

Université IBM i 2017

17 et 18 mai – IBM Client Center de Bois-Colombes

S45 – Solutions de disponibilité sur IBM i.

Jeudi 18 mai – 15h15-16h45

Joseph Esnault
IBM CTS Power
joseph_esnault@fr.ibm.com
GSM : 06 84 81 50 06



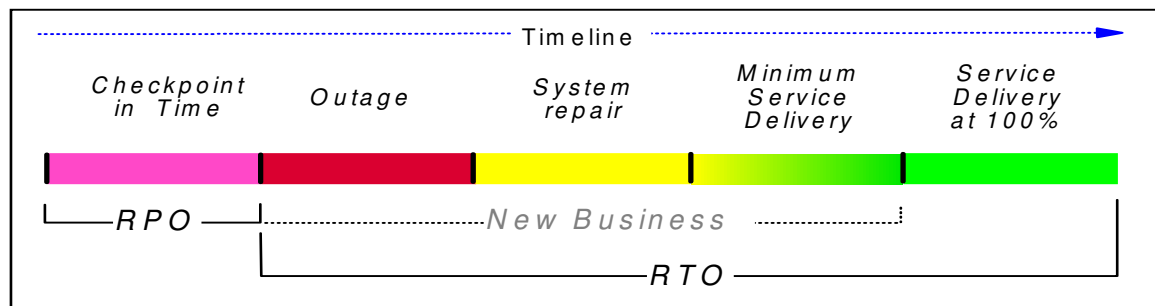
Agenda

- IBM i Full System vs. IASP Replication
- IBM i Full System Replication Solutions
 - Metro Mirror with optional HyperSwap
 - Global Mirror
 - Multi-Target PPRC
 - SVC Split-Cluster
 - FlashCopy
- IBM i Clustering and IASP Replication with PowerHA SystemMirror for i
 - Metro Mirror & Global Mirror
 - FlashCopy
 - Multi-Target PPRC with IASP Manager
 - LUN Level Switching with SVC Split-Cluster
 - DS8000 IASP HyperSwap

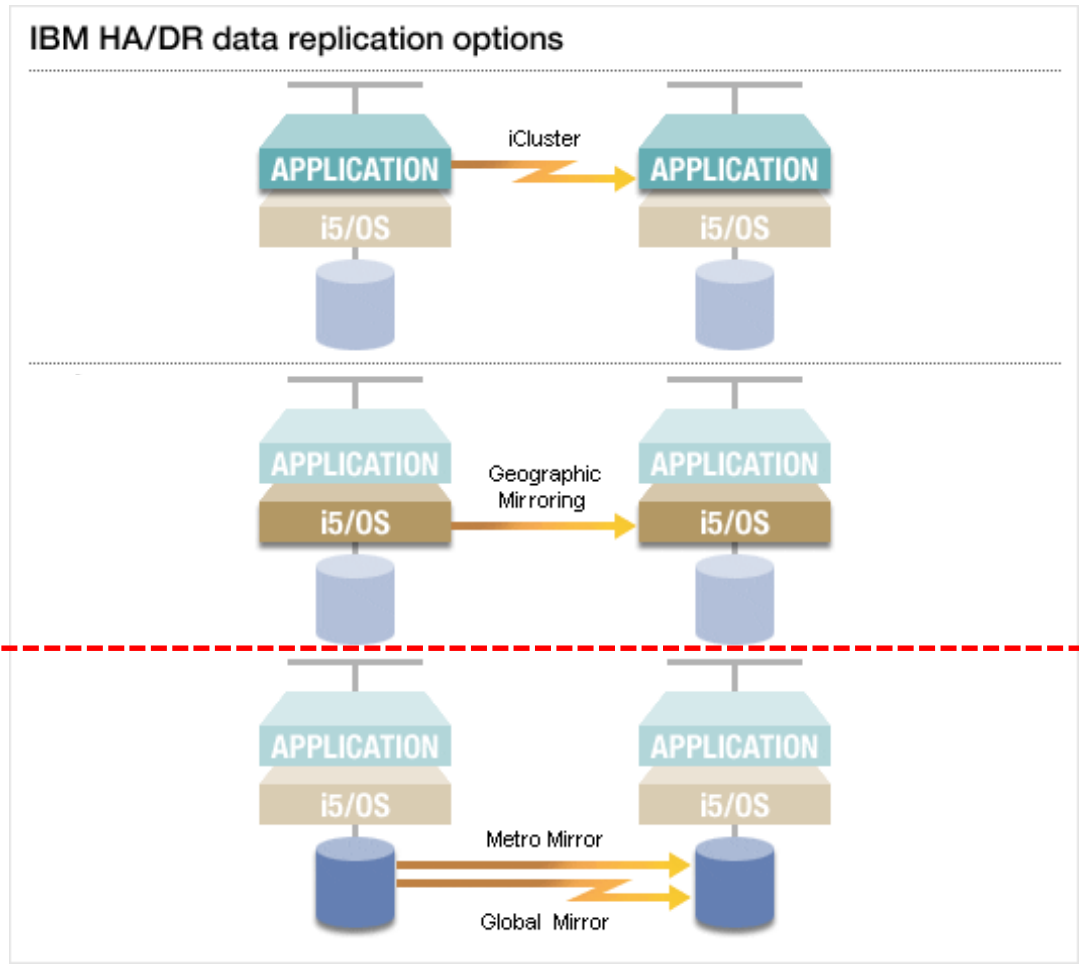
Votre objectif.

- **Assurer une plus grande disponibilité des serveurs Power Systems IBM i face à différents type de situation :**
 - **Arrêts planifiés** : par exemple une mise à jour système d'exploitation ou logiciels...
 - **Arrêts non planifiés** : par exemple, une défaillance matérielle.
 - **Réduire les fenêtres de sauvegarde** : par exemple, sauvegarde sur base de données répliquées ou via snapshot.
 - **Sinistres** : par exemple, une perte d'un Data Center, coupure générale d'alimentation électrique,...

- **Déterminer les objectifs de disponibilité de votre système d'informations basé sur Power i :**
 - Recovery Time Objective – **RTO** : Temps maximum supportable pour remettre à disposition des utilisateurs un environnement applicatif.
 - Recovery Point Objective – **RPO** : Niveau acceptable de perte de données.
 - Network Recovery Objective – **NRO** : Temps nécessaire pour basculer le réseau



IBM i Data Resiliency Solutions



Logical replication

Asynchronous journal-based replication over TCP/IP
Disk subsystem agnostic

OS based replication

Synchronous or asynchronous IBM i based replication over TCP/IP
Disk subsystem agnostic

Storage-based replication

Synchronous or asynchronous storage-based replication over Fibre Channel or IP with DS8000 or SVC/Storwize

Focus of this presentation

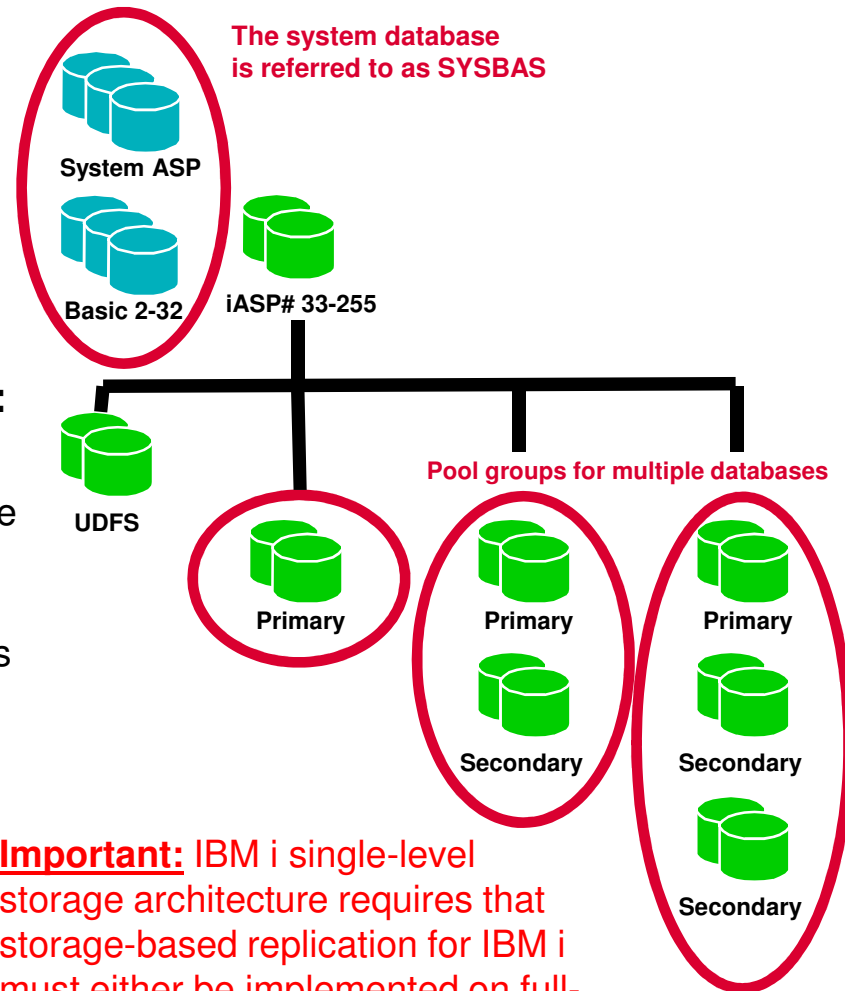
IBM i Storage Architecture – Auxiliary Storage Pools

Auxiliary Storage Pools (ASPs):

- ASPs are created from one or more disk units to isolate data / workload to the same storage area
- System ASP (ASP 1) contains the load source disk unit and exists on every IBM i system
- Optionally basic user ASPs 2-32 or/and independent user ASPs (IASPs) 33-255 can be created

Independent Auxiliary Storage Pools (IASPs):

- a primary disk pool and zero or more secondary disk pools which are grouped together and share the same database
 - similar as System ASP and basic ASPs
 - e.g. primary IASP for libraries and database files and secondary IASP for journals and journal receivers
- IASPs can be varied on/off (as groups together)
- IASPs can be made switchable to be switched (as groups together) between IBM i cluster nodes
- IASPs do not overflow (unlike user ASPs which overflow into the system ASP)



Important: IBM i single-level storage architecture requires that storage-based replication for IBM i must either be implemented on full-system or IASP level.

IBM i Full System vs. IASP Replication

- IBM i full system replication means replicating the complete IBM i SYSBAS disk storage (ASPs 1 to 32) without any independent ASPs
 - considered disaster recovery (DR) solution with standby IBM i backup system at the remote DR site
 - to work with IBM i backup system it needs to be manually IPL-ed after a storage failover
 - recovery time objective (RTO) mainly determined by abnormal IPL time

- IBM i independent ASP (IASP) replication means replicating the IASP(s) only between two IBM i systems in an IBM i cluster
 - considered high availability (HA) solution with IBM i backup system up and running at the remote HA/DR site
 - requires IBM PowerHA SystemMirror for i high availability software
 - provides a high degree of automation due to IBM i cluster and storage failovers managed by PowerHA
 - RTO shorter than with full system replication and mainly determined by abnormal IASP vary-on time
 - allows for staged IBM i OS upgrades or PTF installations due to IBM i production and backup server(s) having their „independent“ SYSBAS storage space

IBM i Logical Replication / Geographic Mirroring

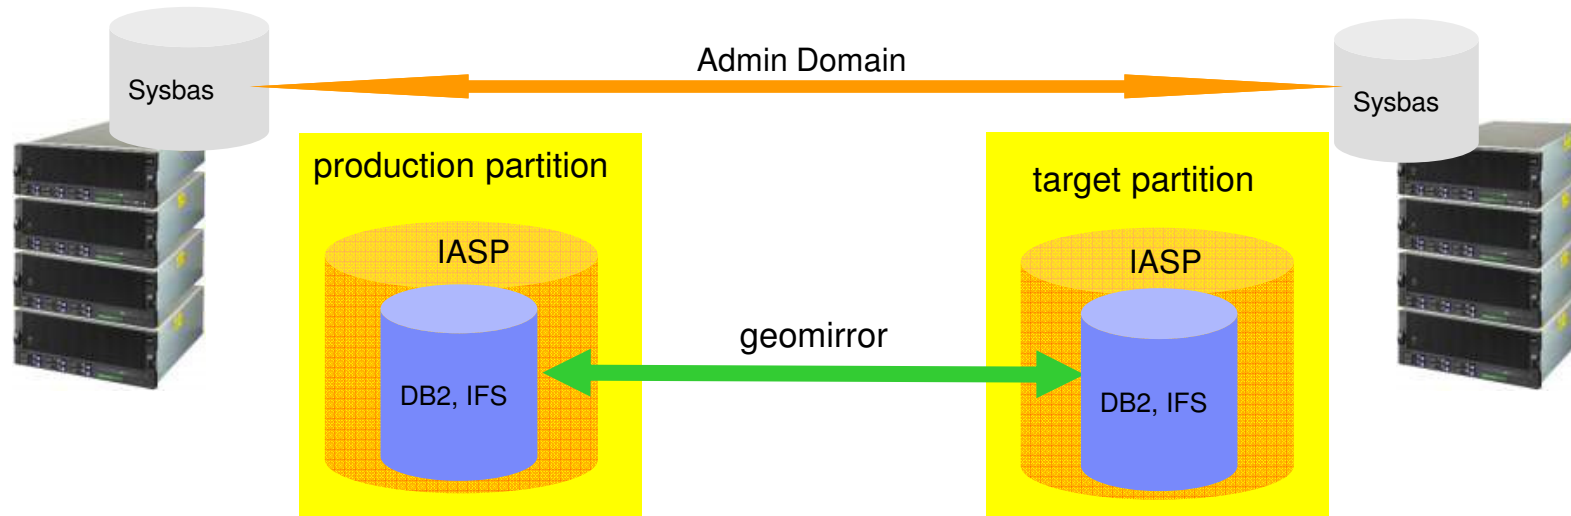
■ IBM i Logical Replication

- disk subsystem agnostic traditional ISV software solution for replication of selective SYSBAS objects to one or more IBM i backup systems
- synchronous or asynchronous object level replication over TCP/IP typically using IBM i remote journaling with journal apply jobs running on the backup system
- data on backup system usable for read-only type of work like queries or data backup (with prior stopped apply jobs)
- needs some operator monitoring for objects to be replicated and their up-to-dateness on backup system as journal apply jobs may lag behind or data simply wasn't replicated due to transmission errors

■ IBM i PowerHA XSM Geographic Mirroring

- synchronous or asynchronous IASP based IBM i memory page segment replication by IBM i System Licensed Internal Code over TCP/IP to another IBM i system in the cluster
- data on backup system usable only while IASP is detached (replication stopped) (can be combined with FlashCopy e.g. for online backup)
- full resync required after failure of online IASP
- rather suited for smaller IASP environments (up to ~2 TB) due to required time for full synchronization which much depends on the number of files / objects

PowerHA – Geomirroring – HA/DR clustering



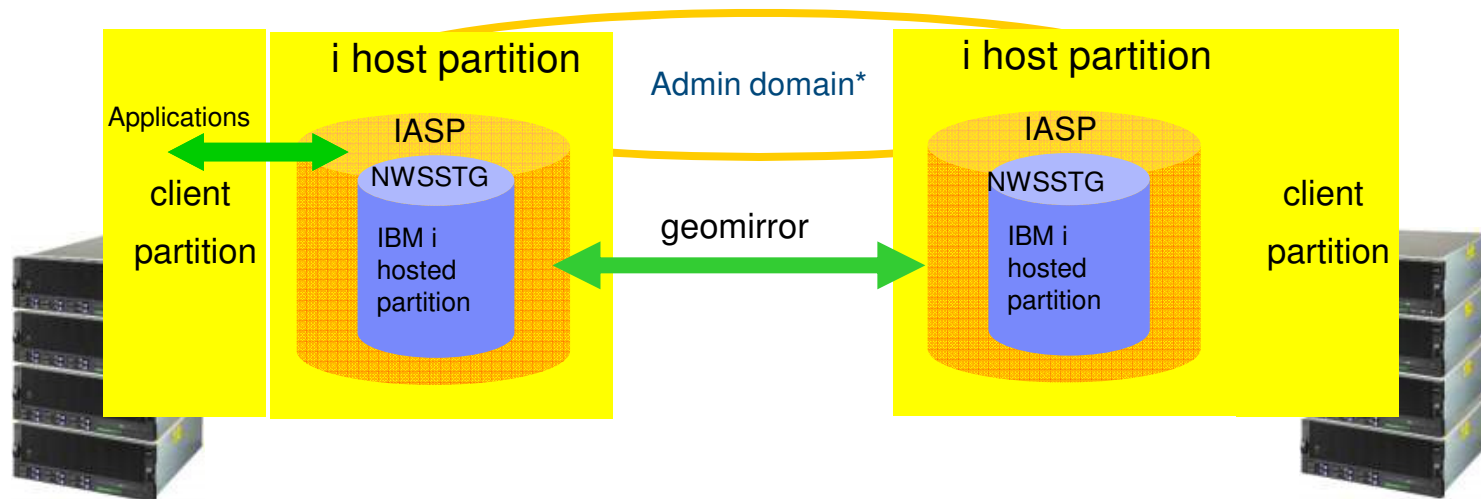
- PowerHA geomirror cluster (typically with internal disk and < 4 Tbytes)
- Complete HA/DR coverage for all outage types (hardware, middleware, operator error)
- Off line back-up followed by source side /target side tracking change resynchronization
- Both bandwidth and network quality are important.
- Synchronous mode up to 40 KM, production and target always identical
- Asynchronous mode unlimited distance, production and target ordered and consistent

PowerHA with geomirroring

- PowerHA with geographic mirroring is popular with our smaller IT shops
 - Its easy to use and requires little day to day monitoring and maintenance.
 - Typically used by shops with under 2 Tbytes and on internal disk
 - After an unplanned outage the target data is replicated back to the source therefore the larger the IASP the longer the resync time
 - Doing a tape backup off of the target works great, *Quiesce the IASP on the current production copy system if you want to get a “clean transaction boundary. If you choose to do the quiesce (recommended), do it during a quite time.*

- Geomirroring is IBM i mirroring of the data writes to the IASP source and target over IP
 - Referred to as host based replication (as opposed to storage based replication)
 - IBM i storage management sends the memory pages to both the remote and local server either synchronously or asynchronously
 - Synchronous distance is typically under 30 or 40 KM but distance is driven by application response time requirements
 - Bandwidth and quality of service are both important. When we’ve seen issues they usually come down to network quality.

Geomirroring – i hosting i remote VM restart for DR



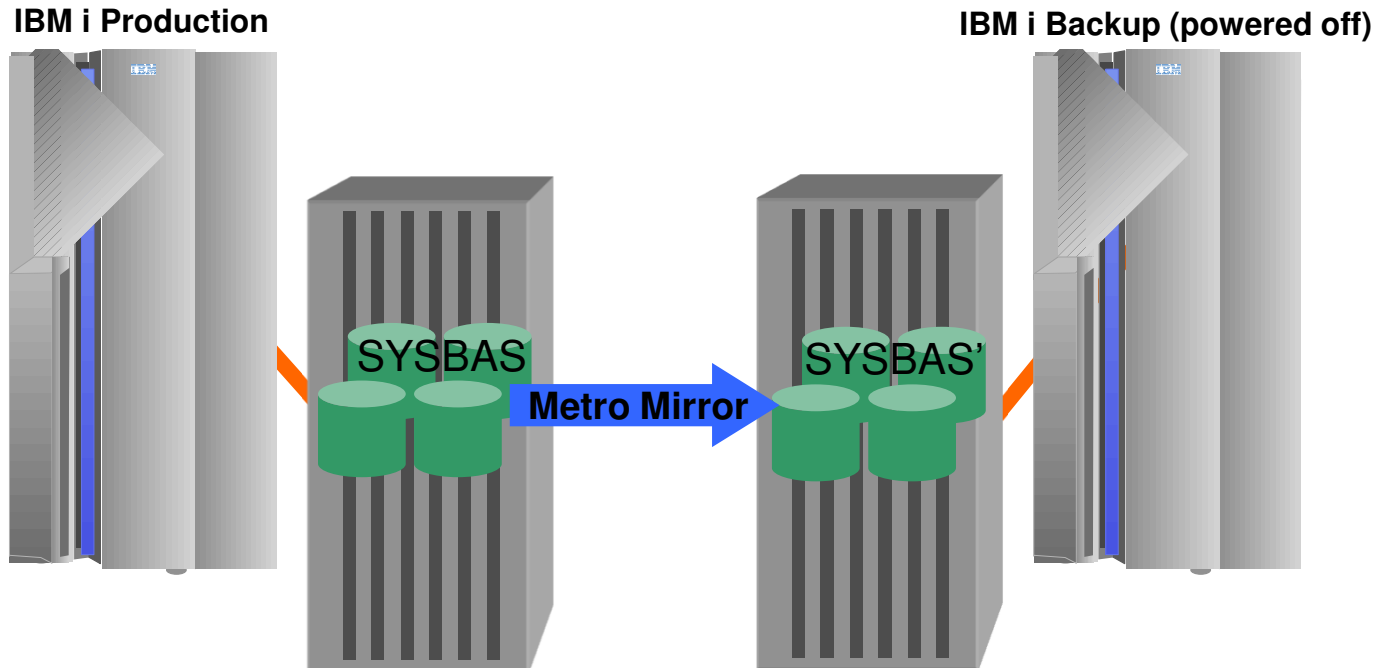
* Admin Domain on host is optional

- Non-cluster PowerHA configuration, full system replication, this is disaster recovery setup, this is not a HA solution
- IBM client placed into a network storage space which is placed into an IASP
- Guest and host partition must be shut down before remote host and client can be restarted
- Limitation: no heart beating, can't do concurrent OS upgrades, is more resource intensive than a PowerHA cluster
- Note that everything is being replicated so network bandwidth and quality is critical
- Any roll over or failover requires an abnormal IPL of target

Agenda

- IBM i Full System vs. IASP Replication
- **IBM i Full System Replication Solutions**
 - Metro Mirror with optional HyperSwap
 - Global Mirror
 - Multi-Target PPRC
 - SVC Split-Cluster
 - FlashCopy
- IBM i Clustering and IASP Replication with PowerHA SystemMirror for i
 - Metro Mirror & Global Mirror
 - FlashCopy
 - Multi-Target PPRC with IASP Manager
 - LUN Level Switching with SVC Split-Cluster
 - DS8000 IASP HyperSwap

IBM i Full System Synchronous Replication




DR Solution with IBM i Full System Metro Mirror Replication

- synchronous replication up to 300 km distance
- providing “zero data loss” (RPO of zero, i.e. last transaction)
- remote copy not available with active replication (backup system powered off)
- RTO determined by abnormal IPL time
- write performance degraded due to synchronous replication delays
- minimal day-to-day administration
- managed by CLI scripts or TPC-R / CSM
- support for DS8000, SVC, Storwize HyperSwap with IBM i 7.2 and later


PowerHA for i 7.2 DS8000 Full System HyperSwap

- Provides ability to ‘instantly’ switch access from production IBM i DS8000 instance to remote DS8000
- HyperSwap based on DS8000 Metro Mirror replication of IBM i production system
- Hardware availability solution for storage outages, not DR
- Automatic host-managed storage failover in case of primary DS8000 failure
 - I/O consistency ensured by IBM i multi-path driver
 - Control of DS8000 HyperSwap by IBM i in-band SCSI commands
- Switch can be manually triggered for planned maintenance: Change HyperSwap Status (CHGHYSSTS) command.



PowerHA DS8000 HyperSwap with PowerVM Partition Mobility

- Affinity can be defined between a physical system and a storage server
- Once affinity is defined, a partition mobility migration will also trigger a HyperSwap switch when appropriate
- Eliminates any I/O performance penalty of switching to a server geographically further from the storage server

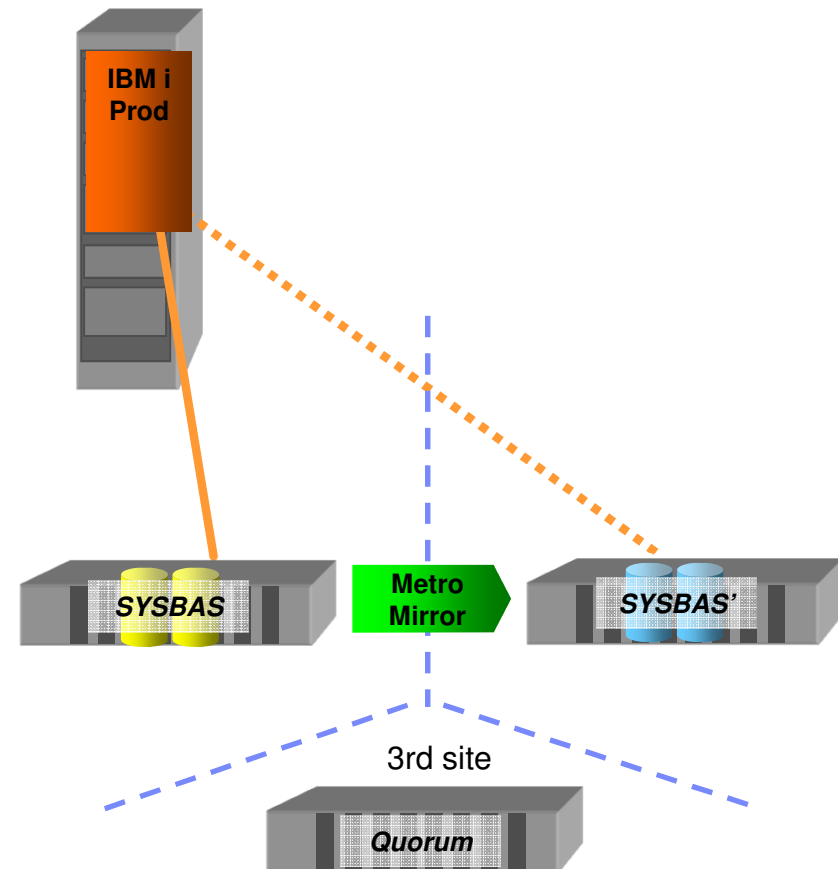


PowerHA SystemMirror for i Express Edition

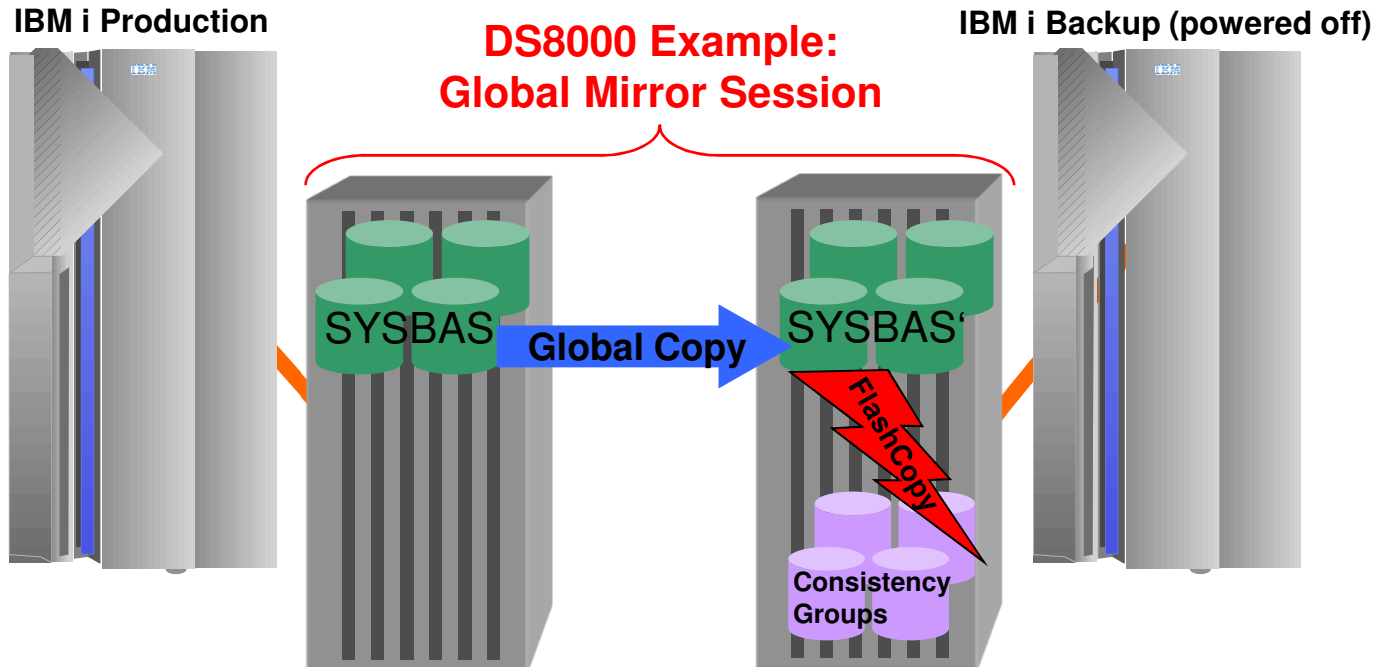
- Provides new support for DS8000 HyperSwap in full system replication environments with IBM i 7.2 and later
 - requires DS8880 / DS8870 R7.3 / DS8800 R6.3 SP7 code (bundle 86.31.97) or later
- HyperSwap by itself is a hardware availability solution
 - Single partition solution
 - Not a disaster recovery solution
 - No protection against software planned or unplanned outages
- No system downtime is associated with a HyperSwap switch
 - consistency prior to a HyperSwap reached by quiescing I/O for all volumes at the multi-path driver level
- Once configured, HyperSwap switch will occur automatically in the case of a DS8000 failure
- HyperSwap switch can be triggered manually via the Change HyperSwap Status (CHGHYSSTS) command
- Affinity between a partition and a storage server can be defined via the Add HyperSwap Storage Description (ADDHYSSTGD) command
 - Storage affinity used to designate local storage for partition mobility cases

IBM i 7.2 TR3 SVC/Storwize Full System HyperSwap

- Provides ability to ‘instantly’ switch access from production IBM i SVC/Storwize instance to remote SVC/Storwize
- HyperSwap based on Metro Mirror replication of IBM i production system
 - SVC/Storwize intra-cluster HyperSwap support only
 - same UID used for master and auxiliary volumes and available for read/write access via two I/O groups
 - local SVC/Storwize with preferred paths designated by host site attribute
- Automatic storage-managed failover in case of primary storage outage or if majority of volume pair I/O is received on secondary volume
- No user initiated manual HyperSwap switch-over



IBM i Full System Asynchronous Replication


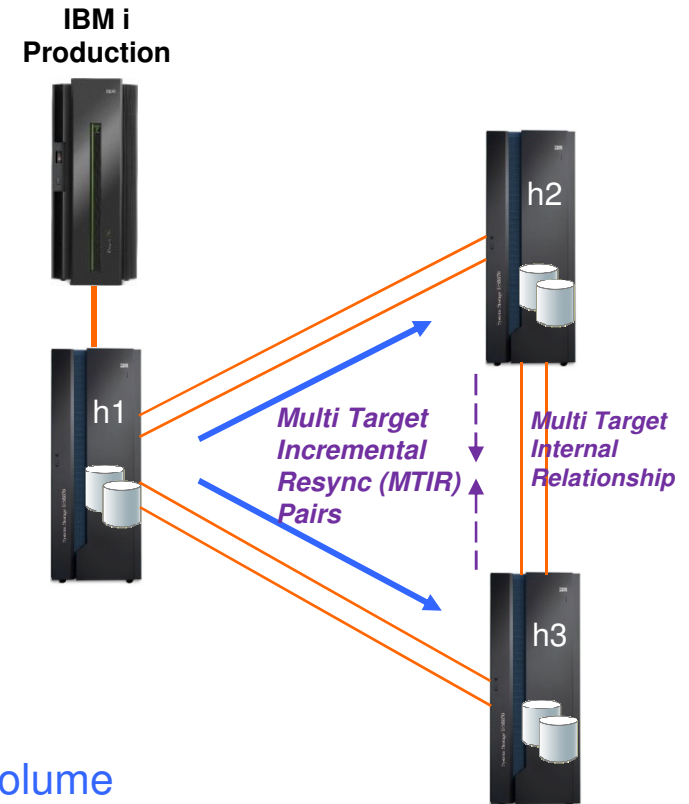


DR Solution with IBM i Full System Global Mirror Replication

- asynchronous replication providing “unlimited” distance
- Global Mirror consistency groups ensuring data consistency
- RPO as low as near zero
(determined by replication bandwidth or if using SVC/Storwize change volumes by cycling period)
- usually no host performance degradation due to asynchronous replication
- remote copy not available with active replication (backup system powered off)
- RTO determined by abnormal IPL time
- minimal day-to-day administration
- managed by CLI scripts or TPC-R / CSM

IBM i Full System DS8000 Multi-Target PPRC

- Multi-target PPRC provides mirroring protection even after loss of one site
 - requires DS8870 R7.4 w/ feature #7025, or later
 - all involved DS8000 must meet this requirement
- Set up by establishing first PPRC pair, then second pair
 - internal PPRC pairs enabling change tracking for incremental resync between secondary volumes are automatically created
 - if second pair gets removed, internal pairs are also removed
- MT-PPRC sessions supported by CSM 6.1 or later:



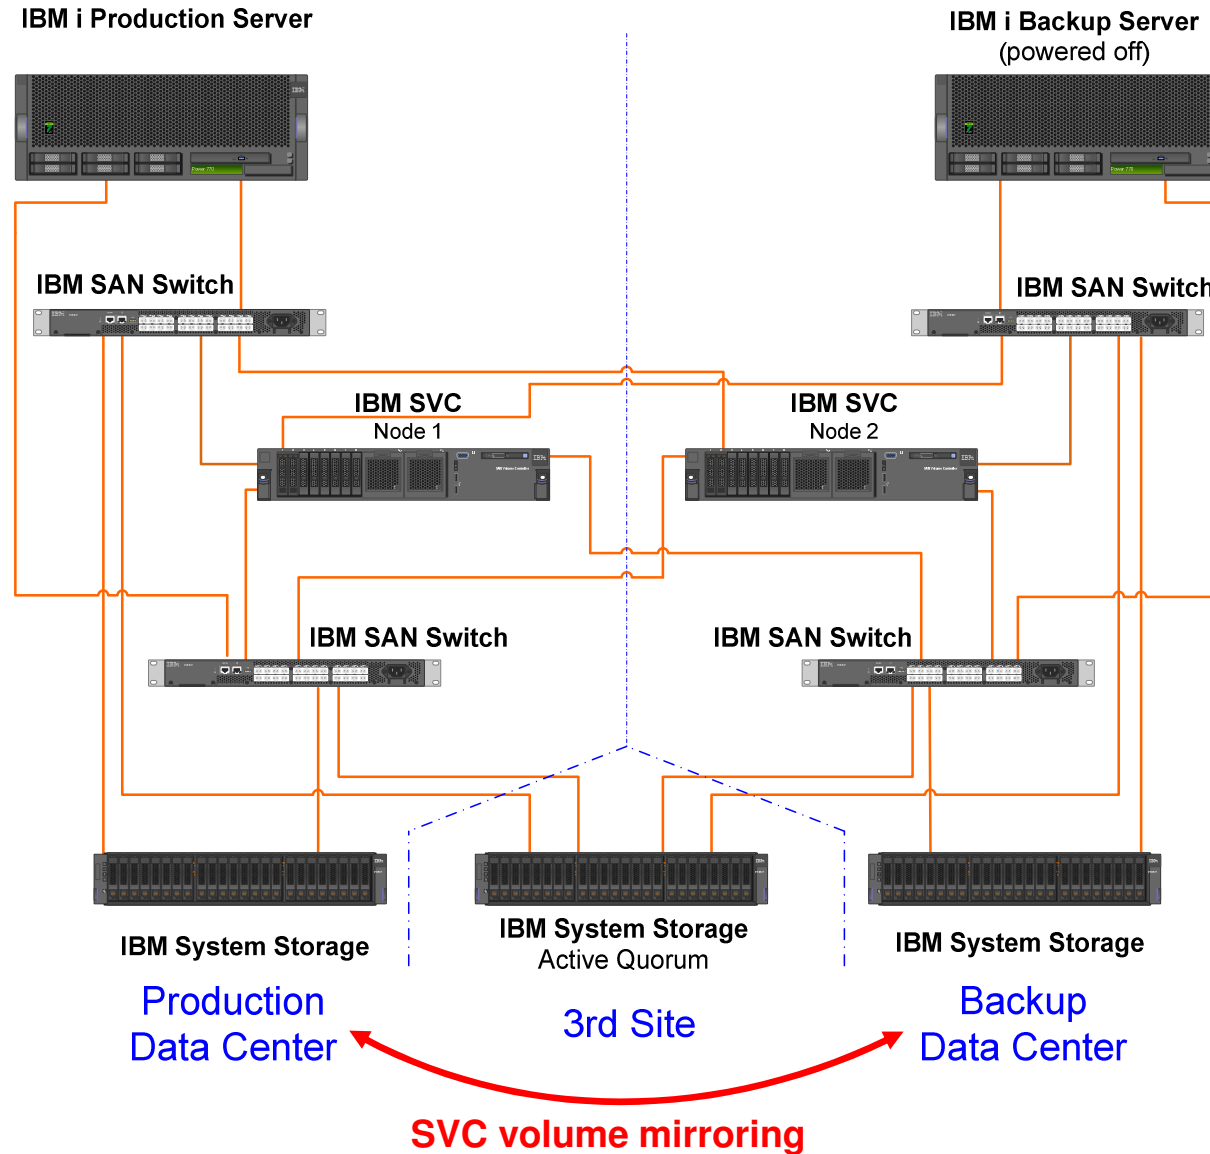
CSM = Copy Services Manager
 RedBook SG24-8367 DS8000 Copy Services

- alternatively asynchronous Global Copy is supported in any leg
- CSM 6.1.3 added support for MM-GM with a return to sites H1 or H2 in GM mode

Multi Target positioning

- Many companies and businesses require their applications to be continuously available and cannot tolerate any service interruption. A loss of a disaster recovery site is often considered to be a severe impact on their business.
- If their local production site fails, swapping to a Metro Mirror target allows applications to continue running. However, without another target to act as a backup for disaster protection, business applications are left unprotected from a subsequent failure.
- MultipleTarget PPRC addresses these needs by allowing an additional replication target.
- Multiple Target PPRC also complements the use of cascaded 3-site mirroring topologies, such as Metro/Global Mirror, by adding flexibility and simplifying certain recovery scenarios

IBM i Full System Replication with SVC Split-Cluster



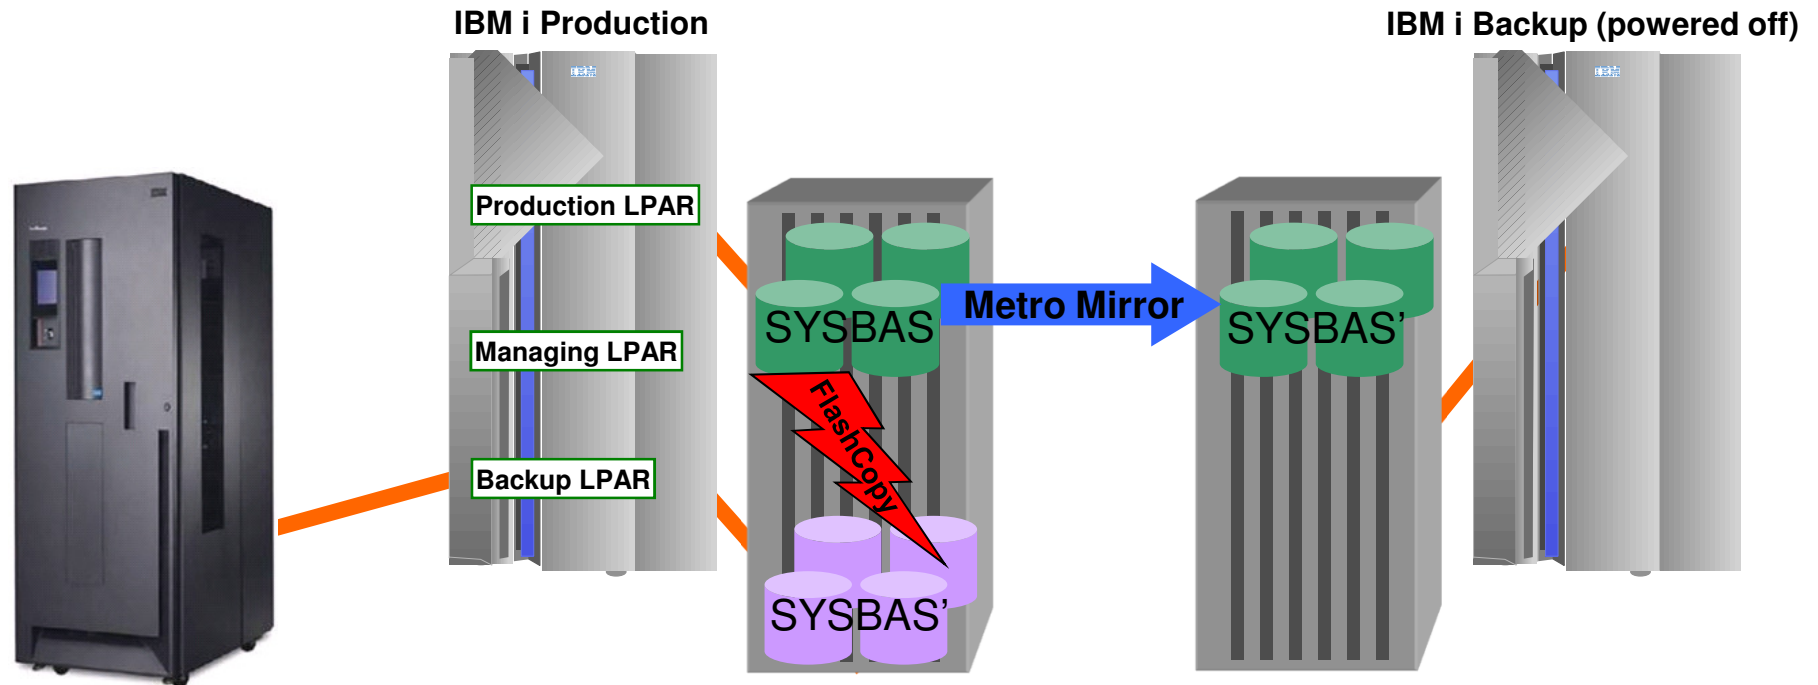
SVC Split-Cluster Considerations

- Storage HA solution with IBM SVC split-cluster and SVC volume mirroring
 - synchronous volume mirroring for up to 40 km distance without ISLs (up to 300 km with ISLs)
 - transparent storage “failover” provided by SVC volume mirroring
 - simultaneous host access to both copies must be prevented (backup system powered off)
 - RTO determined by abnormal IPL time
 - minimal day-to-day administration
 - HA testing and configuration changes more challenging than with remote copy
 - e.g. manual assignment needed for preferred node to enable local reads
 - *Enhanced Stretched Cluster* introduced with SVC V7.2 adds site awareness feature (reads always locally) and DR capability in case of simultaneous site and active quorum failure

Remote Copy Bandwidth Considerations

- Synchronous replication bandwidth requirement dependent on host peak write I/O throughput
- Asynchronous replication bandwidth requirement:
 - for DS8000 Global Mirror:
 - required bandwidth depending on tolerable RPO
 - RPO self-adapting to available bandwidth as low as 2-3 sec.
 - RPO Estimator spreadsheet available for sizing
 - for SVC/Storwize Global Mirror:
 - same bandwidth requirements like for synchronous replication
 - RPO near zero with sufficient bandwidth
 - cache fill-up and host I/O group slow-down if bandwidth remains too low
 - for SVC/Storwize with Global Mirror change volumes:
 - used primarily for cases of high RPO tolerance / low bandwidth
 - RPO typically 2x cycling period
 - FlashCopy change volumes using grain size of 256 KB
 - bandwidth requirement depending on host write locality and usually hard to predict unless prior tested with monitoring FlashCopy change rates

IBM i Full System FlashCopy

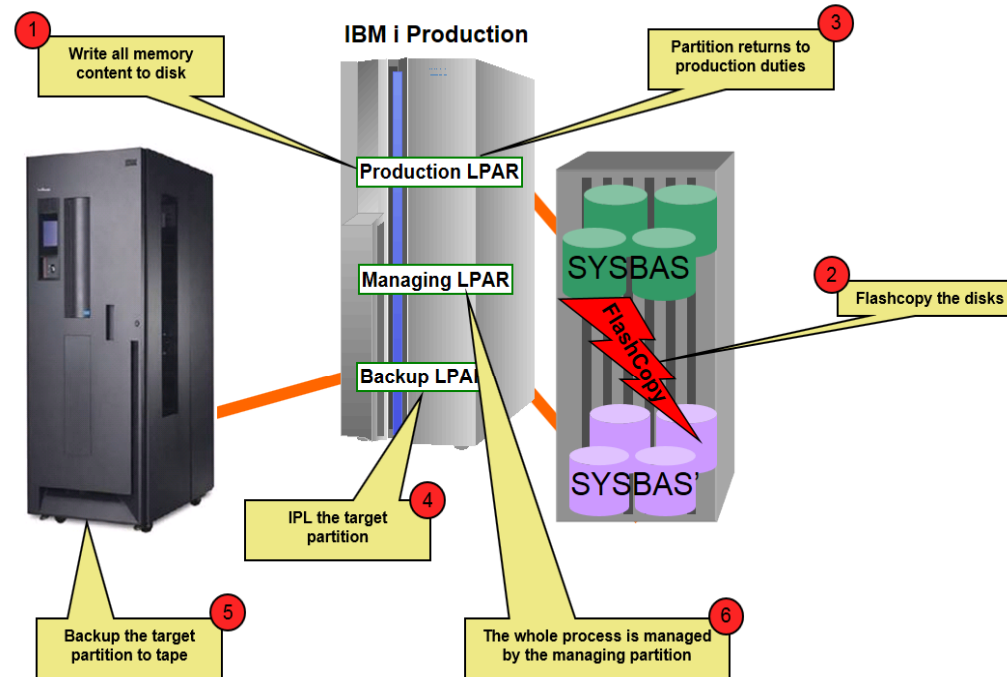


Non/low disruptive Backup Solution with IBM i Full System FlashCopy

- instant IBM i production system cloning with FlashCopy for online backup, testing, dev.
- only minimal disruption (~1 min.) for system quiesce with IBM i 6.1 or later
- FlashCopy integration with BRMS Network
- FlashCopy automation supported either from Metro Mirror primary or secondary volumes
- full automation by *Full System Copy Services Manager* IBM Lab Services offering

IBM i Full System FlashCopy Process

IBM i Full System FlashCopy Process



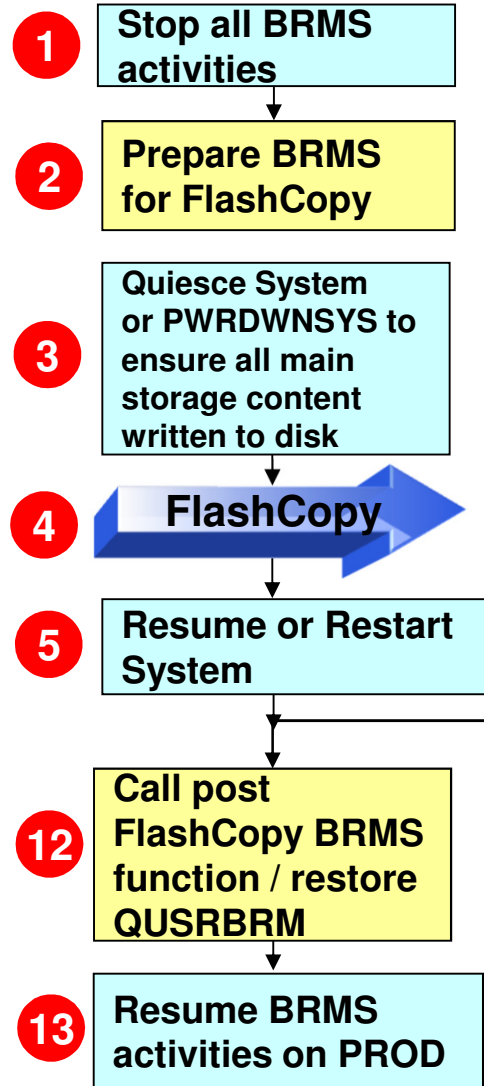
Managing LPAR with IP connection to production/backup LPARs, HMC and storage CLI is key to manage the full system FlashCopy process as FlashCopy cannot be invoked from a quiesced or powered off production LPAR

BRMS Network FlashCopy Integration

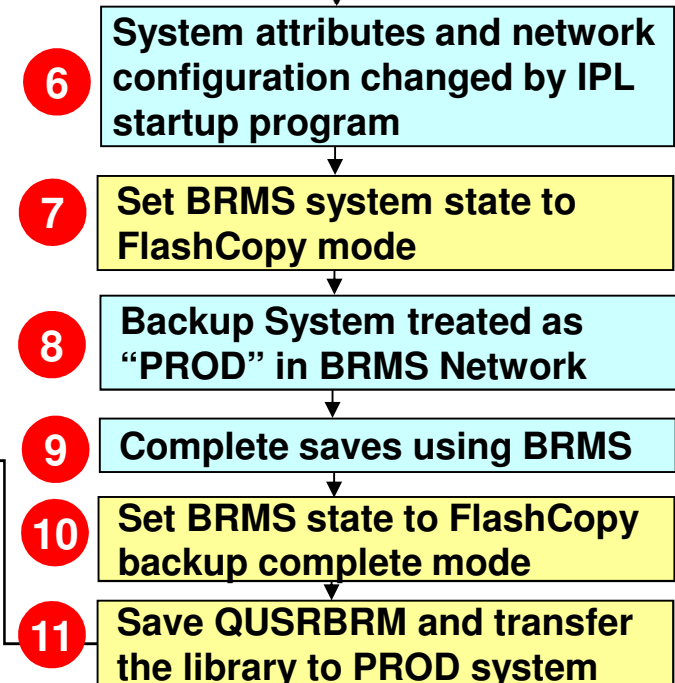
Production System Name: PROD
BRMS System Name: PROD



- Allows to perform a backup on a system that has been copied by FlashCopy with a BRMS history appearing as if the backup was performed on the production system
- *Full System Copy Services Manager* automates these steps
 - DS8000 storage support
 - SVC/Storwize storage support (with FSFC Toolkit version 7.4 or later)



Backup System Name: PROD
BRMS System Name: PROD



For more information refer to *IBM System Storage Copy Services and IBM i: A Guide to Planning and Implementing*, SG24-7103

IBM i Full System Copy Services Manager

Allows you to perform full system backups with minimal downtime

Features

- Automates managing IBM i partitions, performing a FlashCopy and starting your backups.
- Communicates with the DS8000 or SVC/Storwize to automate FlashCopy and changing host connections.
- Communicates with the HMC to IPL partitions.
- Source partition IPL is not required for FlashCopy at IBM i 6.1 and higher.
- All releases support 'warm FlashCopy' for testing purposes.
- Integrated with BRMS to maintain the integrity of the database, including transferring QUSRBRM back to the source partition.
- Supports Space Efficient FlashCopy.
- Flexible architecture allows customization work with many environments and scenarios.
- Administrative and Operator tasks separated for ease of use and security.
- Can perform sequential operations on shared target disks or parallel operations on dedicated disks.
- Manages changing host adapter connections on the storage to allow one partition to attach to different sets of disks
- Software, support, and services are provided through a services contract.

Typical Benefits

- Perform full system backups without a system outage (IBM i 6.1 and higher)
- Perform full system backups in the time it takes to IPL (all releases).
- Replace 10-20 manual steps on multiple systems with a single command issued from an IBM i partition.
- Create copies of partitions for any purposes.
- Minimize potential user errors that cause system outages.
- Reduce personnel training requirements.

Why IBM?

- **Deep skills in IBM i implementation and integration.**
- **Experience in application architecture and design gained from thousands of engagements across many industries.**
- **Ability to deliver skills transfer as part of service engagement.**

www.ibm.com/systems/services/labservices

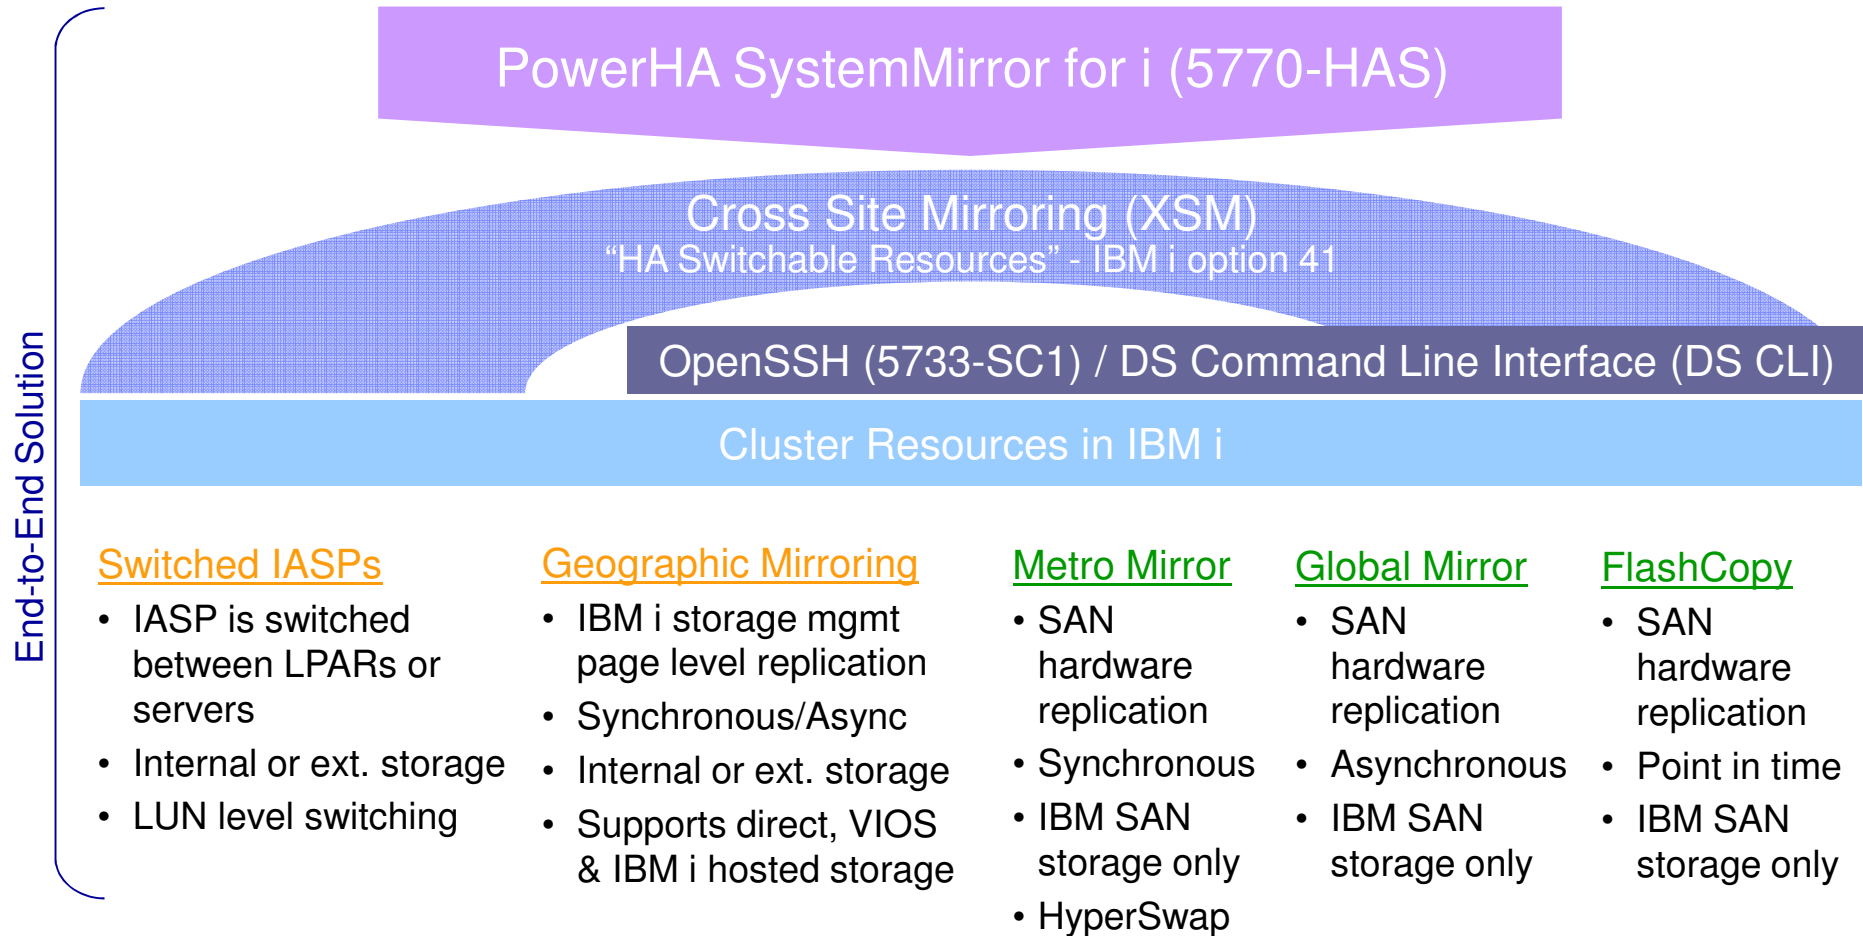
stgls@us.ibm.com

Agenda

- IBM i Full System vs. IASP Replication
- IBM i Full System Replication Solutions
 - Metro Mirror with optional HyperSwap
 - Global Mirror
 - Multi-Target PPRC
 - SVC Split-Cluster
 - FlashCopy
- IBM i Clustering and IASP Replication with PowerHA SystemMirror for i
 - Metro Mirror & Global Mirror
 - FlashCopy
 - Multi-Target PPRC with IASP Manager
 - LUN Level Switching with SVC Split-Cluster
 - DS8000 IASP HyperSwap

IBM PowerHA SystemMirror for i

Provides an end-to-end solution for management of IBM i IASP and DS8000, SVC / Storwize family resiliency and replication technologies for HA and DR



PowerHA System Mirror for i Packaging

■ IBM PowerHA SystemMirror is offered in three editions:

- *Express Edition* (5770-HAS *BASE & option 3) for full-system DS8000 HyperSwap
- *Standard Edition* (5770-HAS *BASE & option 2) for local data center replication
- *Enterprise Edition* (5770-HAS *BASE & option 1) supporting also multi-site replication

PowerHA SystemMirror for i	Express Edition	Standard Edition	Enterprise Edition
Centralized cluster management		✓	✓
Cluster resource management		✓	✓
Centralized cluster configuration		✓	✓
Automated cluster validation		✓	✓
Cluster admin domain		✓	✓
Cluster device domain		✓	✓
Integrated heartbeat		✓	✓
Application monitoring		✓	✓
IBM i event / error management		✓	✓
Automated planned switch over		✓	✓
Managed unplanned fail over		✓	✓
Centralized Flash Copy		✓	✓
LUN level switching		✓	✓
GeoMirror sync delivery		✓	✓
GeoMirror async delivery			✓
Multi-Site HA/DR management			✓
SVC / Storwize / DS8000 Metro Mirror			✓
SVC / Storwize / DS8000 Global Mirror			✓
Full-System DS8000 HyperSwap	✓	✓	✓
IASP DS8000 HyperSwap			✓



PowerHA SystemMirror for i Storage Support

	Internal SAS/SSD	DS8000	SVC V840 V9000	V7000 V5000 V3700	Flash System 840/900	XIV	Other Storage
Geographic mirroring	✓	✓	✓	✓	✓	✓	✓
Metro Mirror		✓	✓ ¹	✓ ¹			
Global Mirror		✓	✓ ¹	✓ ¹			
Metro Global Mirror (MGM)		✓ ²					
Multi-Target PPRC		✓ ³					
LUN level switching		✓ ¹	✓ ¹	✓ ¹			
SVC split-cluster			✓ ⁶				
FlashCopy		✓ ⁴	✓ ¹	✓ ¹			
HyperSwap		✓ ⁵	✓ ⁷	✓ ⁷			

¹ requires IBM i 7.1 or later

² requires IASP Copy Services Manager 3.1 for PowerHA and TPC for Replication or CSM

³ requires IBM i 7.1, 3x DS8870 R7.4, IASP Manager 4.1 for PowerHA and CSM 6.1.4 or later

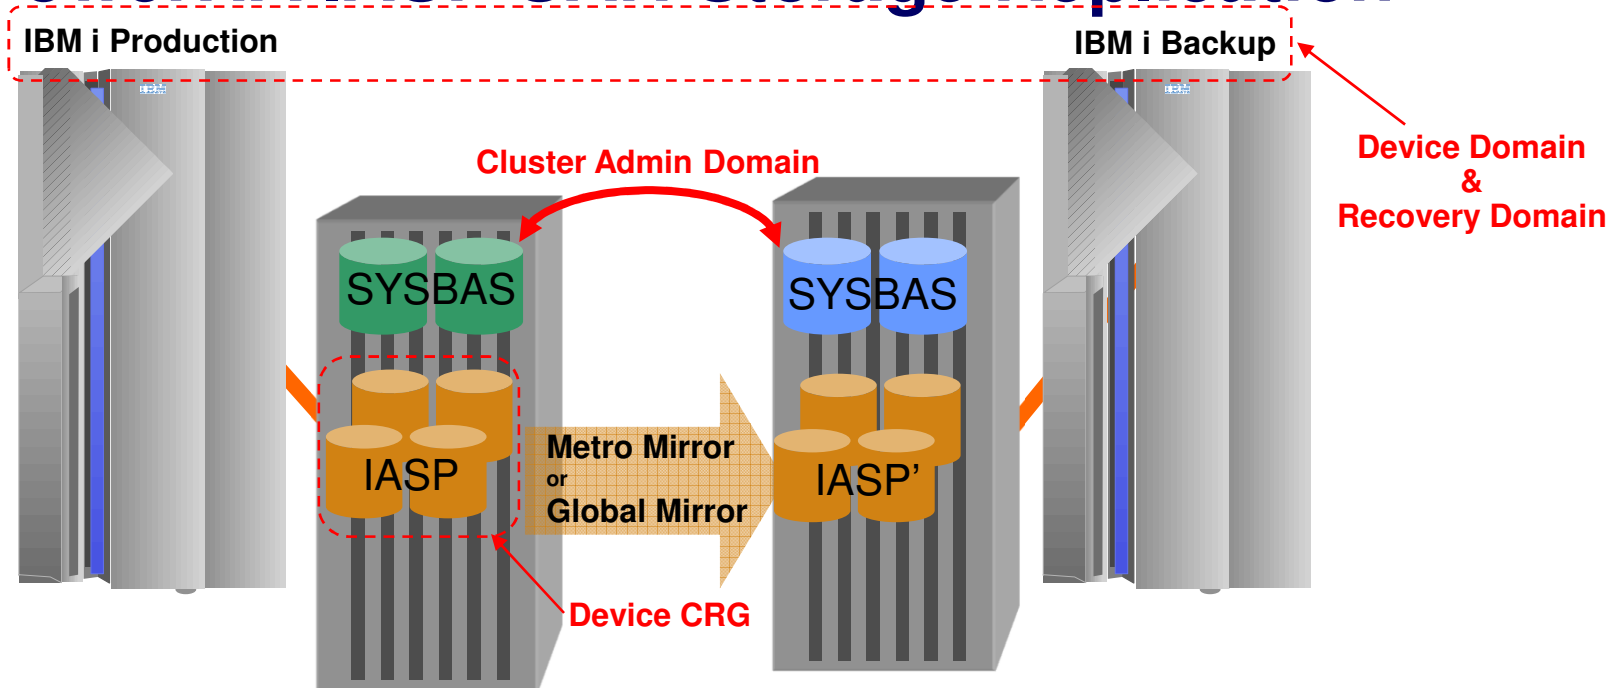
⁴ DS8000 FlashCopy Space-Efficient requires IBM i 7.1 or later

⁵ requires DS8870/DS8800 or later and IBM i 7.2 or later, IASP HyperSwap requires IBM i 7.2 TR4 or later

⁶ SVC split-cluster not supported with V840 / V9000

⁷ IASP HyperSwap requires IBM i 7.2 TR5 / i 7.3 TR1 or later, no V3700 support

IBM i PowerHA IASP SAN Storage Replication

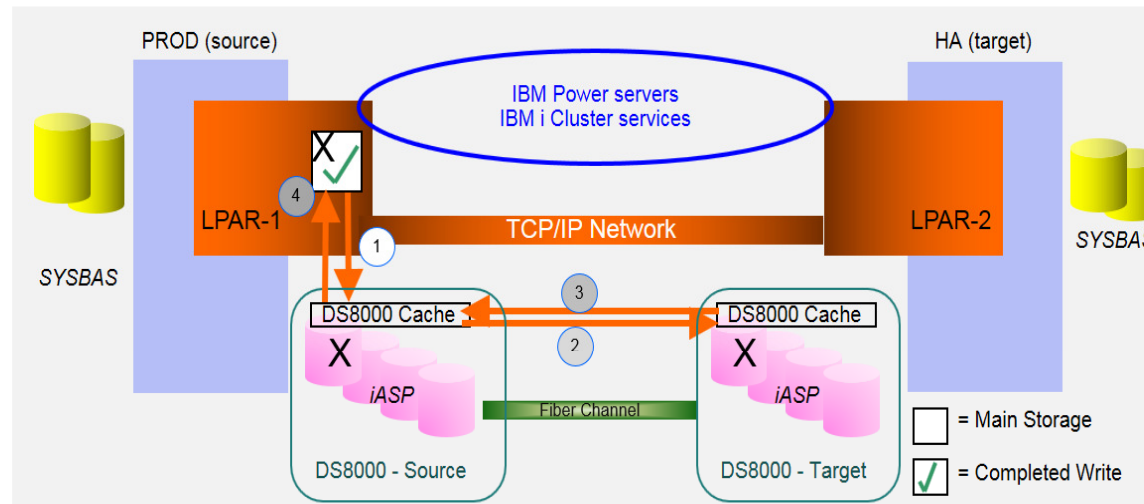


HA/DR Solution with IBM i IASP SAN Copy Services Replication

- synchronous replication of IASP only (up to 300 km)
 - providing “zero data loss” (RPO = 0), write performance degraded by replication delays
- asynchronous replication of IASP only for “global” distances
 - still consistent data but RPO > 0 depending on bandwidth, no write performance degradation
- lower bandwidth requirements and better performance due to non-replicated *SYSBAS temporary writes
- remote copy of IASP not connected with active replication
- RTO determined by abnormal IASP vary-on time
- environment resilience enabled by using *cluster administrative domain*
- minimal day-to-day administration
- managed with IBM PowerHA System Mirror for i or *IASP Manager* for PowerHA IBM Lab Services offering

PowerHA Metro Mirror

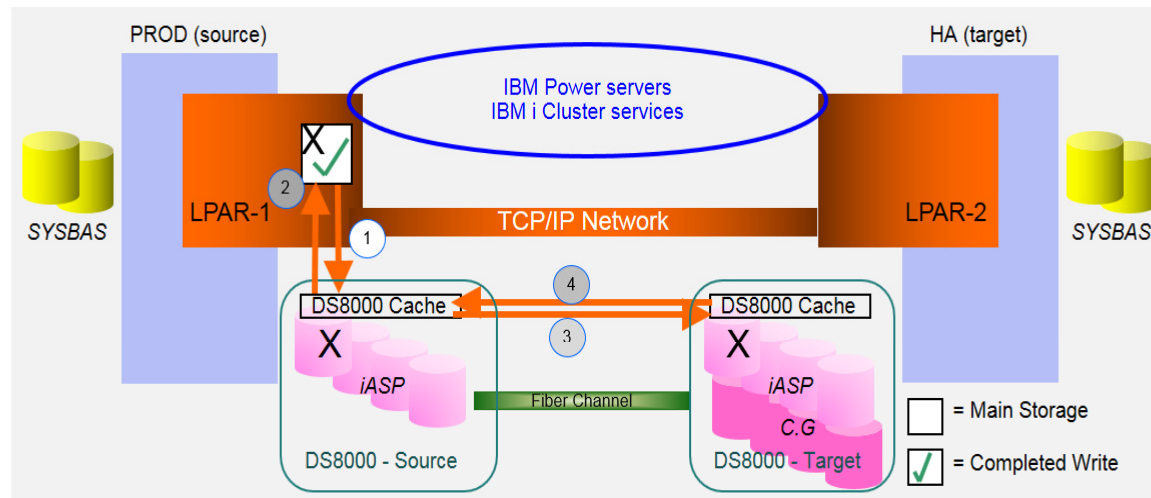
Synchronous for Local or short distances HA



- Metro Mirror provides external storage data replication between remote DS8000 through Storage Area Network (SAN) either using Fiber Channel connection or FC over Ethernet connection
- Synchronous write to target copy means the change must be completed on target DS8000 non-volatile storage before it can complete on source
 - the source and target copy of the iASP are in consistent state when unexpected failure occurs on any of the cluster nodes
- Very good RPO – changes are confirmed on target copy before source
- To limit application performance impact, the distance between the replication sites is limited to short distances, which is about 30 kilometers

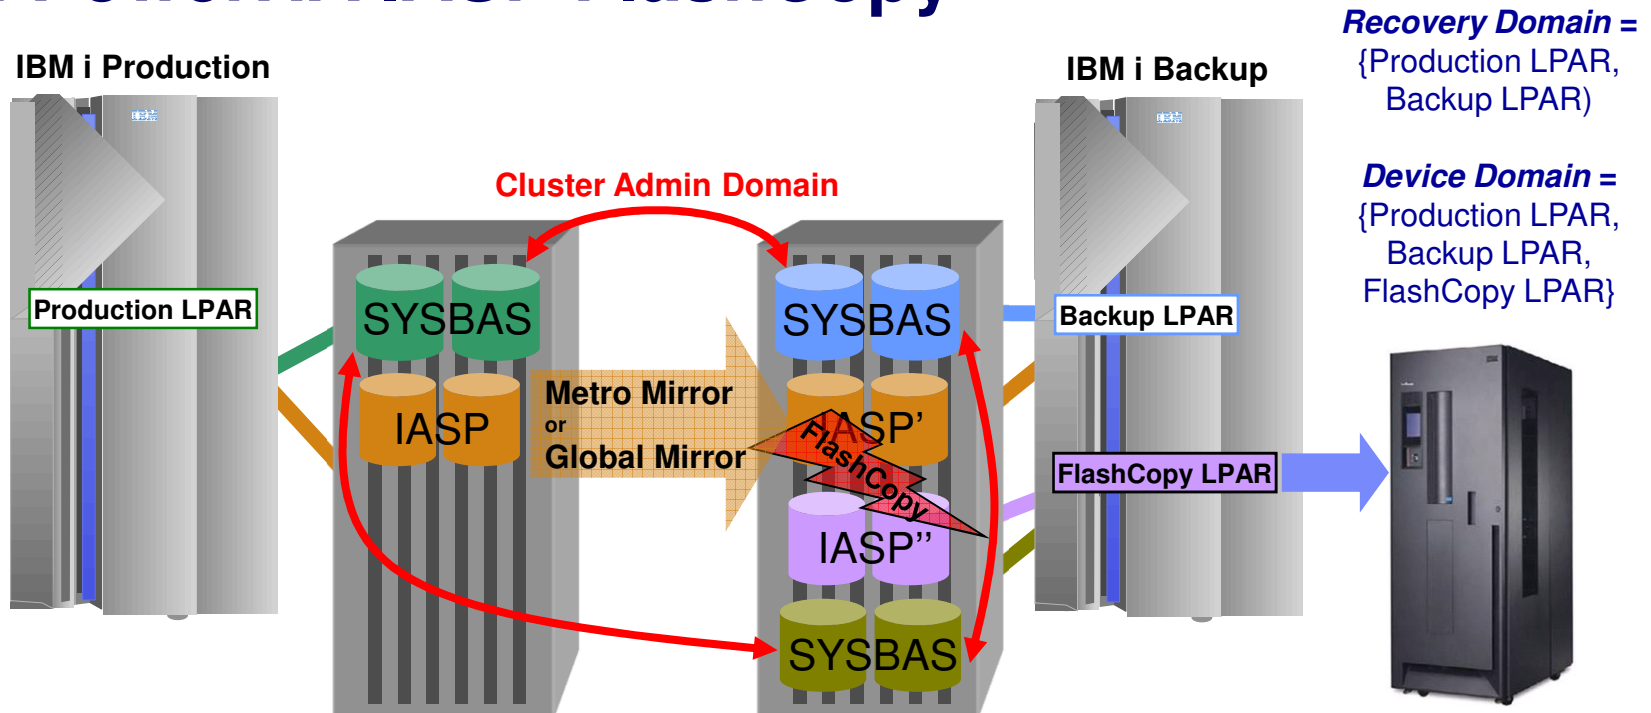
PowerHA Global Mirror

Asynchronous for Long distances HA or DR



- Global Mirror provides external storage data replication between remote DS8000 through Storage Area Network (SAN) either using Fiber Channel connection or FC over Ethernet connection
- Asynchronous write to target copy means the change is written on source DS8000 non-volatile storage to be complete on source partition
 - the source and target copy of the iASP may not be in consistent state when unexpected failure occurs on any of the cluster nodes
- Global Mirror combines two DS8000 techniques, which are Global Copy and FlashCopy.
 - A 3rd set of volumes is required on target storage for Consistency Group used for FlashCopy
- RPO not null – changes are confirmed on source copy then replicated on target
- No application performance impact, the distance between the replication sites is not limited

IBM i PowerHA IASP FlashCopy



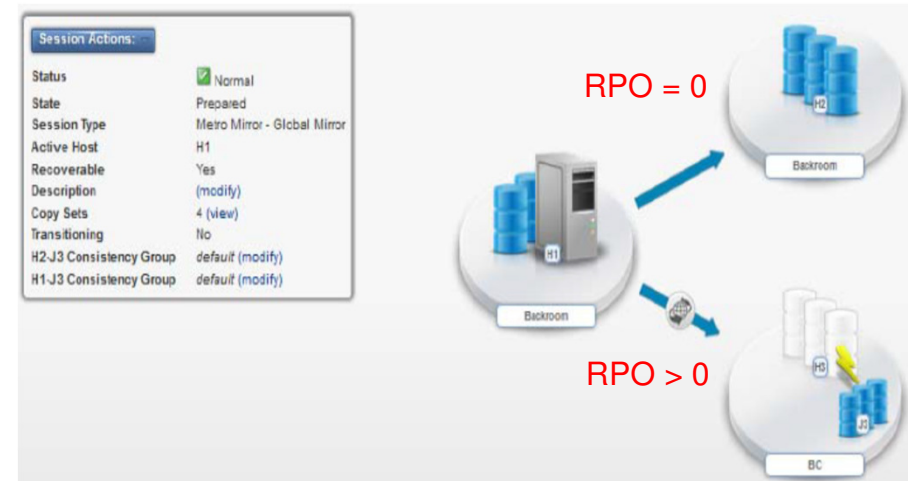
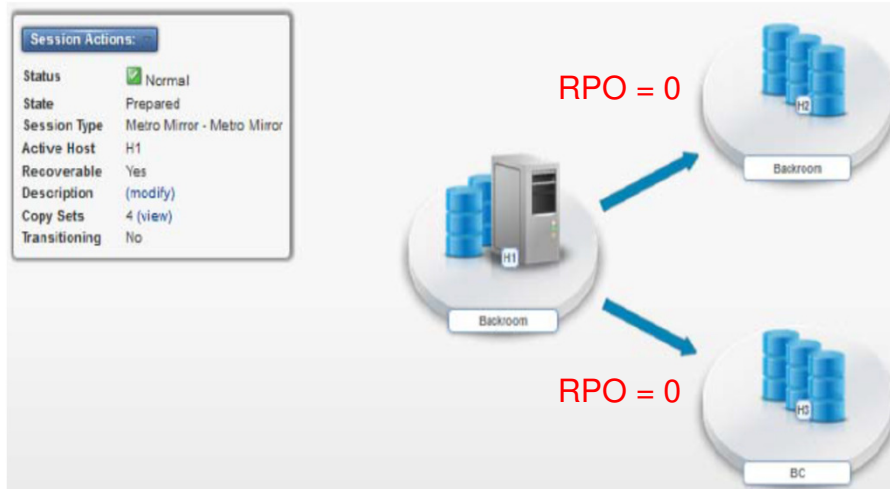
Non/low disruptive Backup Solution with IBM i IASP FlashCopy

- FlashCopy creates a point-in-time copy of an IASP for online backup, testing, development etc.
- IASP activity is quiesced via CHGASPACT CL command before starting FlashCopy
- with Metro Mirror: FlashCopy of the IASP can be implemented either on production or backup site
- DS8000 Global Mirror: PowerHA support for FlashCopy on the backup site requires IBM i 7.1 TR6 or later
- support for DS8000 FlashCopy no-copy reverse requires IBM i 7.1 TR6 or later
- managed with PowerHA or *IASP Manager* for PowerHA IBM Lab Services offering

Systems Lab Services – IASP Manager 4.1 for MT-PPRC

Multi-Target: Metro Mirror / Metro Mirror

Multi-Target: Metro Mirror / Global Mirror




MT-PPRC 3-Site Solution IBM i IASP Replication Support with IASP Manager 4.1

- provides 2x synchronous (1x sync. + 1x async.) replicated IASPs pairs with MT-PPRC MM-MM (MM-GM)
- requires 3x DS8870 R7.4 or later and CSM 6.1.4 or later
- requires IBM i 7.1 or later (due to CSM client Java 8 requirement)
- managed with PowerHA and *IASP Manager 4.1* IBM Systems Lab Services offering

* Note: Cascaded Metro-/Global Mirror (MGM) support has been discontinued with IASP Manager 4.1 in favour of MT-PPRC

PowerHA LUN Level Switching

- Local HA solution providing server level redundancy for protection from server outage but not disk outage
 - IBM i 7.1 or later required
 - Requires NPIV or native attachment of DS8000 or SVC/Storwize storage device
 - Single-copy, i.e. non-replicated, IASP in a device domain can be switched between two systems in a cluster (can be combined though with Metro or Global Mirror)
 - IASP accessible only from single system at a time
 - PowerHA remaps host connections within the storage to switch



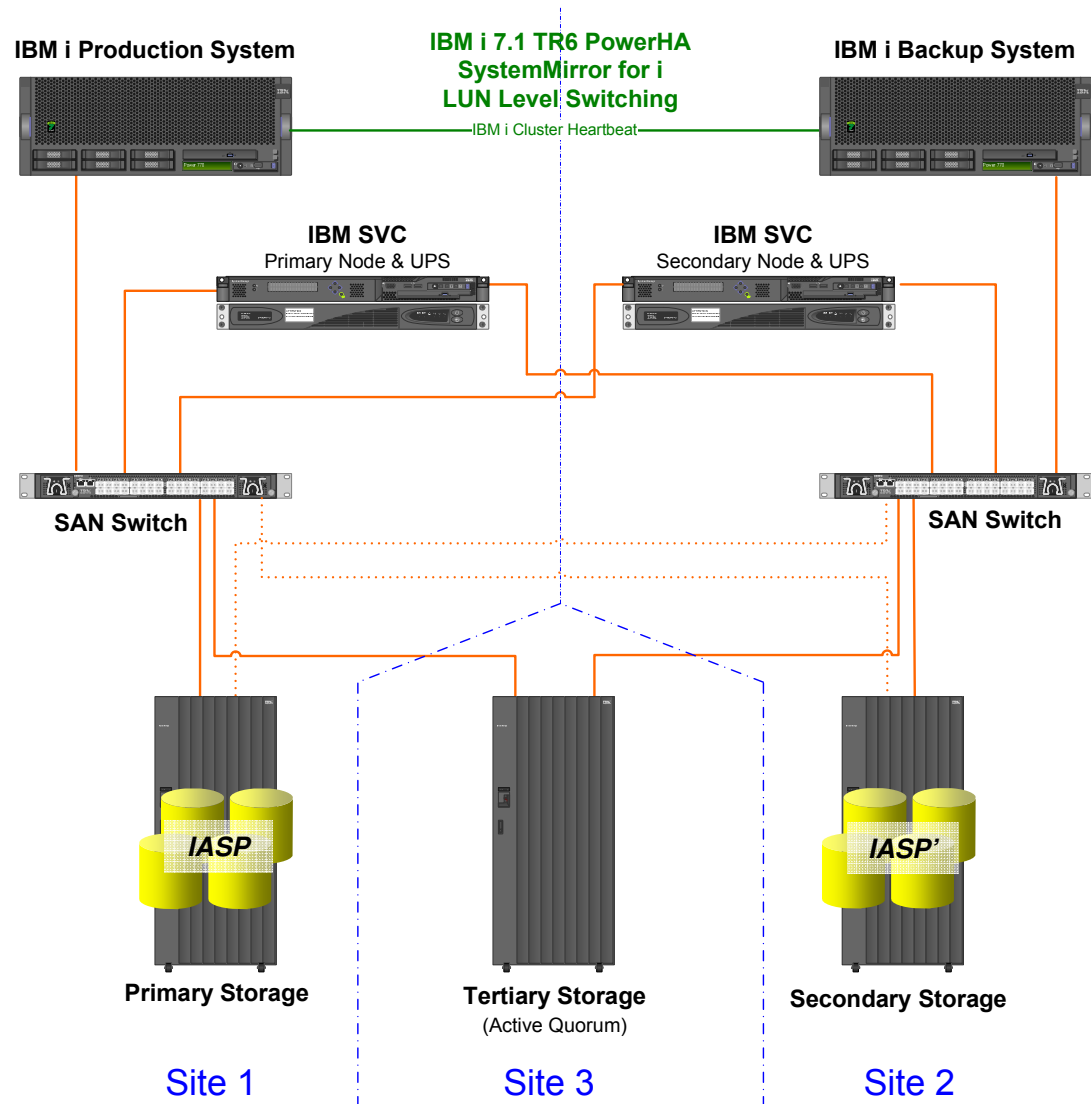
PowerHA LUN Level Switching for SVC Split-Cluster (1/2)

- Supported two-site HA solution using a combination of IBM i PowerHA LUN level switching for server redundancy with SVC split-cluster volume mirroring for storage redundancy
 - Available via PowerHA group PTF SF99706 level 5
 - IBM i 7.1 TR6 or later required
 - Requires NPIV or native attachment of SAN Volume Controller
 - Enhanced protection against rolling disasters with SVC V7.2+ Enhanced Split-Cluster requires RPQ support for PowerHA

- Provides the following benefits over alternative IBM i full-system replication SVC support
 - High degree of automation for planned and unplanned site-switches respectively failovers
 - Shorter recovery times by using IASP technology
 - Reduced mirroring bandwidth requirements by using IASP technology (temporary writes in SYSBAS like for index builds are not mirrored)

PowerHA LUN Level Switching for SVC Split-Cluster (2/2)

- Schematic view of IBM PowerHA System Mirror for i LUN level switching with SVC split-cluster and SVC volume mirroring of IASP volumes between site 1 & site 2



Legend: Dotted lines denote FC paths required only for SVC non-split cluster configuration

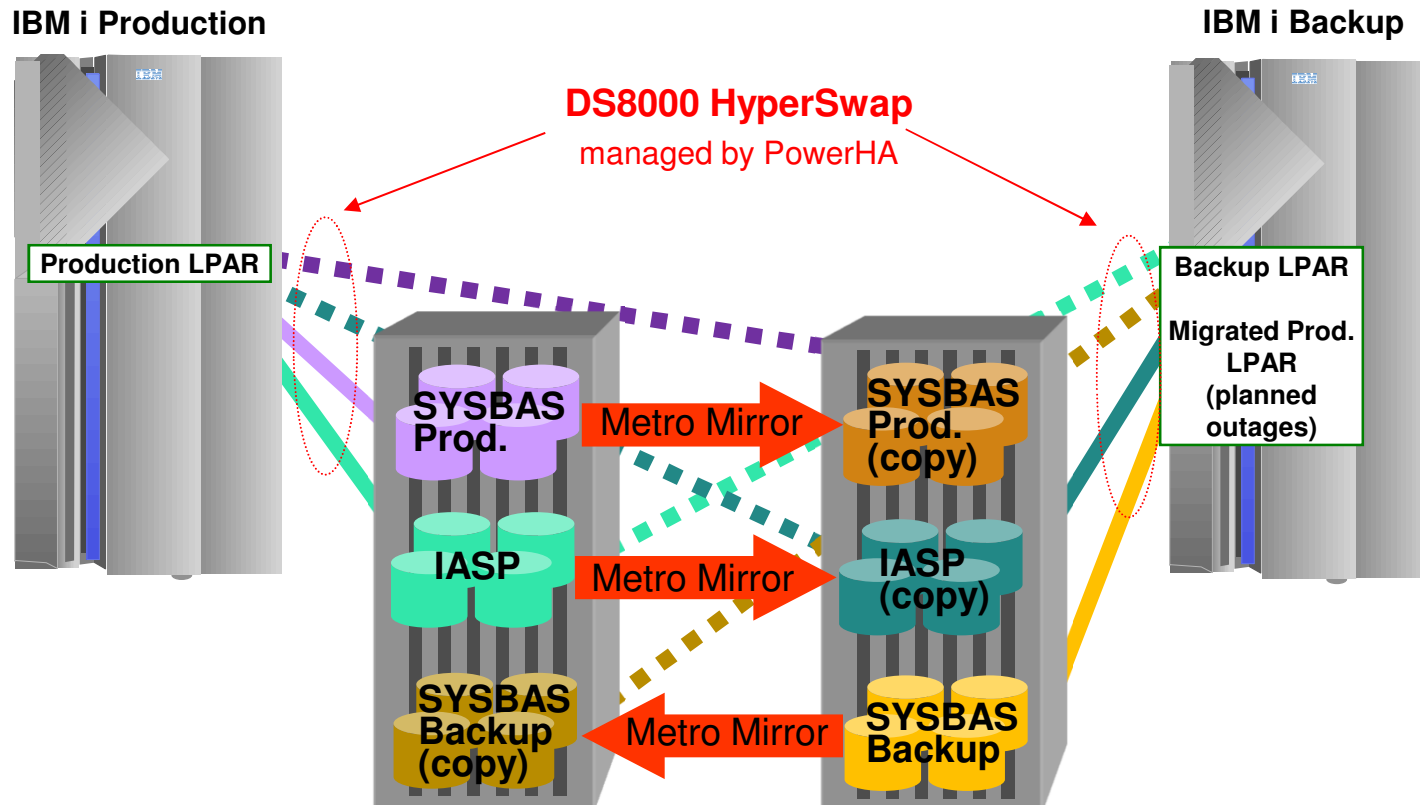
SVC Split-Cluster vs. HyperSwap for IBM i

Criteria	SVC Split-Cluster	SVC Enhanced Stretched-Cluster	SVC HyperSwap	DS8000 HyperSwap
Consistent data in case of rolling disaster	no	yes	yes	yes
Preserve consistency at both sites during resynchronization	no	no	yes	no
Consistency group support	no	no	yes	yes ¹
Write caching after site failure	no	no	yes	yes
Scalability (max. # of volumes)	4096 5000 (V7.7.1+)	4096 5000 (V7.7.1+)	1024 1250 (V7.7.1+)	65280
Min. # of storage HW	2x SVC nodes	2x SVC nodes (4x recommended)	4x SVC nodes	2x DS88xx
Remote Mirroring license required	no	no	yes	yes
Quorum device required	yes	yes	yes	no
IBM i IASP replication support	yes	RPQ	yes	yes
PowerVM Live Partition Mobility Support	yes	yes	yes	yes

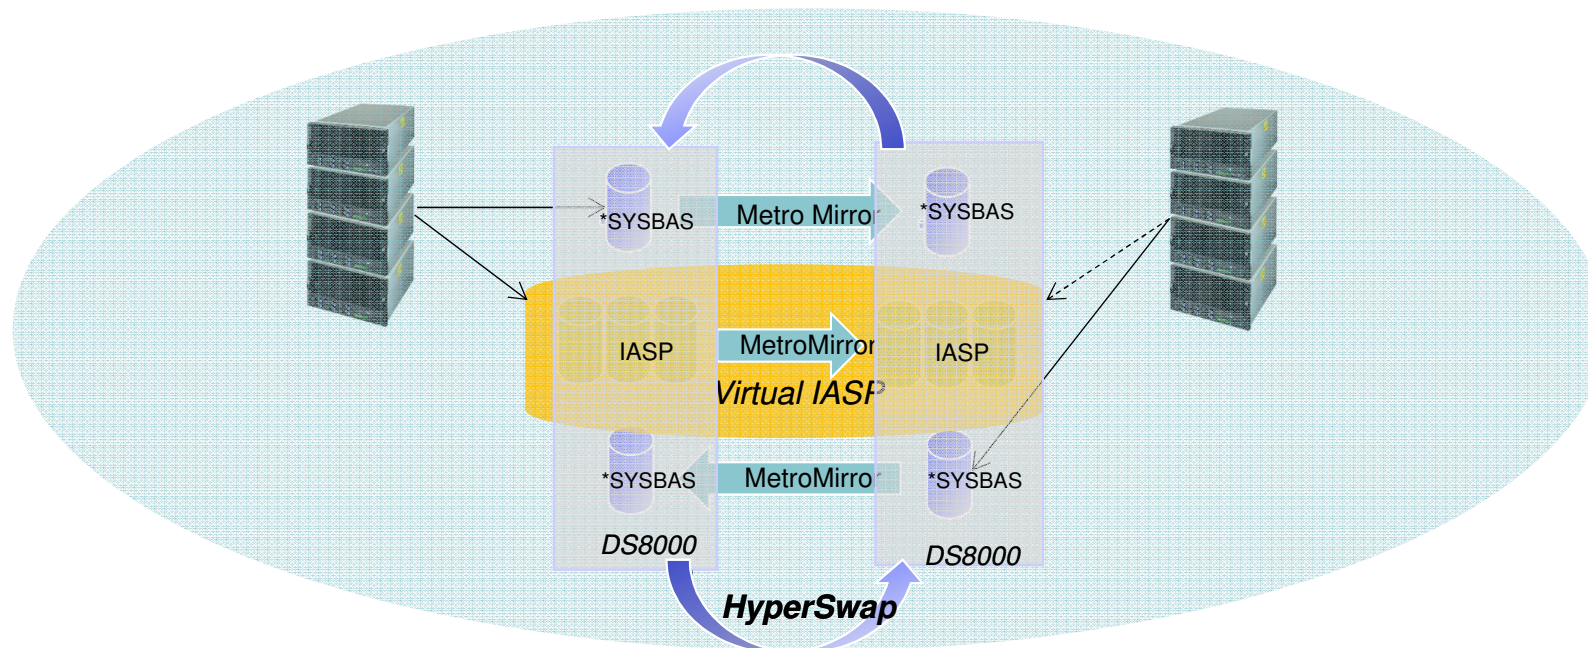
¹ IBM i system level consistency enabled via IBM i multi-path driver

IBM i DS8000 IASP HyperSwap with PowerHA

- Combination of HyperSwap with PowerHA LUN level switching technology provides near zero downtime for planned storage outages, and minimal downtime for unplanned storage and planned or unplanned server outages
- supported with IBM i 7.2 TR4 / i 7.3 with PowerHA Enterprise Edition



Power HA Enterprise edition HyperSwap cluster.

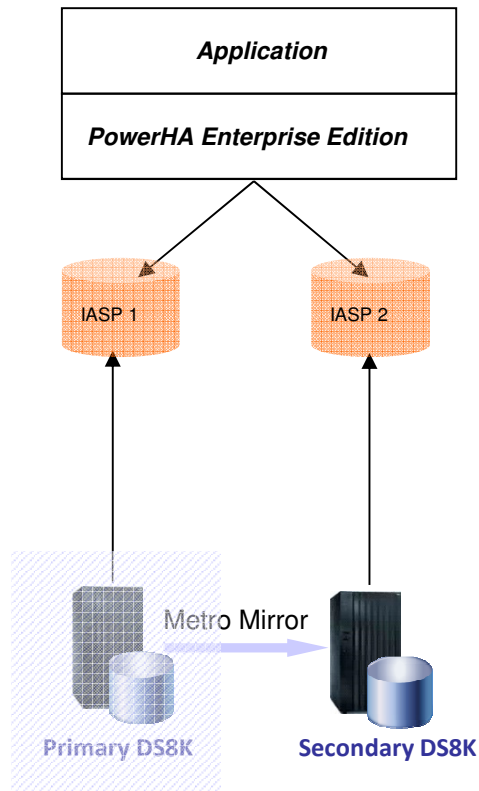


- HyperSwap has the effect of making the replicated pair of IASPs appear as a single virtual IASP
 - HyperSwap switches the source and target IASP and SYSBAS in the event of a storage outage
 - The source IASP is mirrored to the target IASP via Metro Mirror
 - The source and target SYSBAS data are mirrored via Metro Mirror
- In the event a storage outage event, the source system switches to the mirrored IASP and the mirrored SYSBAS
- In the event of a production server outage, PowerHA conducts a failover to the target production server, (the virtual IASP is switched to target) Metro Mirror will reverse direction of replication and production resumes on the secondary power server
- If VIOS is deployed, LPM can be used for firmware updates, load balancing etc
- DS8800 and above (TPC-R not utilized)

PowerHA HyperSwap clustering technology (Principe)

Applications halts while cluster switches

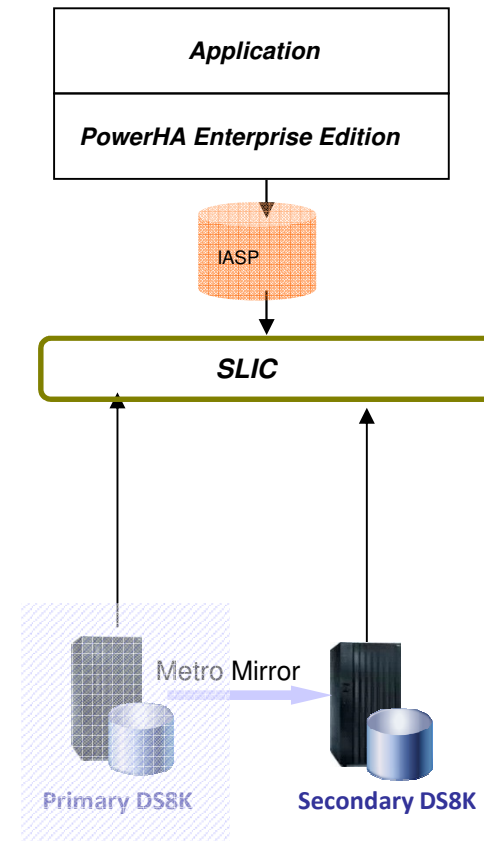
- Storage resumes after cluster switch



Traditional Metro Mirror Cluster

Applications continue to use the IASP

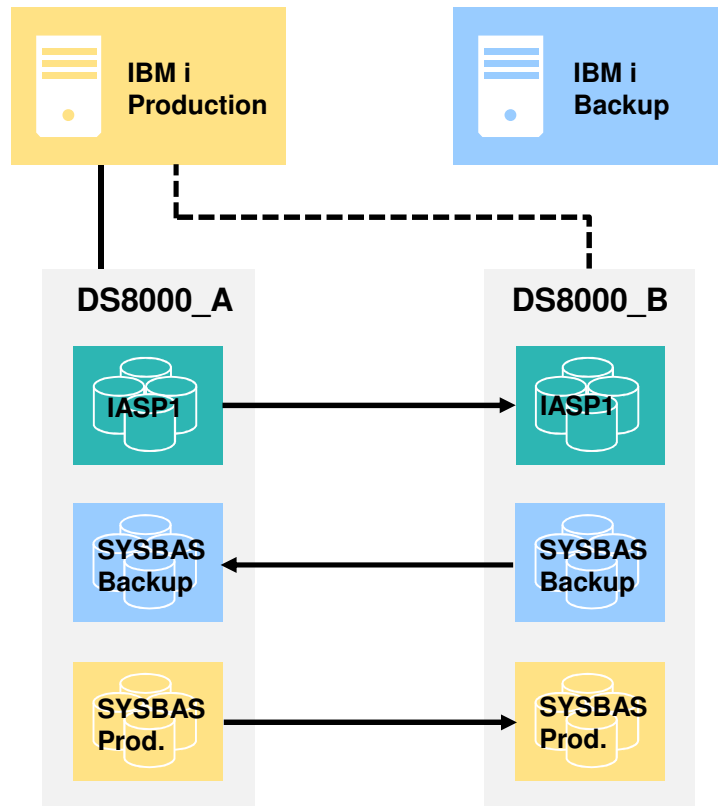
- Storage switched via SLIC



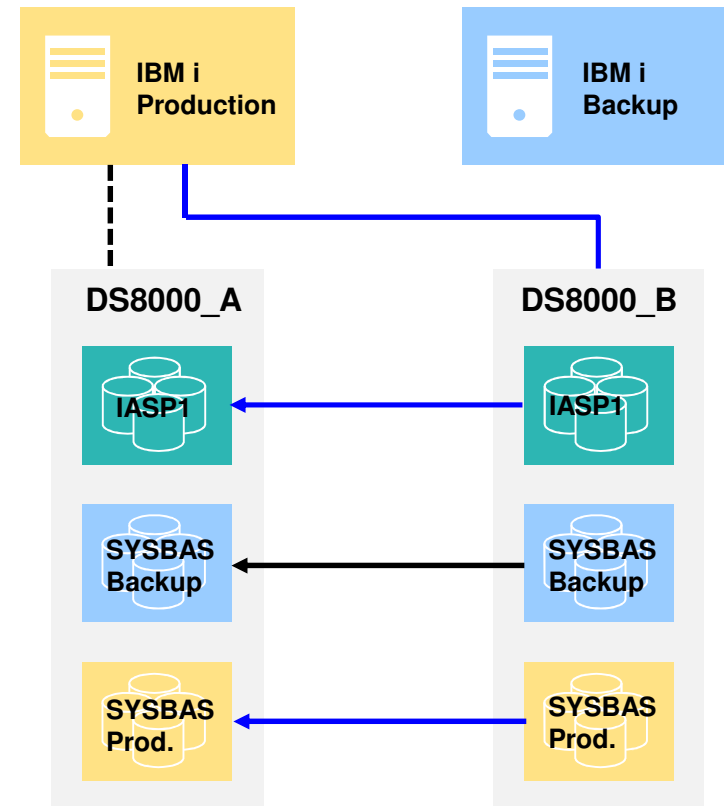
HyperSwap Cluster

IBM i DS8000 IASP planned HyperSwap

1) Normal Operation

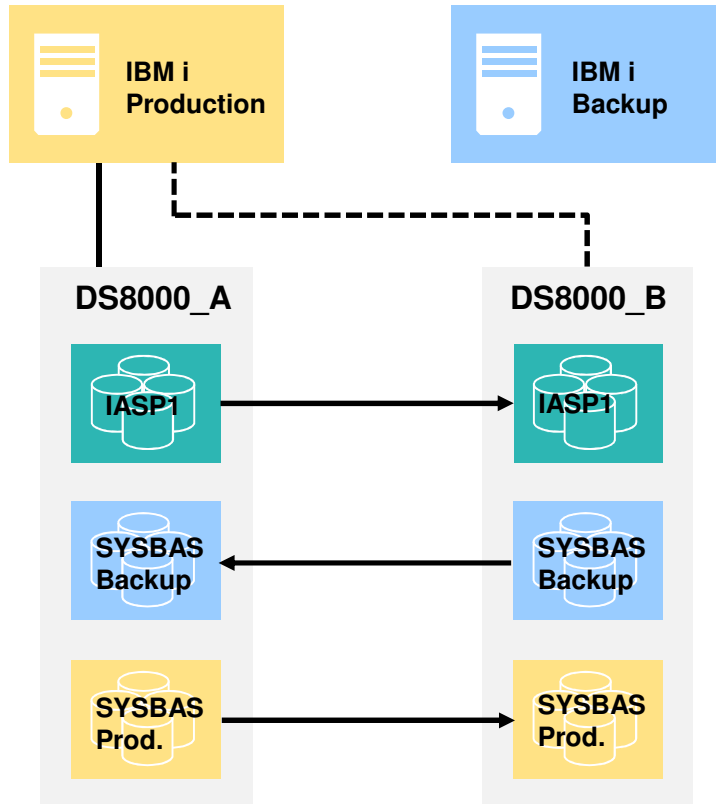


2) Planned HyperSwap to secondary DS8000 with near-zero downtime

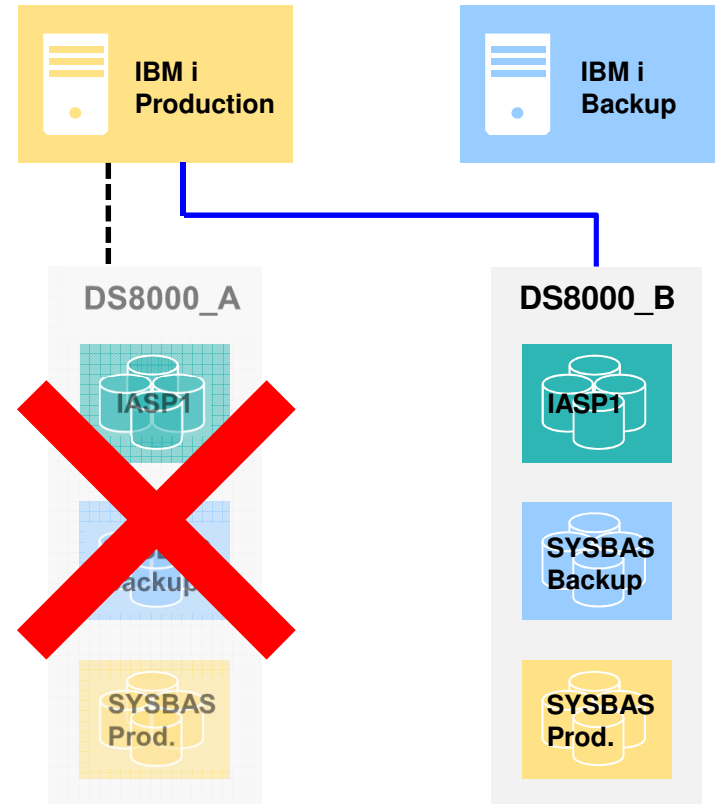


IBM i DS8000 IASP unplanned HyperSwap

1) Normal Operation

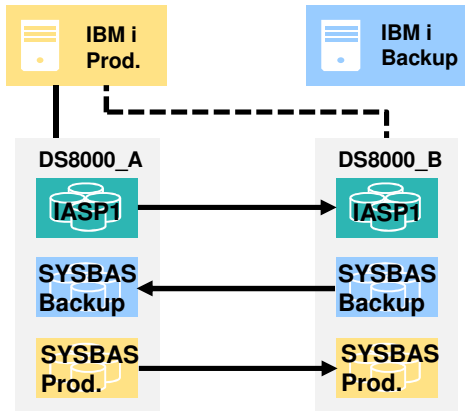


2) Unplanned HyperSwap to secondary DS8000 w/ minimal downtime (seconds to minutes; no IASP vary-off)

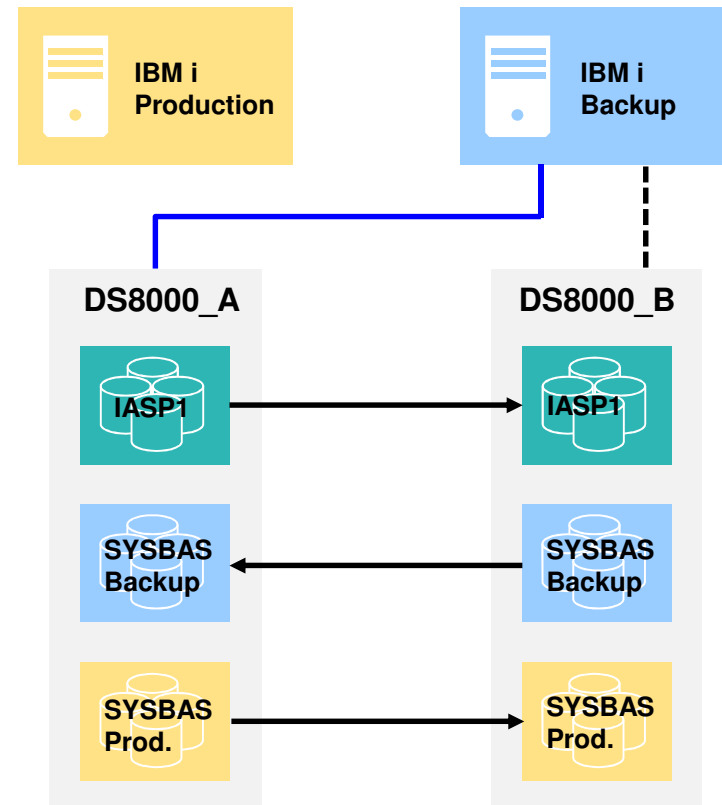


IBM i DS8000 IASP HyperSwap Planned CRG Switchover

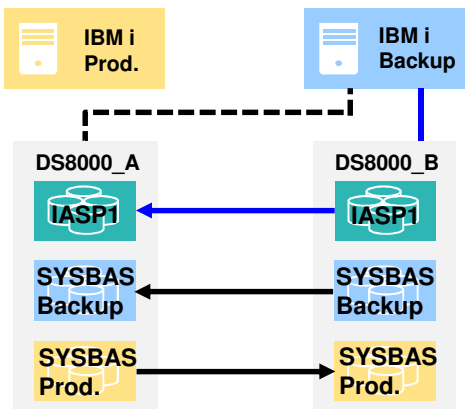
1) Normal Operation



2) Planned IASP switchover w/ minimal downtime (IASP vary-off/on)

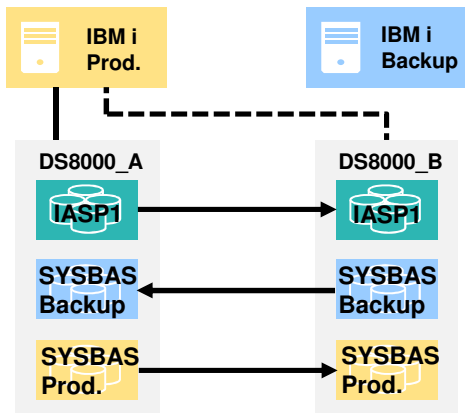


3) Automatic HyperSwap to correct storage affinity

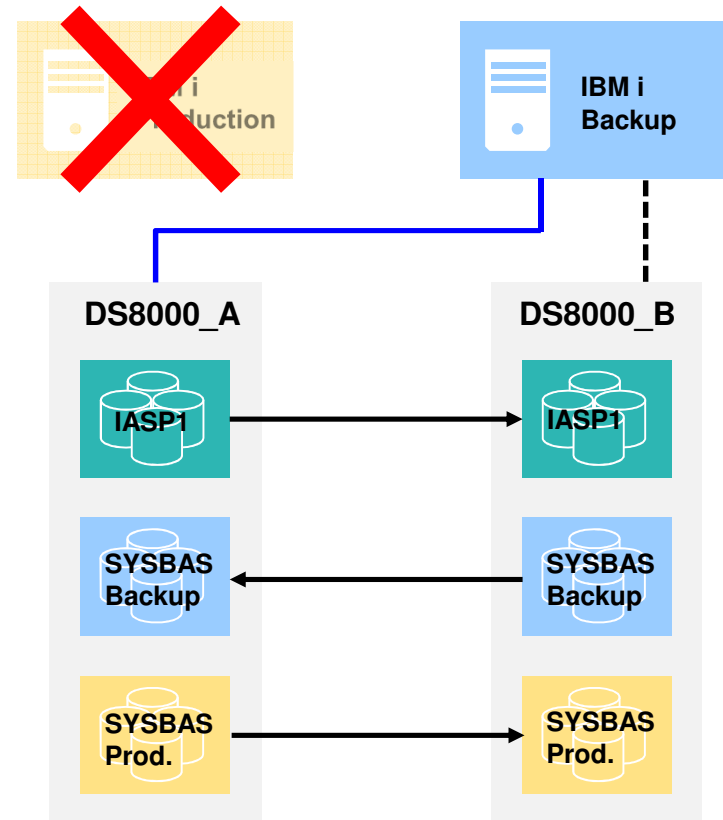


IBM i DS8000 IASP HyperSwap Unplanned CRG Failover

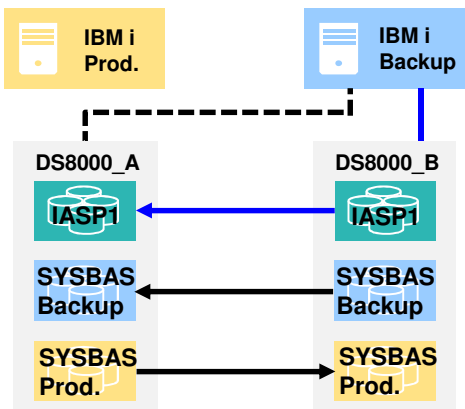
1) Normal Operation



2) Failure of production system triggering IASP failover w/ minimal downtime (IASP vary-off/on)

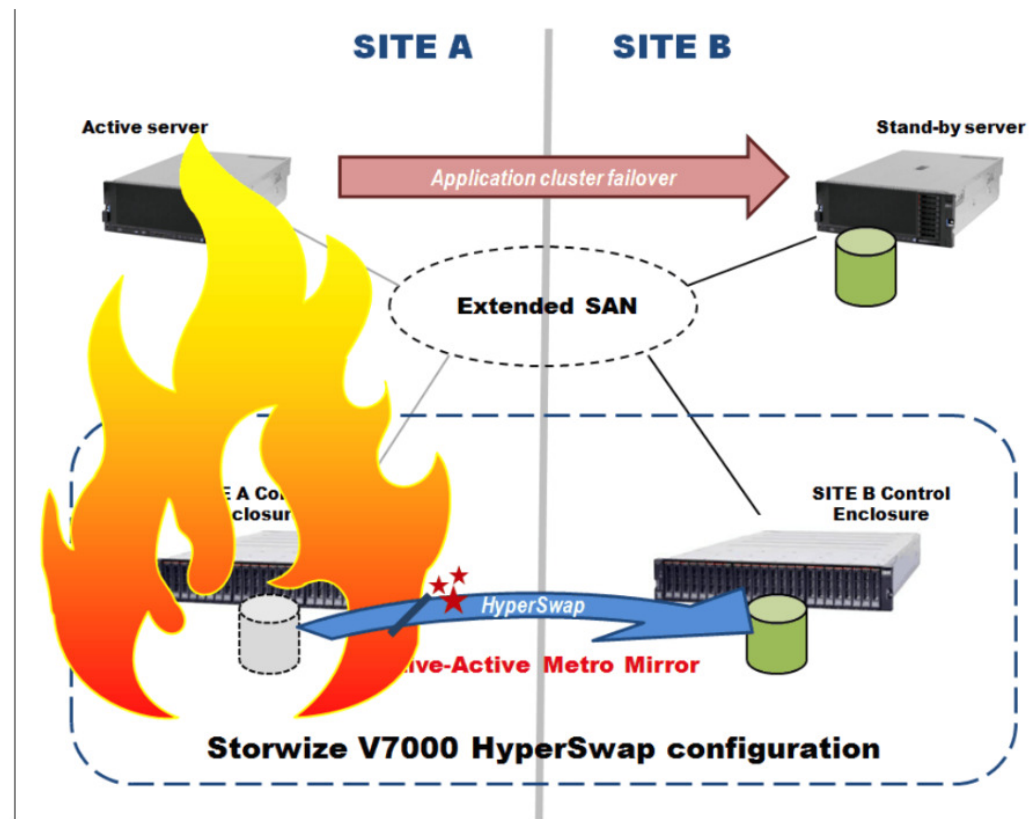


3) Automatic HyperSwap to correct storage affinity



SVC & Storwize Hyperswap with LUN Level Switching

- PowerHA for V7R2 & V7R3 support Hyperswap with IASP.



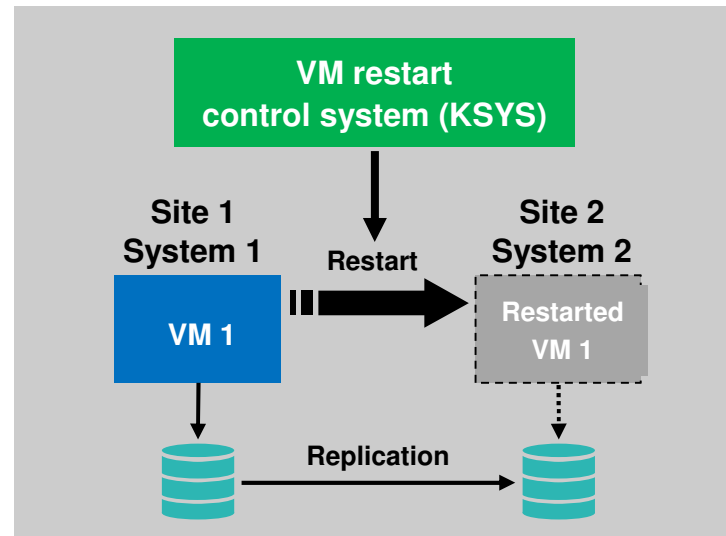
GDR for Power Systems

Geographically Dispersed Resiliency

What is the Geographically Dispersed Resiliency (GDR) for Power Systems DR solution?


A simplified way to manage DR

- Automated DR management
- Improved economics by eliminating the need for hardware and software resources at the backup site
- Easier deployment for disaster recovery operations; unlike clustering or middleware replication technologies, VM restart technology has no operating system or middleware dependencies.



- Support for IBM POWER7® and POWER8® Systems
- Support for heterogeneous guest OSs
 - AIX
 - Red Hat
 - SUSE
 - Ubuntu
 - **IBM i with GDR v1.1 SP1 (Ann May 9 – GA June 23)**

GDR for Power Systems – how it works



- The storage subsystem at the backup host is prepared and mapped to VIOS and then VM1 and VM2 are booted up
- VMs from site 1 are now restarted on the backup host in site 2
- The underlying mechanism that enables this to happen is the KSYS orchestrator at site 2
- From a customer perspective, this operation is accomplished with a single command

The End

References

- IBM PowerHA SystemMirror for i website,
<http://www-03.ibm.com/systems/power/software/availability/i>
- IBM i developerWorks PowerHA wiki,
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20PowerHA%20SystemMirror%20for%20i>
- IBM i Advanced Copy Services wiki,
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20i%20Advanced%20Copy%20Service%20s>
- PowerHA SystemMirror for i Cookbook,
<http://www.redbooks.ibm.com/abstracts/sg247994.html?Open>
- IBM PowerHA SystemMirror for i with IBM System Storage SAN Volume Controller Split-Cluster Setup Example
<http://w3.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD106126>
https://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/tech_TD106126
- IBM STG Lab Services website,
http://www-03.ibm.com/systems/services/labservices/platforms/labservices_power.html
- IBM Systems Magazine website “System FlashCopy Toolkit Could be a Lifesaver”,
http://www.ibmssystemsmag.com/ibmi/fsfc_toolkit/34841p1.aspx
- IBM System Storage Copy Services and IBM i: A Guide to Planning and Implementation (SG24-7103),
<http://www.redbooks.ibm.com/abstracts/sg247103.html?Open>
- IBM System Storage DS8000 Series: Implementing Copy Services in Open Environments (SG24-6788),
<http://www.redbooks.ibm.com/abstracts/sg246788.html?Open>

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

AS/400®, e business(logo)®, eServer, FlashCopy®, IBM®, IBM (logo)®, iSeries®, Power Systems, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, System Storage®, BladeCenter®

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 therefrom.

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.

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.

* 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 partly 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.