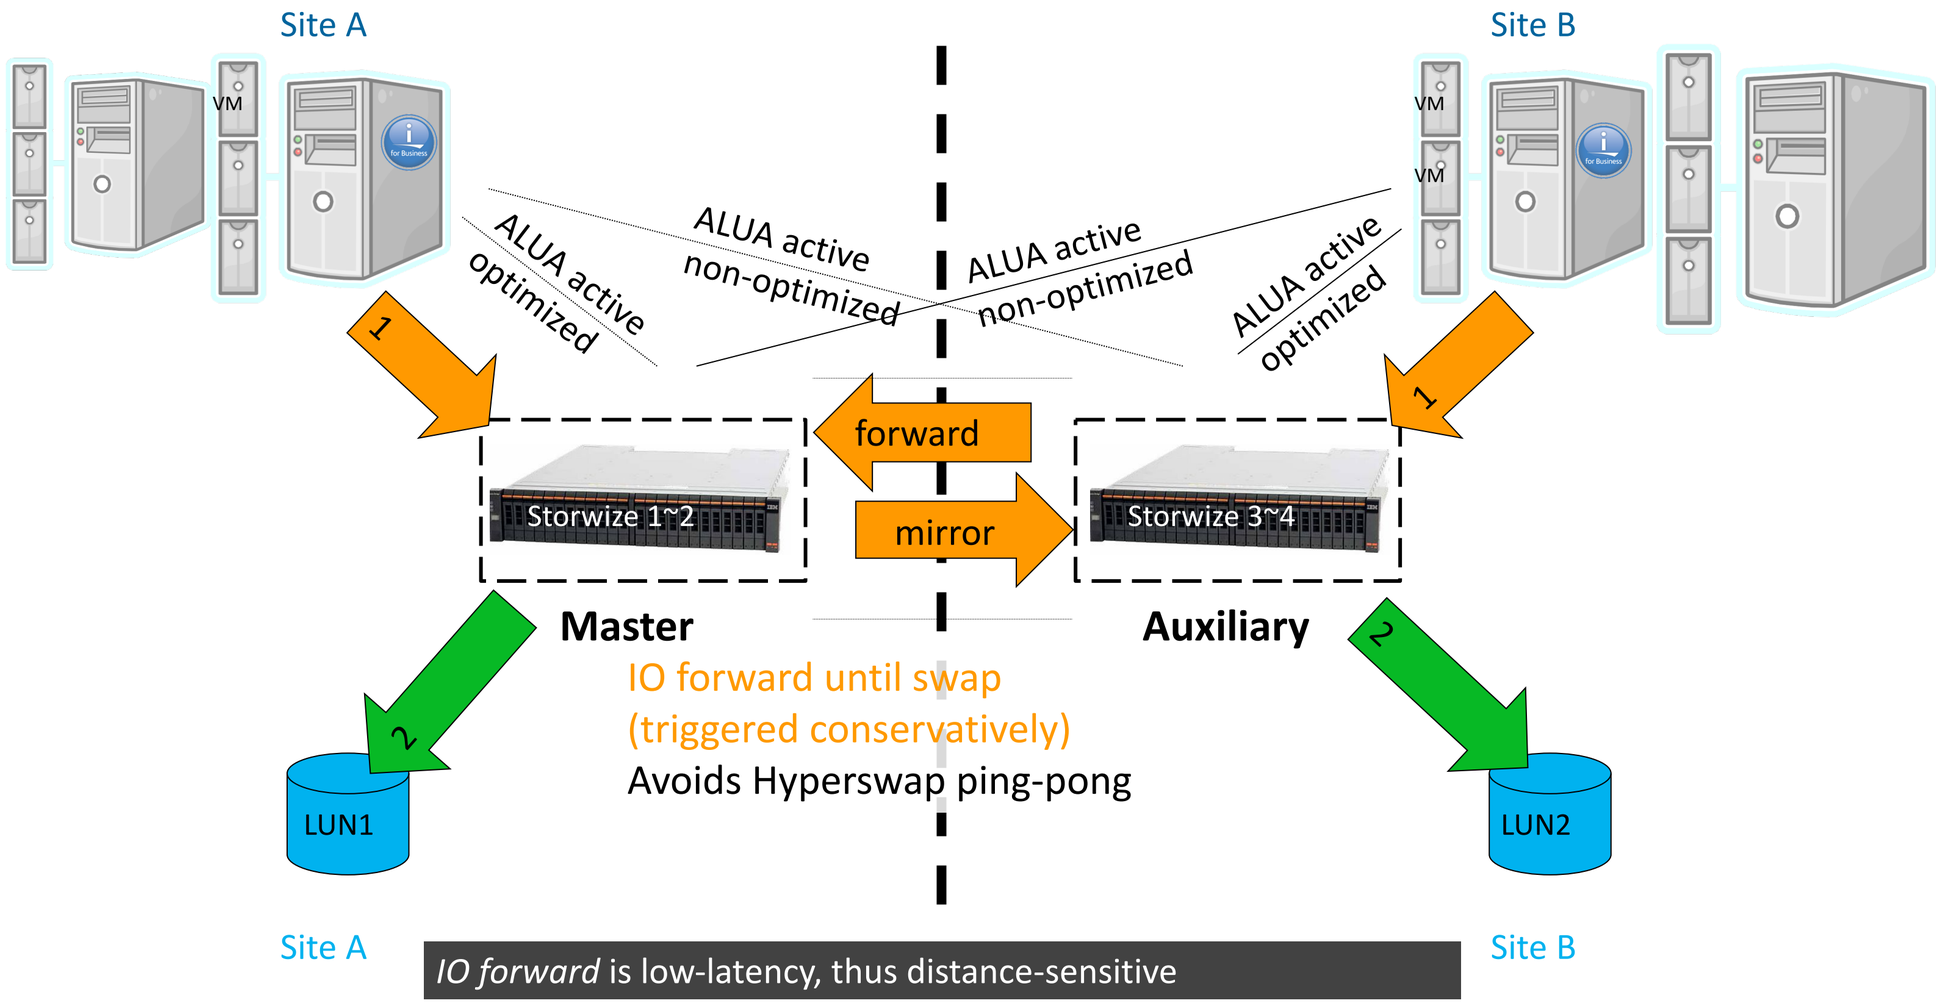

V7.7. limitations: http://www-01.ibm.com/support/docview.wss?uid=ssg1S1005825

HyperSwap reconciliation uses snapshots for undo purpose

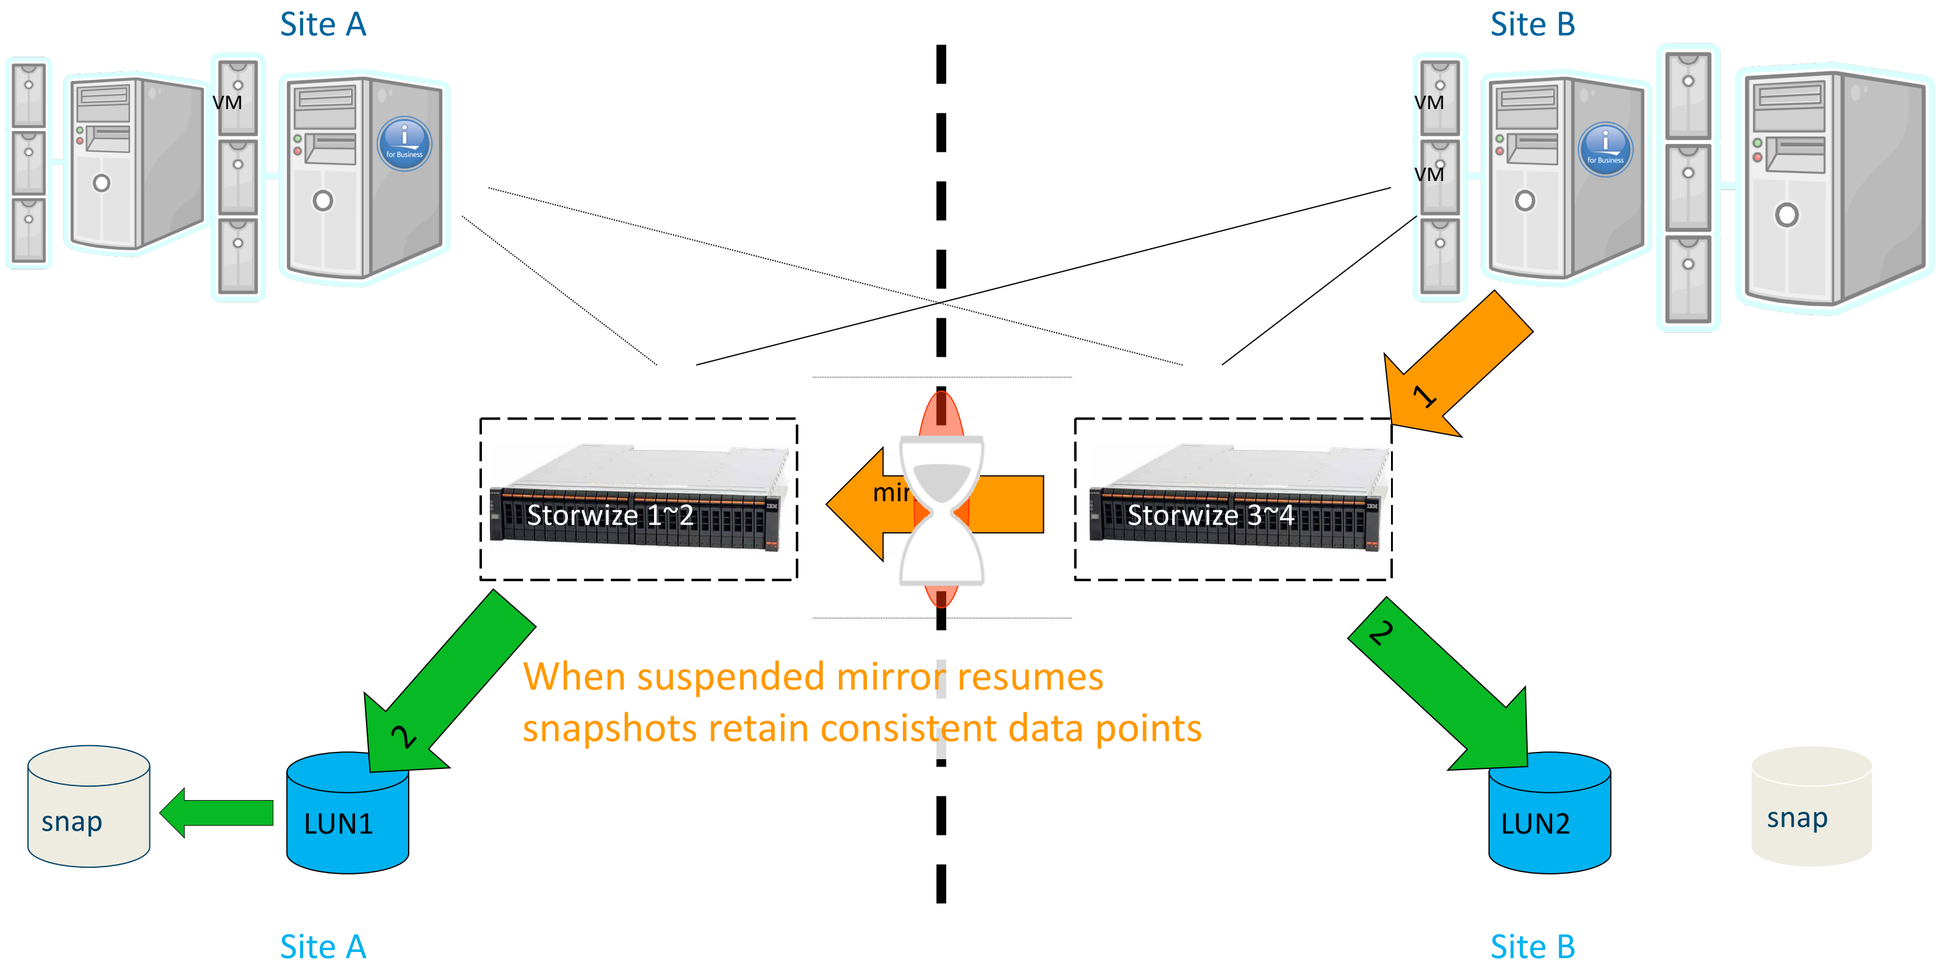

HyperSwap volumes consume two vdisks + two snapshots

IBM Systems

Benefits and limitations of the HyperSwap

- **Benefits:**
  1. Best DR resiliency, adds shadow snapshots to mirror
  2. 'Consistency Groups' concept inherited from MetroMirror
  3. No performance impact to other site during site disaster
  4. Runs natively on Storwize
  5. No manual intervention required: Automatic Failover & Failback
  6. Incremental resynchronization

http://www.redbooks.ibm.com/abstracts/sg248317.html?Open

- **Limitations:**
  1. Two IOgroups / 4 nodes strictly recommended minimum
  2. Four vdisks consumed, MetroMirror consumed (no 3-site replication)
  3. Currently* only supported intracluster = 8 nodes / 4 Storwize

* restriction to be lifted in the future (intercluster)

IBM Systems

HyperSwap : Limits and Restrictions

- **Limits and Restrictions**
  - Max of 1024 HyperSwap volumes per cluster
  - Max capacity is 1PB per I/O group or 2PB per cluster
    - Much lower limit for Gen1 Storwize V7000
    - Run into limit of remote copy bitmap space
  - Can’t replicate HyperSwap volumes to another cluster for DR
  - Limited FlashCopy Manager support
    - Can’t do revFC to HyperSwap volumes
  - **Max of 8 paths per HyperSwap volume same as regular volume**

- Requirements
  - Remote copy license
  - Size public/private SANs as we do with ESC today
    - Only applicable if using ISLs between sites/IO groups

IBM Systems

HyperSwap / SVC Enhanced Stretched Cluster comparison

	SVC Enhanced Stretched Cluster	SVC, V9000 and Storwize HyperSwap
Products that function is available on	SVC only	SVC with 2 or more I/O groups; **V9000**, **Storwize V7000 and Storwize V5000**
Complexity of configuration	CLI or GUI on single system; simple object creation	Simple object creation through GUI and CLI
Sites data stored on	2	2
Distance between sites	Up to 300km	Up to 300km
Independent copies of data maintained	2	2 (4 if additionally Volume Mirroring to two pools in each site)
Technology for host to access multiple copies and automatically fail over	Standard host multipathing driver	Standard host multipathing driver
Cache retained if only one site online?	No	**Yes**
Host-to-storage-system path optimization	Automatic based on host site	Automatic configuration based on host site
Synchronization and resynchronization of copies	Automatic	Automatic
Stale consistent data retained during resynchronization for disaster recovery?	No	**Yes**
Scope of failure and resynchronization	Single volume	One or more volumes, user configurable
Ability to use FlashCopy together with High Availability solution	Yes (though no awareness of site locality of data)	Limited: can use FlashCopy maps with HyperSwap volume as source, avoids sending data across link between sites
Ability to Remote Copy with High Availability solution	One remote copy, can maintain current copies on up to four sites	**No**
Maximum highly available volume count	4096	**1024**
Licensing	Included in base product	**Requires Remote Mirroring license** for volumes. Exact license requirements may vary by product.

IBM Systems

Hyperswap Demonstration

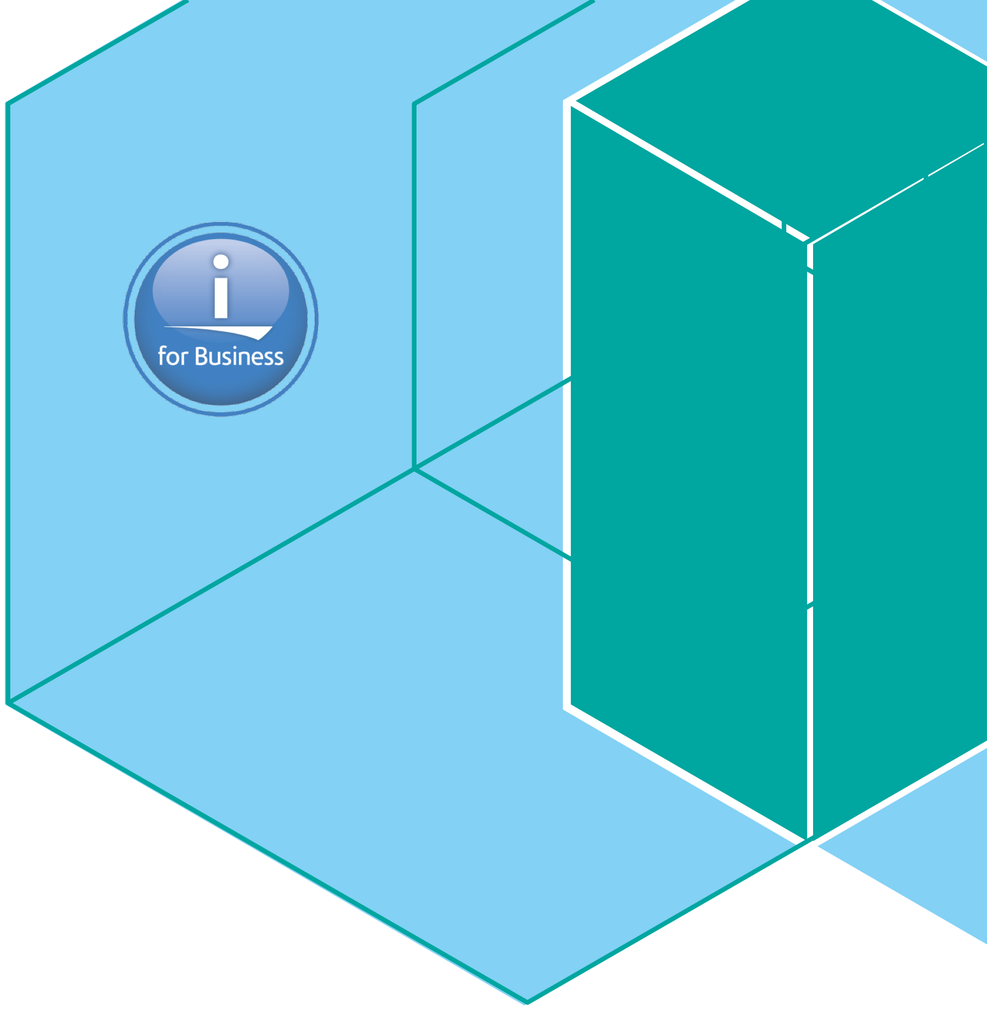

IBM Systems

GDR

Geographically Dispersed Resiliency

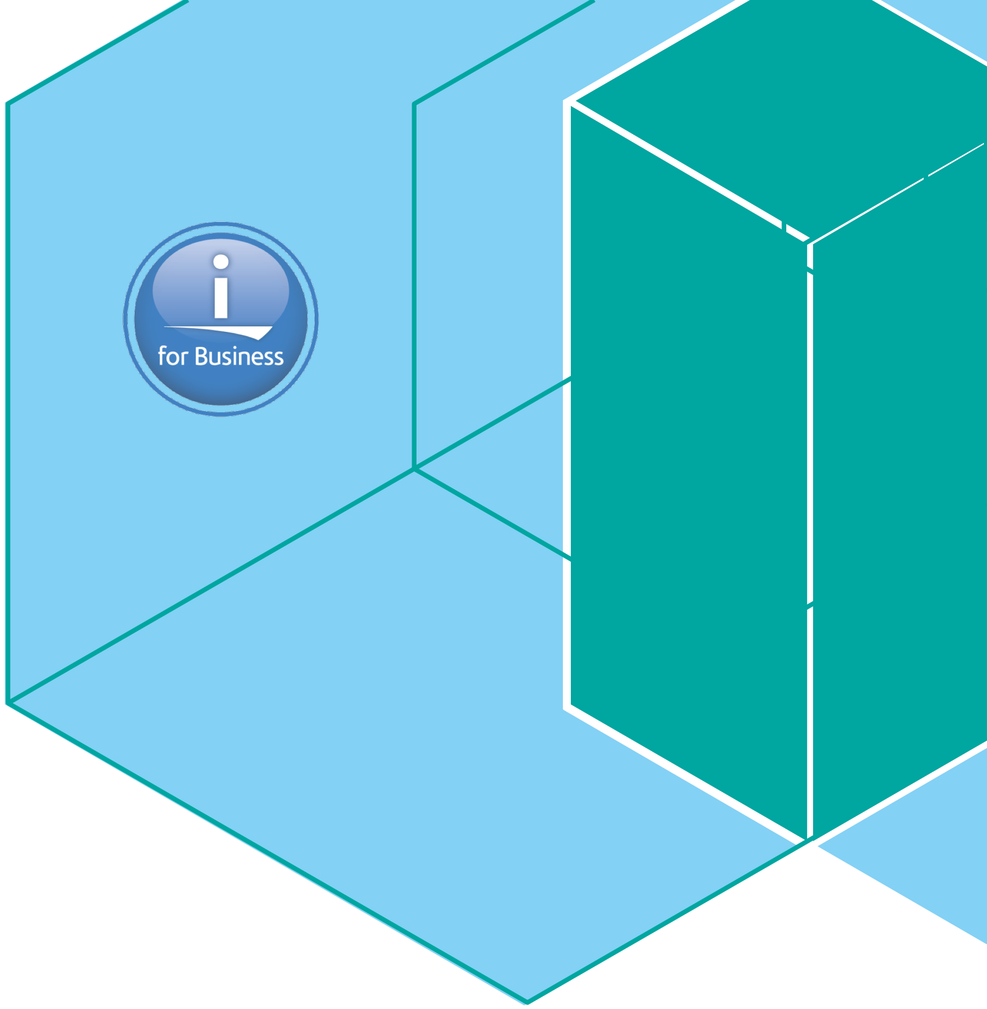

IBM Systems

GDR V1.2 for Power Systems

GDR provides disaster recovery solutions based on virtual machine restart technology

KSYS is the orchestration controller (Kontroller) LPAR

- manages the real-time replication of the VMs between storage servers
- manages the VM restart operations between source and target systems
- installed on an AIX partition

GDR supports PowerVM for Power 7, 8, & 9

- Supports AIX, i and Linux
- DS8000, Storwize, EMC and Hitachi

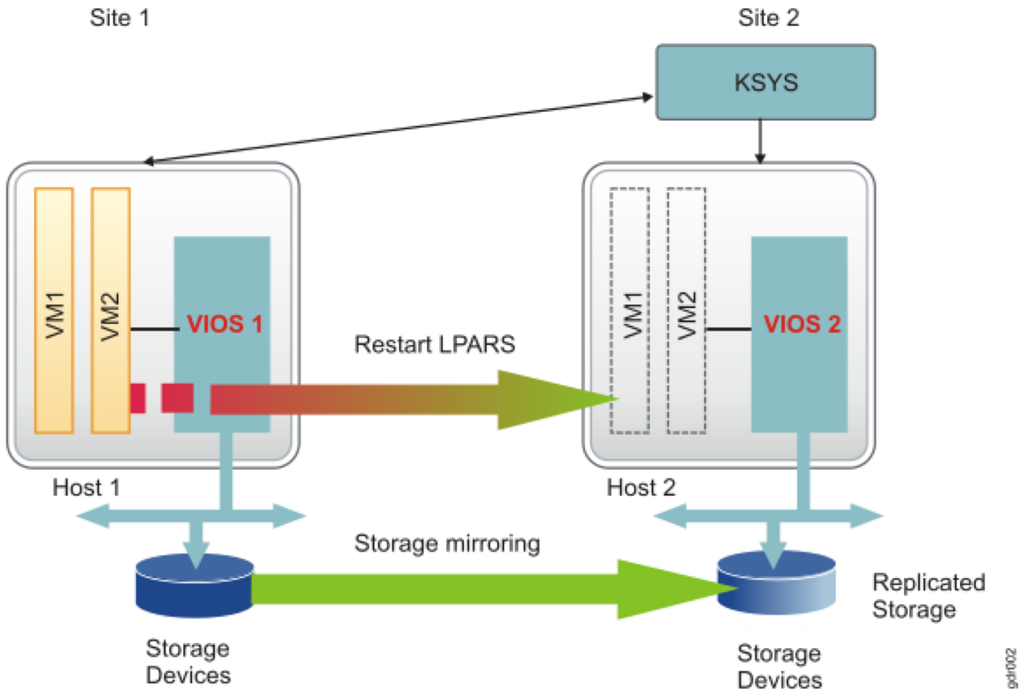

- VMs are replicated real-time via SAN storage
  - synchronous mode replication: RPO of 0
  - asynchronous mode: RPO seconds/minutes
- GDR licenses reside in Ksys partition

Disaster recovery operations made easy & economical

Automation
- Administrator initiated end to end automation
- Eliminates human intervention and error
- Reliable consistent recovery time
- Non disruptive disaster recovery testing
- Adapts to configuration changes (eg: VMs or disks added)

Single point of control
- Centralized status reporting
- Simple uni-command administration

Validation
- Auto discovery of changes to environment
- Daily verification testing
- Email & SMS alerts sent to admin
- Scripting support

Capacity management
- Enterprise pool
- CBU for Enterprise Systems

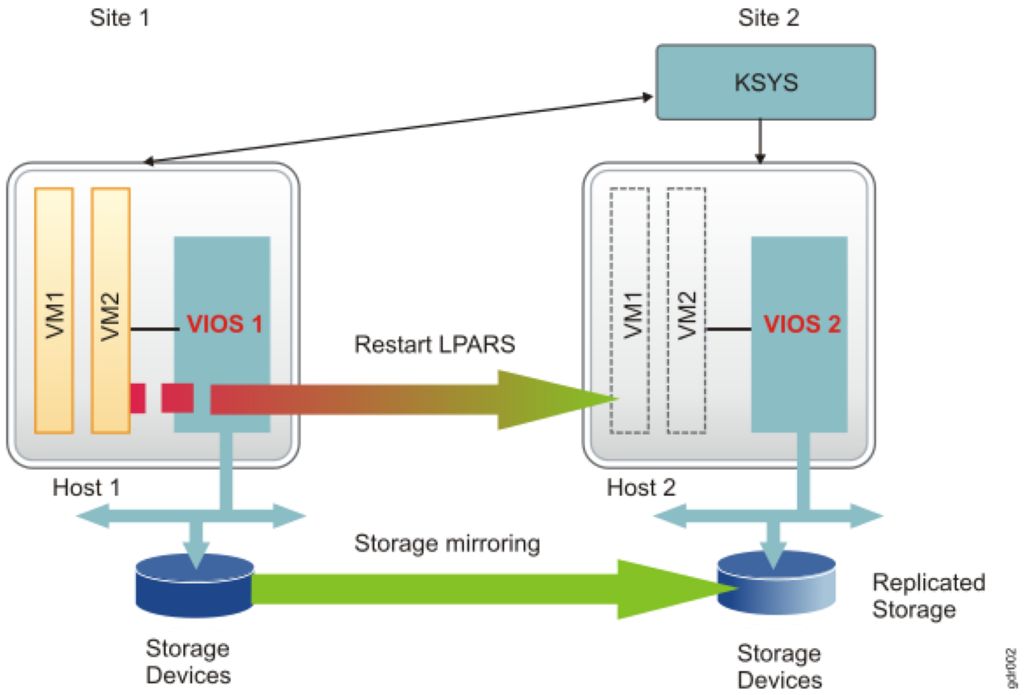

- Active/Inactive (cold standby) topology therefore target DR server requires no licensing (IBM LPPs)
- GDR licenses reside in a KSYS partition running on an instance of AIX typically deployed in the DR center

GDR 1.2 Prerequisites

Guest OS in VMs	1.AIX: V6 or later 2.IBM i: V7.2 or later 3.Linux: • RedHat(LE/BE): 7.2 or later • SUSE(LE/BE): 12.1 or later • Ubuntu: 16.04
VIOS	VIOS 2.2.6 (2017) + FIXES
HMC	V8 R8.7.0 (2017) + SP1 V9 R9.1.0 (2018)
1.EMC Storage: SRDF 2.DS8K: Global PPRC 3.SVC/Storwize: Metro or Global 4.Hitachi VSP, G1000, G400: Universal Copy	1.VMAX family, Solutions Enabler SYMAPI V8.1.0.0 2.DS8700 or later DS8000® storages (DSCLI-7.7.51.48 or later) 3.SVC 6.1.0 (or later) Or Storwize (7.1.0 or later) 4.Universal Copy (CCI version 01-39-03 or later)
KSYS LPAR	AIX 7.2 TL1 SP1 or later

GDR V1.2

GDR Target Markets

- Power Systems customers desiring a low cost highly automated disaster recovery solution
  - AIX, IBM i and Linux environments
  - PowerHA for AIX environments with manual operations and/or IP based replication
  - IBM i environments deployed using logical replication software
  - SAP HANA on Power Systems

Offering construct

- Based on PowerVM and SAN storage replication of the production VMs
- Two tiers small and medium
- The number of GRD licenses = number production cores supporting replicated VMs

GDR for Power Systems – 5765-DRG (orderable via AAS on Aug 11)		
Software tier	small processor group/core	medium processor group price/core
List price for managed cores (US pricing example)	**$1020**	**$1575**

U.S. prices

- Priced per processor core
- First year SWMA included
- Second year and later SWMA: 20% of license cost/year

GDR V1.2 for Power Systems

Virtual Machine replication & restart solution for disaster recovery

- AIX, i and Linux Power System customers

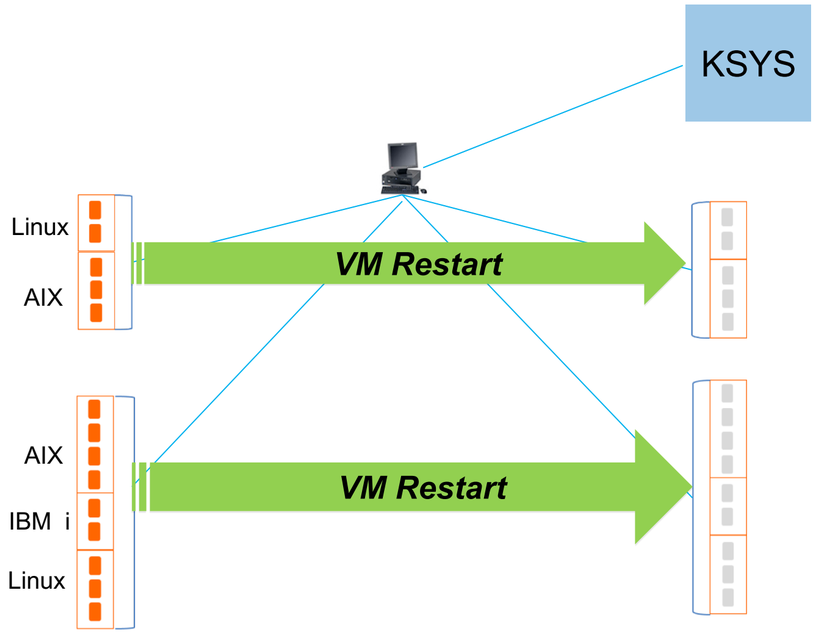

Production Site
14 managed cores

Recovery Site
0 managed cores

- Storage based VM replication
- Ksys (orchestrator partition) manages replication and VM restart and resides on an AIX partition
- GDR licenses (contractual, not key based) are tied to the KSYS partition, not on the source or the target production systems

GDR for Power Systems – 5765-DRG (orderable via AAS on Aug 11)		
Software tier	small processor group/core	medium processor group price/core
List price for managed cores (US pricing example)	**$1020**	**$1575**

U.S. prices

Priced per managed production core on the production system
No IBM license required on recovery system
First year SWMA included
SWMA renewal 20% of list price

GDR Pricing Example

Using our example of 14 managed cores on the production servers

- Production servers in this example are in the small processor group therefore $1020/core
- 14 managed cores = 14 GDR entitlements to be installed on the KSYS partition
- Total price; $15,300 (U.S. prices)
- DRaaS hosted customers can conduct disaster recovery testing 4 times a year for 72 hours per test

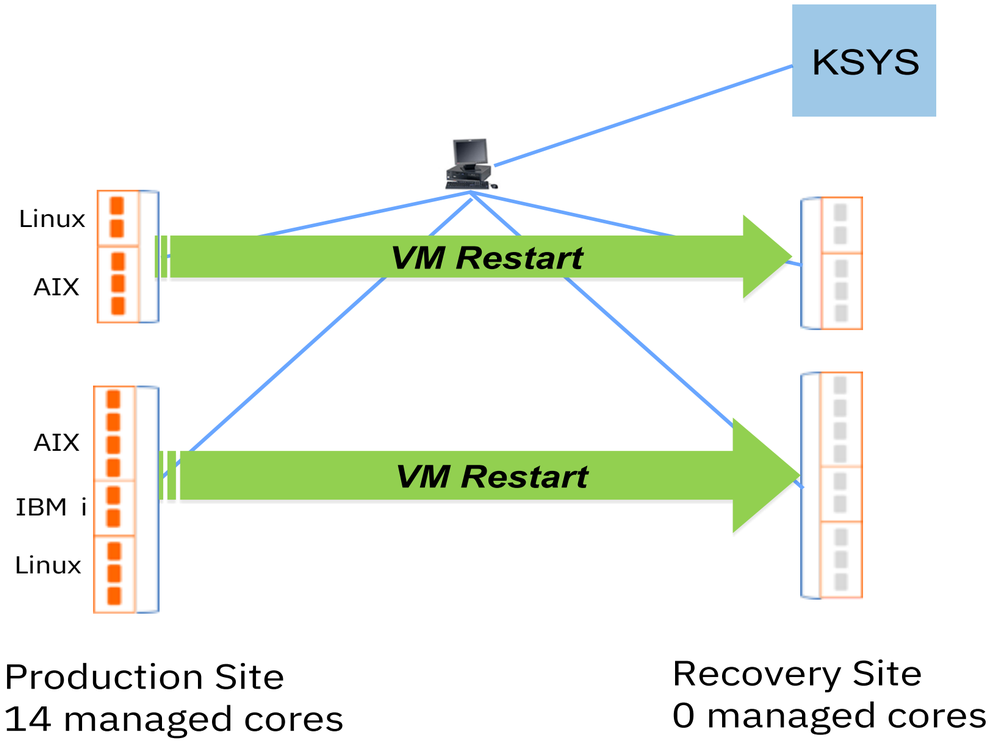

GDR for Power Systems – 5765-DRG (orderable via AAS on Aug 11)		
Software tier	small processor group/core	medium processor group price/core
List price for managed cores (US pricing example)	**$1020**	**$1575**

PowerHA vs GDR

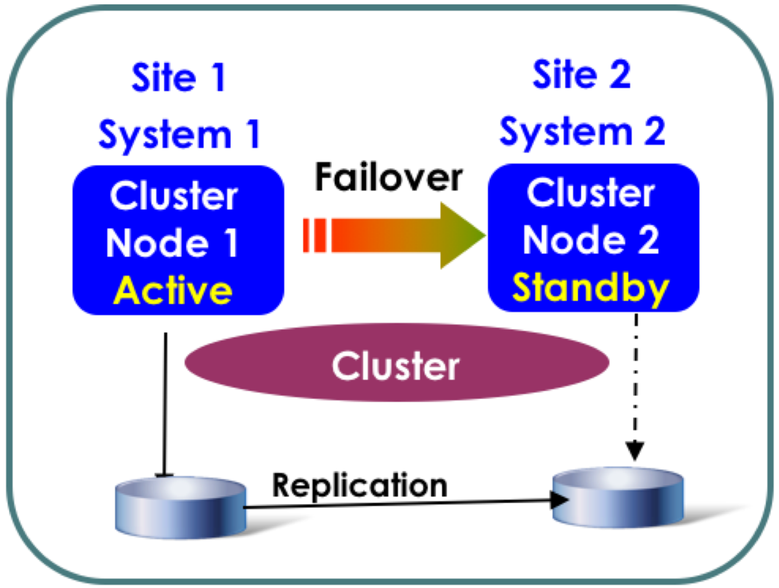

Fig 1: **Cluster DR** Model

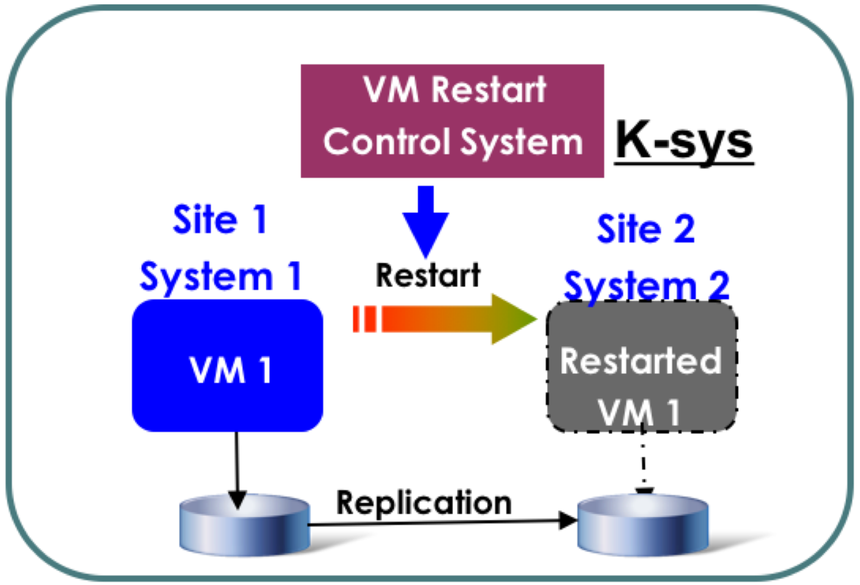

Fig 2: **VM Restart DR** Model

Deployment Approach	**Deployment inside each VM (complex)**	**Deployment outside VMs (simpler)**
Workload Failover Time	**Fast**	**Fast Enough (VM Reboot)**
Cost	**High (duplicate SW & HW)**	**Low (No SW duplication)**

Terms for DRaaS deployed GDR

GDR will prove to be the most cost effective easy to use solution for DRaaS operations

- Optimized for multi-tenant DRaaS hosting
- DRaaS GDR servers may only be used for disaster recovery operations, no production workload
- No AIX, IBM i/LPPs required on the DRaaS system
- Hosted customers can conduct disaster recovery testing 4 times a year for 72 hours per test
- In the event of an actual disaster, client may run on DRaaS server for up to 70 days
- DRaaS provider can acquire the required GDR LPPs to run on their KSYS partition
- DRaaS server may be in a software tier greater than the customer production servers *

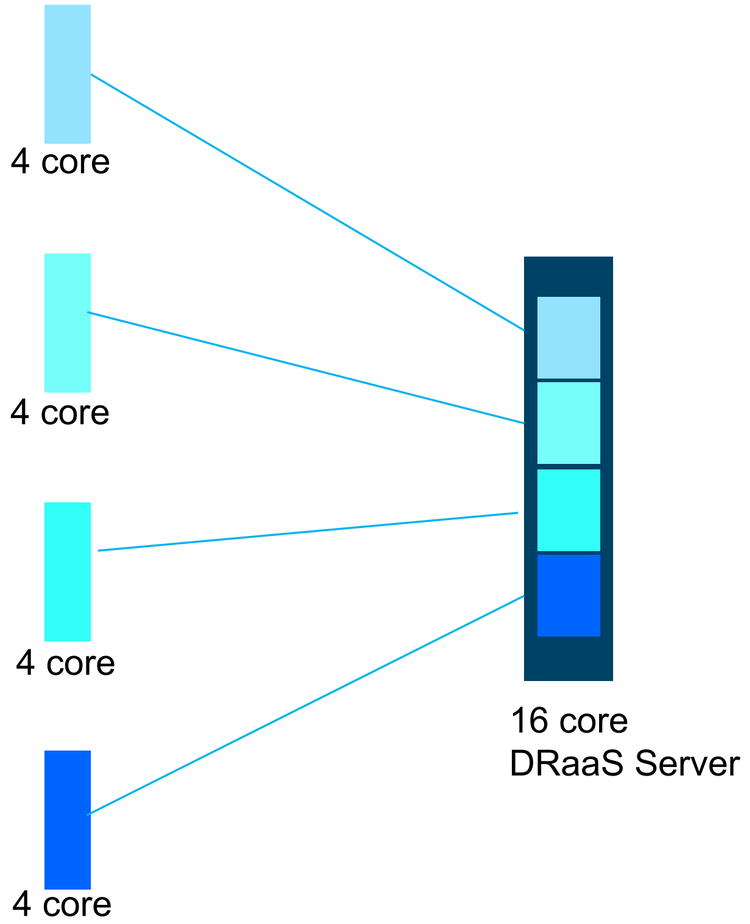

* subject to specific terms and conditions

PowerHA for IBM i + GDR for DR (animation)

automated disaster recovery operations with a PowerHA IBM i standard edition cluster

Host level VM restart of production PowerHA cluster in Madison to the Chicago disaster recovery center

Why do DR testing manually when GDR automates it for you ?

Switch VLAN ids for VMs on the back up site

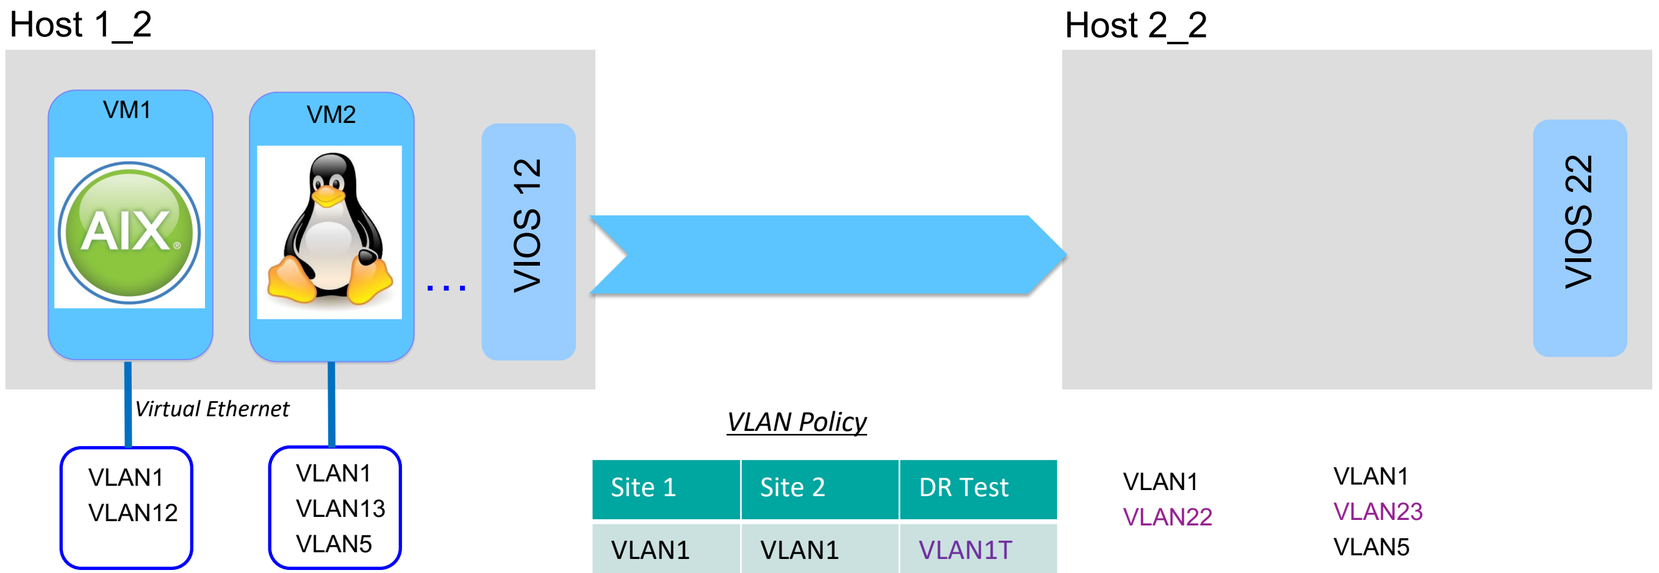

VLAN Policy

Site 1	Site 2	DR Test
VLAN1	VLAN1	VLAN1T
VLAN12	VLAN22	VLAN2T
VLAN13	VLAN23	VLAN3T
VLAN5	VLAN5	VLAN5T

- Enables VLAN based network customization across sites
- Pre-defined policy table that admin sets up
- As part of an admin initiated failover, VLANS are changed when LPARs are started at backup site

Failover Rehearsal: non disruptive disaster recovery testing

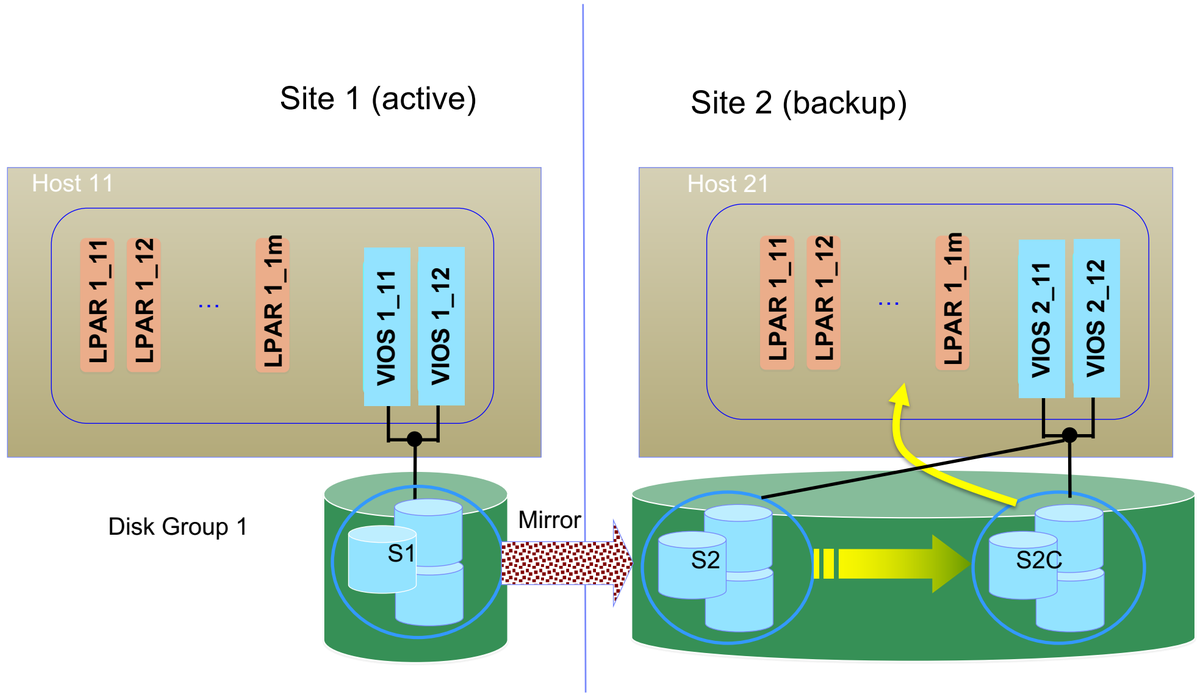

A point in time copy (ie Flash Copy) is created to start VMs on the back up system for DR testing or backup operations
Enables IT operations to validate disaster recovery compliance without disrupting production
Network isolation needs to be established by the administrator (admin can design and use the test VLANs for test VMs)

SAP HANA Disaster Recovery

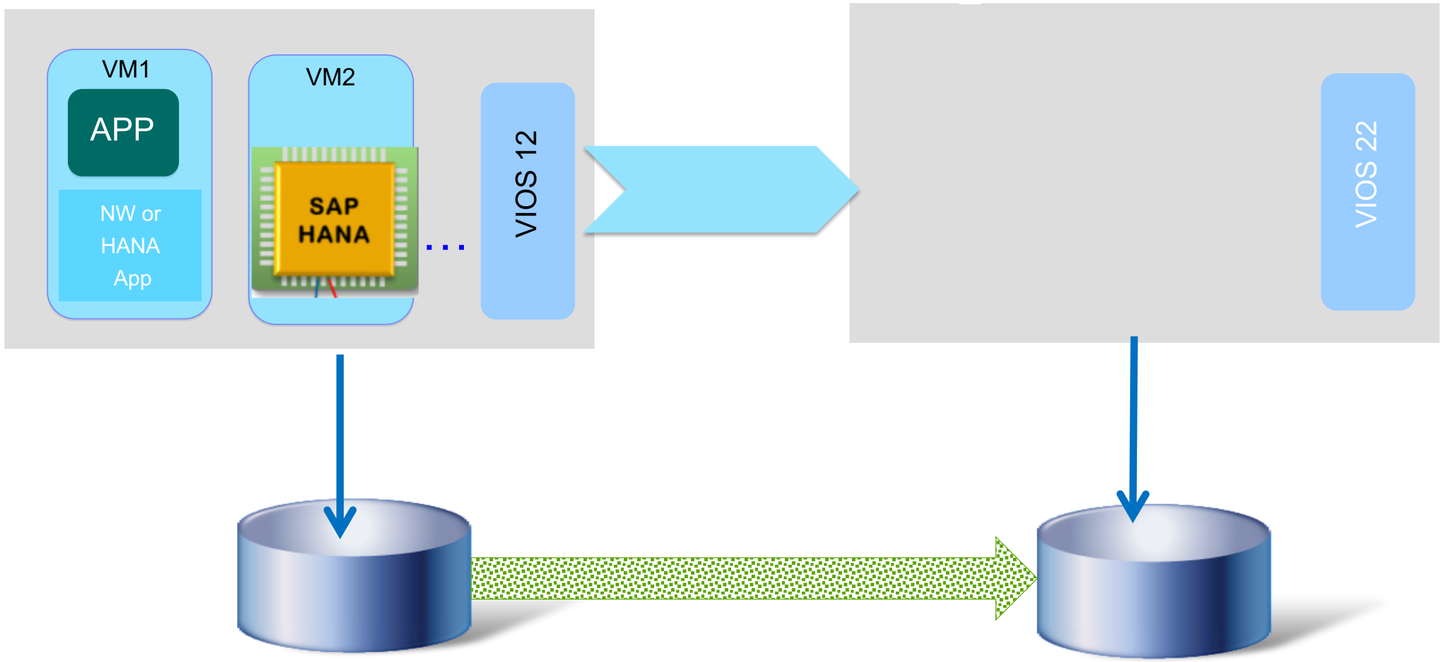

SAP HANA VMs + NetWeaver VM back & forth DR failovers

SAP HANA and its workload will be checked for recovery + functionality (not performance)

Guidelines: http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102502

GDR support for ”shared” storage configurations

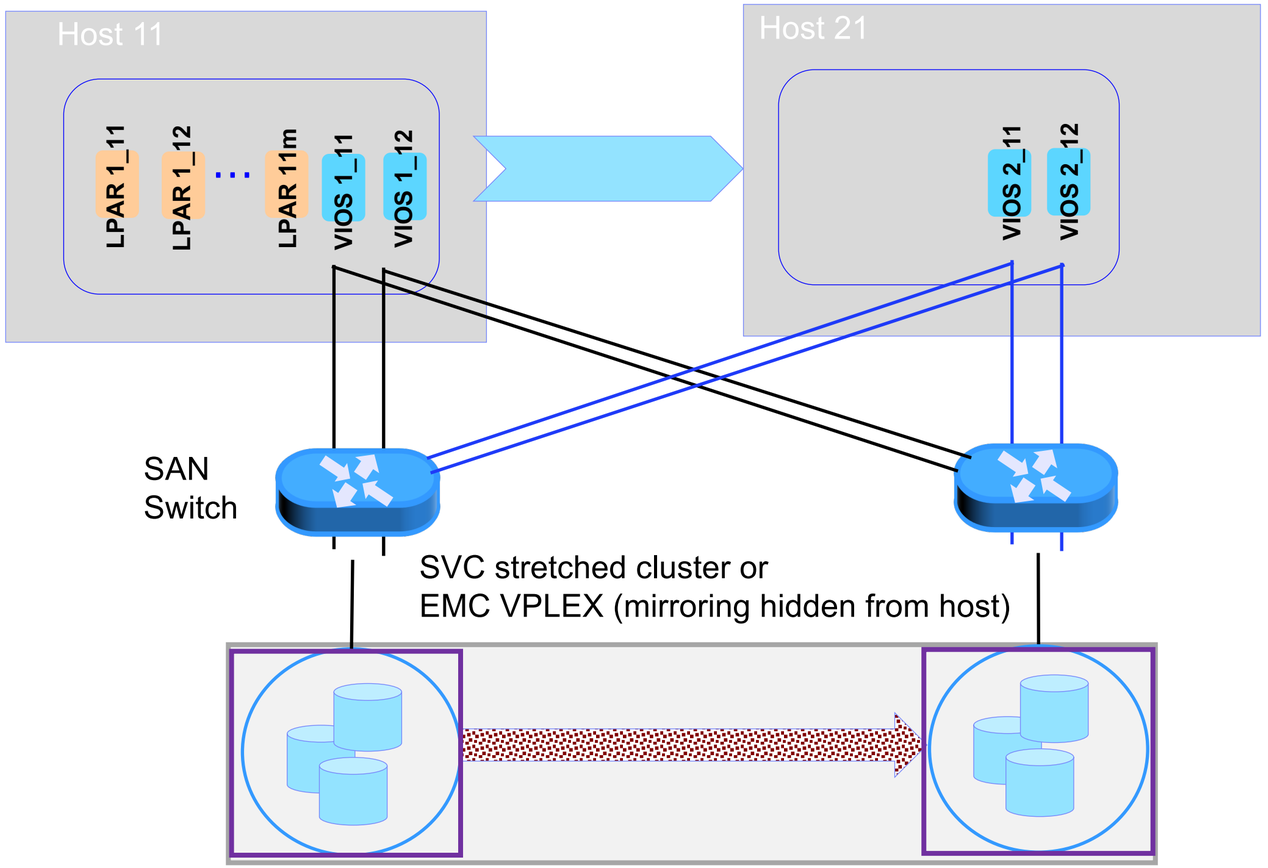

- Mirror management hidden from Host/VIOS, mirrored storage subsystem appears as a shared storage configuration
- Planned and unplanned failovers
  - KSYS thinks it is un-mirrored shared storage and will start VMs on backup site etc
- Limitations
  - Can not support the DR rehearsal capability or snapshots for backup operations
  - Short distances, synchronous mirroring

Merci de votre attention

N’oubliez pas de remplir
le questionnaire
de satisfaction !