Crede Team

Tworzenie bezpiecznych kopii danych metodami replikacji fizycznej i logicznej (IBM Flash/Copy, Mimix)

Systemy IBM Power & Flash Systems Portfolio 2026



Jesienna Konferencja Common Polska 2026

Wisła

Power11 Portfolio Specs Overview

Model(s)	Footprint Capacity	Notes
High-End (4S+): IBM Power E1180	• 4 nodes (5U CEC) with 4 processors each • Up to 16 cores per socket/module • 64 TB – DDR5 based memory	• In-place upgrades with serial number preservation.
Mid-Range (4S4U): IBM Power E1150	• 2-4 processors in a 4U form factor • Up to 30 cores per socket/module • Up to 16 TB – DDR5 based memory	• In-place upgrades with serial number preservation.
Scale-Out 2S4U: IBM Power S1124 IBM Power L1124	• Up to 2 processors in a 4U form factor • Up to 30 cores per socket/module • Up to 8 TB – DDR5 based memory	
Scale-Out 2S2U: IBM Power S1122 IBM Power L1122	• 2 processors in a 2U form factor • Up to 30 cores per socket/module • Up to 4 TB – DDR5 based memory	• Integrating the eSCM models into the S1122 family
Entry/Edge: IBM Power S1112	• 1 processor in a 2U (1/2 wide) form factor OR tower • Up to 10 cores per socket/module • Up to 512GB – DDR5 based memory	• 1S2U will replace 1S4U model in both rack-mount and tower forms



IBM Power E1180

The industry's most scalable server for consolidating database and mission critical workloads



Single Node



IBM Power E1150

Built for large data & memory intensive workloads.

IBM Power Entry & Edge

Enterprise-class reliability, security





Power9 to Power10 to Power 11 Upgrade Options – IBM i SW Tiers

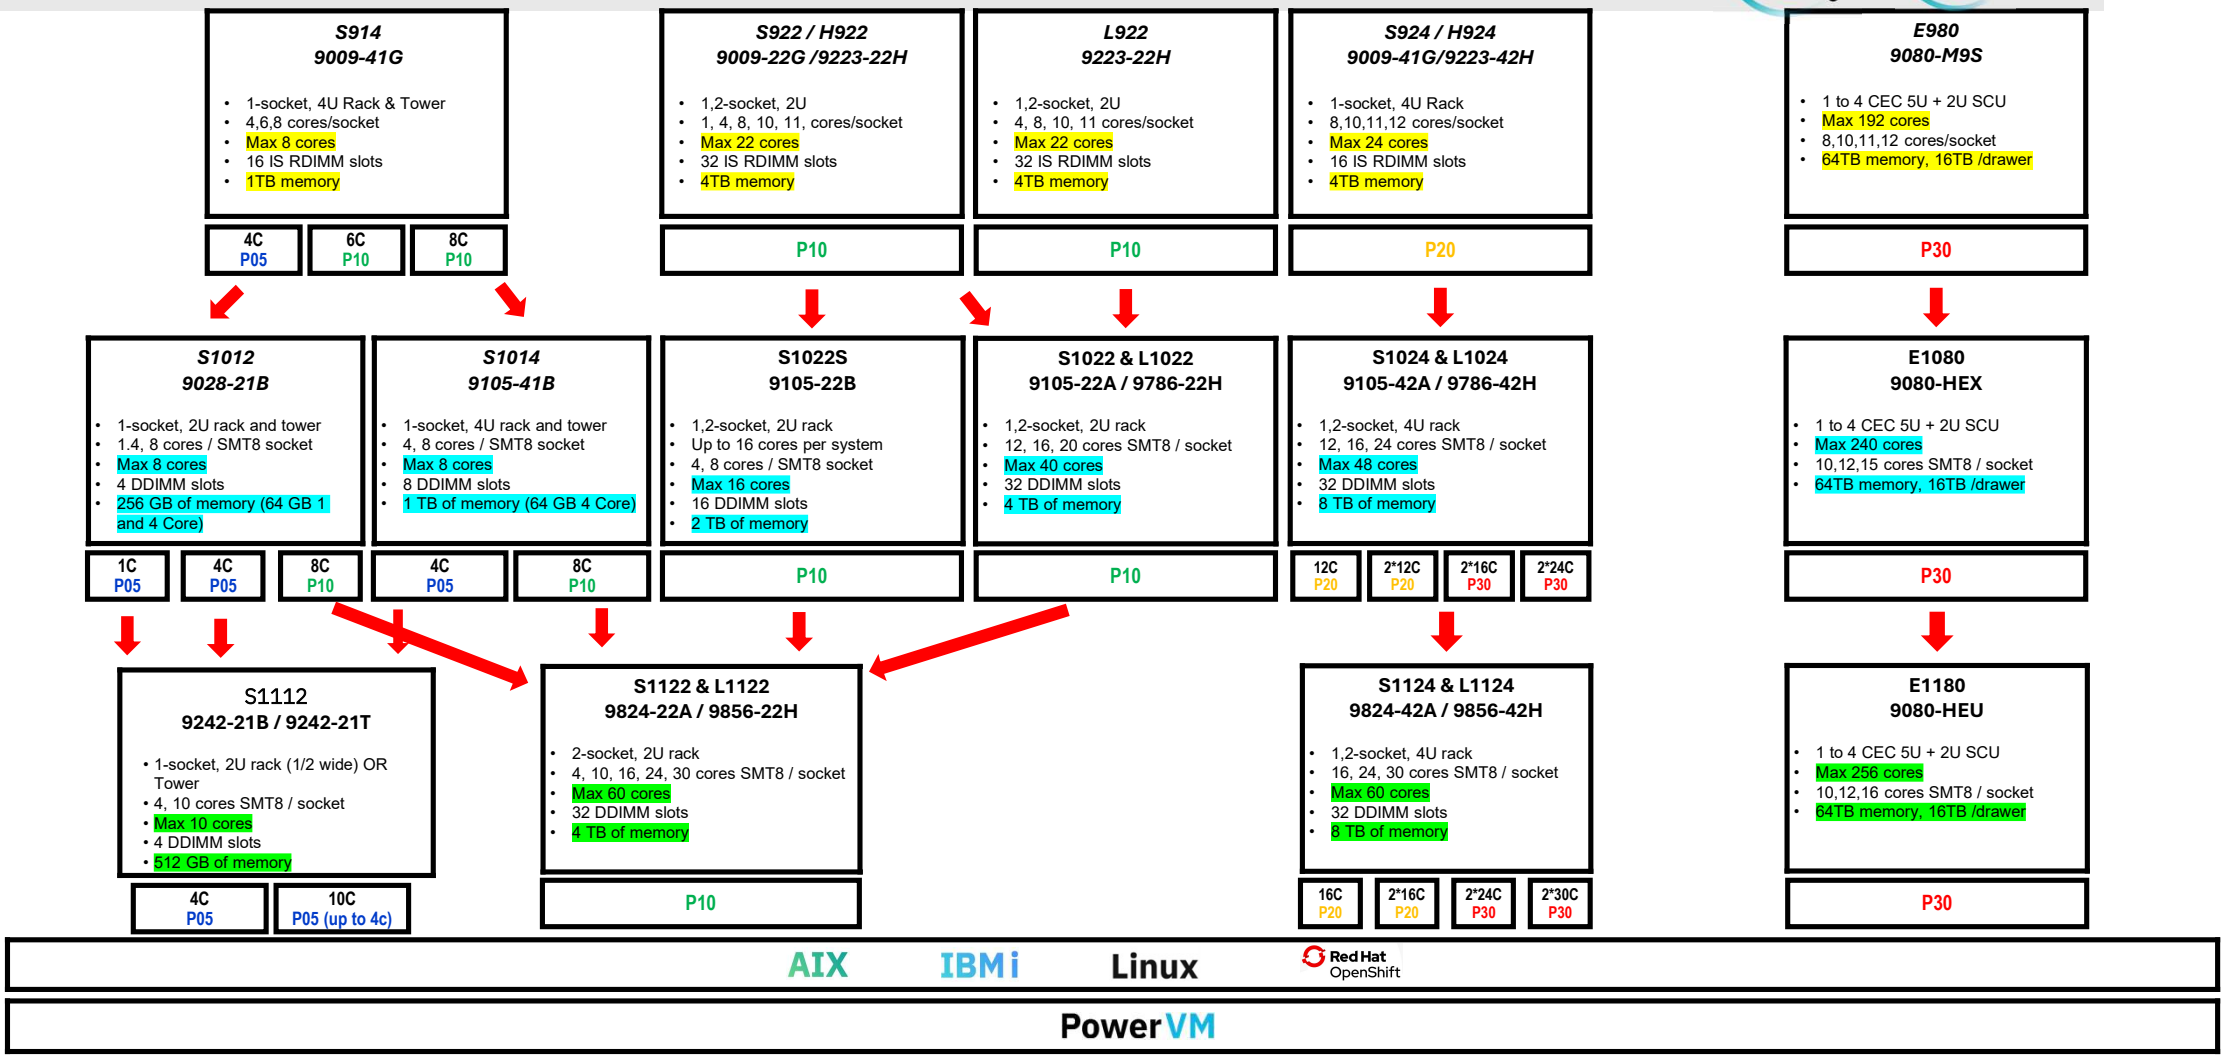

P05 P10 P20 P30 IBM i group processor

@2025 IBM Corporation

Power11: S1124



Hardware

- The S1124 will be offered with up to 2-sockets of DCM's
- 30-core DCM system with 2 sockets and 60 Power11 processor cores (2.8-3.95 GHz)
- 24-core DCM system with 2 sockets and 48 Power11 processor cores (3.05-4.15 GHz)
- 16-core DCM system with 1-2 sockets and 16-32 Power11 processor cores (3.4-4.2 GHz)
- Up to 8TB DDR5 memory

Software Requirements

- AIX 7.2 & 7.3 or later
- IBM i 7.4, 7.5 & 7.6 or later
- RHEL 8.6, 9.4, 9.6, and 10
- SLES15 & 16
- OCP 4.19 and later
- VIOS 3.1.4 (EOS is 4/2026)
  VIOS 4.1.0, VIOS 4.1.1

AIX

Power S1124 rPerf Projections (SMT8)		
Offerings	**Max. Cores**	**2-socket / 1-socket**
S1124/30c	60-core	1737 / NA
S1124/24c	48-core	1529 / NA
S1124/16c	32-core	1117 / 559

IBM i

Power S1124 CPW Projections (SMT8)		
Offerings	**Max. Cores**	**2-socket / 1-socket**
S1124/30c	60-core	1,345,900 / NA
S1124/24c	48-core	1,118,500 / NA
S1124/16c	32-core	823,000 / 433,200

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

Power S1124 (CPW) IBM i Migrations from Power8

- Consolidate 5-6 Power S824's onto 1 Power S1124 with increased core performance and increased core density
  - Reduce energy and cooling costs by 80%
  - Simplify management of infrastructure

- Optimize SW maintenance costs (Oracle, DB2, WAS, etc)
  - 125-150% Core performance increase vs Power S824
  - Repurpose licenses for future growth and new applications

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System (node) Performance Improvement for Power8 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S824 / 12c	6.2x	5.1x	3.8x
S824 / 8c	8.2x	6.8x	5.0x
S824 / 6c	10.9x	9.1x	6.7x

Core Performance Improvement

Per-core Performance Improvement for Power8 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S824 / 12c	146%	156%	182%
S824 / 8c	118%	127%	150%
S824 / 6c	118%	127%	150%

Power S1124 (rPerf) Migrations from Power8

AIX

- Consolidate 4-5 Power S824's onto 1 Power S1124 with increased core performance and increased core density
  - Reduce energy and cooling costs by 80%
  - Simplify management of infrastructure
- Optimize SW maintenance costs (Oracle, DB2, WAS, etc)
  - 65-90% Core performance increase vs Power S824
  - Repurpose licenses for future growth and new applications
- Reduce response times and batch workload times
  - Flat to slightly increased (+6%) Single-Thread Performance improvement (unless coming from the 8-core Power8 processor)

Based on published results as of 7/8/2025 available at https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System (node) Performance Improvement for Power8 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S824 / 12c	4.4X	3.9X	2.8X
S824 / 8c	5.7X	5.0X	3.7X
S824 / 6c	7.8X	6.9X	5.0X

Core Performance Improvement

Per-core Performance Improvement for Power8 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S824 / 12c	75%	92%	111%
S824 / 8c	52%	67%	83%
S824 / 6c	57%	72%	89%

Power S1124 (CPW) IBM i Migrations from Power9

- Consolidate ~3 Power S924's onto 1 Power S1124 with increased core performance and increased core density
  - Reduce energy and cooling costs by up to 50%
  - Simplify management of infrastructure
- Optimize SW maintenance costs
  - 45-50% Core performance increase vs Power S924
  - Repurpose licenses for future growth and new applications
- Reduce response times and batch workload times
  - 2-9% Single-Thread Performance improvement

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System (node) Performance Improvement for Power9 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S924 / 12c	**3.6X**	3.0X	2.2X
S924 / 11c	**3.9X**	3.2X	2.4X
S924 / 10c	4.2X	**3.5X**	2.6X
S924 / 8c	5.0X	4.1X	**3.0X**

Core Performance Improvement

Per-core Performance Improvement for Power9 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S924 / 12c	**45%**	51%	67%
S924 / 11c	**41%**	46%	62%
S924 / 10c	41%	**47%**	62%
S924 / 8c	32%	39%	**53%**

Power S1124 (rPerf) Migrations from Power9

AIX

- Consolidate ~3 Power S924's onto 1 Power S1124 with increased core performance and increased core density
  - Reduce energy and cooling costs by up to 50%
  - Simplify management of infrastructure
- Optimize SW maintenance costs (Oracle, DB2, WAS, etc)
  - 15-30% Core performance increase vs Power S924
  - Repurpose licenses for future growth and new applications
- Reduce response times and batch workload times
  - 5-13% Single-Thread Performance improvement

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System (node) Performance Improvement for Power9 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S924 / 12c	**3.0X**	2.6X	1.9X
S924 / 11c	**3.2X**	2.8X	2.1X
S924 / 10c	3.5X	**3.1X**	2.2X
S924 / 8c	4.1X	3.6X	**2.6X**

Core Performance Improvement

Per-core Performance Improvement for Power9 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S924 / 12c	**19%**	31%	44%
S924 / 11c	**18%**	30%	42%
S924 / 10c	16%	**28%**	40%
S924 / 8c	8%	19%	**31%**

Power S1124 (CPW) Migrations from Power10

IBM i

- Optimize TCO with Power S1124 with increased core performance and increased core density
- Optimize SW maintenance costs (Oracle, DB2, WAS, etc)
  - Up to 18% core performance increase vs Power S1024 for the constant core counts
  - Repurpose licenses for future growth and new applications
- Reduce response times and batch workload times
  - Up to 2% Single-Thread Performance improvement with up to 4% for same core counts

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System Performance Improvement for Power10 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S1024 / 24c	**1.42X**	1.18X	0.87X
S1024 / 16c	1.86X	**1.54X**	1.14X
S1024 / 12c	2.30X	1.91X	**1.41X**

Core Performance Improvement

Per-core Performance Improvement for Power10 to Power11 Migration Paths

	S1124 / 30c	S1124 / 24c	S1124 / 16c
S1024 / 24c	**14%**	18%	30%
S1024 / 16c	-1%	**3%**	14%
S1024 / 12c	-8%	-4%	**5%**

Power S1124 (rPerf) Migrations from Power10

AIX

- Optimize TCO with Power S1124 with increased core performance and increased core density
  - Reduce energy and cooling costs by over 20%
- Optimize SW maintenance costs (Oracle, DB2, WAS, etc)
  - Up to 15% core performance increase vs Power S1024 for the constant core counts
  - Repurpose licenses for future growth and new applications
- Reduce response times and batch workload times
  - Up to 4% Single-Thread Performance improvement with up to 7% for same core counts

Based on published results as of 7/8/2025 available at
https://www.ibm.com/downloads/documents/us-en/10c31775c5d40fed

System Performance Improvement

System Performance Improvement for Power10 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S1024 / 24c	1.30X	1.15X	0.84X
S1024 / 16c	1.71X	1.51X	1.10X
S1024 / 12c	2.14X	1.89X	1.38X

Core Performance Improvement

Per-core Performance Improvement for Power10 to Power11 Migration Paths			
	S1124 / 30c	S1124 / 24c	S1124 / 16c
S1024 / 24c	5%	15%	26%
S1024 / 16c	-9%	0%	10%
S1024 / 12c	-14%	-6%	3%

IBM widzi przyszłość centrów danych opartych na pamięciach flash



FlashSystem 5600

Type of media	NVMe
Rack Units	1U
Drive slots	12 drive slots[1]
Usable / Effective	403TBu / 2.42 PBe
Achieved with:	FCM 5 – 52.8 TB (XL)
CPU info per system	2x12 core Ice Lake D
CPU cores per system	24 cores
High-speed bus	PCIe Gen 4
Max Cache Hit IOPs:	2.6M
Max 70/30/50 IOPs 16KB	0.7M
Max Read Miss Bandwidth:	30GB/s
Max Bandwidth per drive	1.83 GB/s per drive
Max HBA slot per system	4
Max Gb FC ports	16-32Gb FC ports 12-64Gb FC ports
Max GbE ports	4 - 25GbE on-board 16 – 25GbE ports



FlashSystem 7600

Type of media	NVMe
Rack Units	2U
Drive slots	32 drive slots
Usable / Effective	1.22PBu / 7.29 PBe
Achieved with:	FCM 5 – 52.8TB (XL)
CPU info per system	2x16 core AMD Epyc
CPU cores per system	32 cores
High-speed bus	PCIe Gen 5
Max Cache Hit IOPs:	4.3M
Max 70/30/50 IOPs 16KB	1.2M
Max Read Miss Bandwidth:	55 GB/s
Max Bandwidth per drive	1.87 GB/s per drive
Max HBA slot per system	8
Max Gb FC ports	32 – 32Gb FC ports 12 – 64Gb FC ports
Max GbE ports	32 – 25GbE ports



FlashSystem 9600

Type of media	NVMe
Rack Units	2U
Drive slots	32 drive slots
Usable / Effective	2.43PBu / 11.88 PBe
Achieved with:	FCM 5 – 105.6 TB (XXL)
CPU info per system	2x48 core AMD Epyc
CPU cores per system	96 cores
High-speed bus	PCIe Gen 5
Max Cache Hit IOPs:	6.3M
Max 70/30/50 IOPs 16KB	3.00M
Max Read Miss Bandwidth:	86 GB/s
Max Bandwidth per drive	2.68 GB/s per drive
Max HBA slot per system	8
Max Gb FC ports	32 - 32Gb FC ports 12 – 64Gb FC ports
Max GbE ports	32 – 25GbE ports

5th Generation FlashCore Modules

- Nowy format dla centrów danych (EDSFF – Enterprise Datacenter Form Factor)
- Najnowsza technologia pamięci flash NAND typu QLC o pojemności 2 TB
- Większa pojemność fizyczna
- Większa pojemność efektywna (kompresja sprzętowa 6:1)
- Wyższa wydajność
- Większy zakres odciążania obliczeniowego
- Wyższy poziom bezpieczeństwa
- Niezmiennna niezawodność technologii FlashCore



Capacity Points in TB	Small	Medium	Large	XL	2XL
Form Factor thickness	1T	1T	1T	1T	1T
RAW Capacity	6.6	13.2	26.4	52.8	105.6
Usable “physical" Capacity (1:1)	5.6	11.2	22.4	44.8	89.6
Maximum Data Reduction	6:1	6:1	6:1	6:1	4.88:1[1]
Effective Capacity	33.6	67.2	134.4	268.8	438.0
Platform Support	x600	x600	x600	x600	9600

precisely

Understanding IBM i Replication Technologies

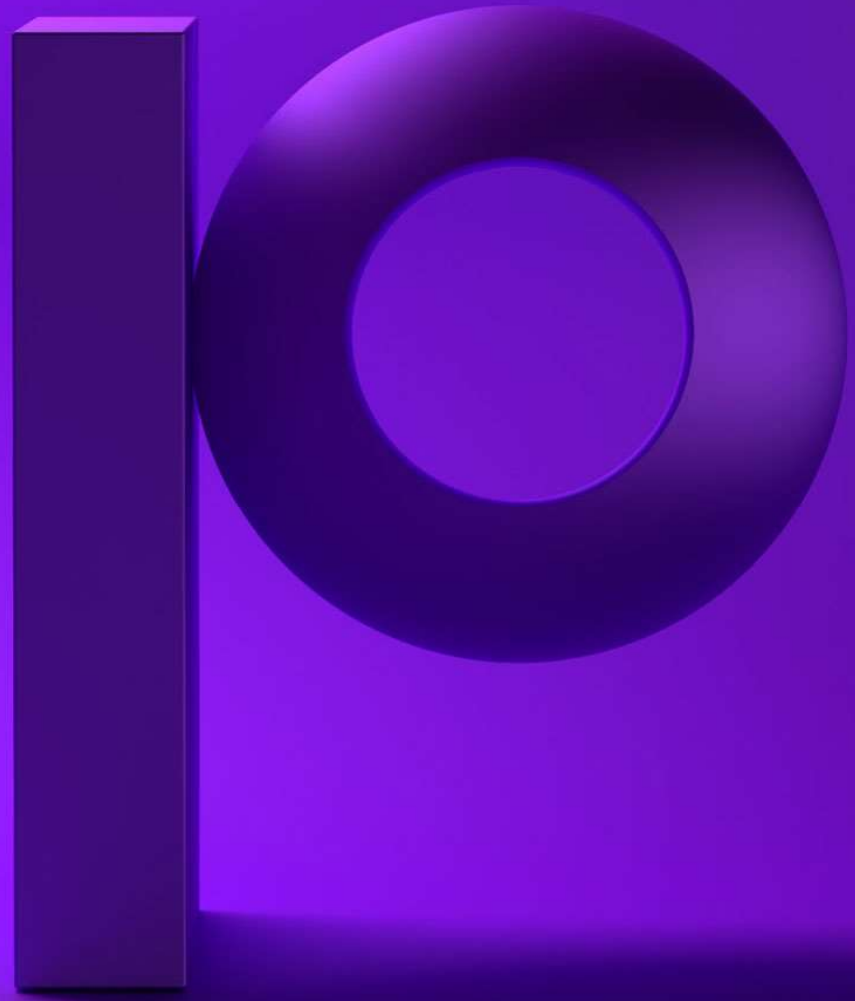

HA Comparison Topics

- Technology considerations
- Common scenarios
- Assure MIMIX offers more
- Q&A



Technology Considerations



1. Underlying architecture
2. Replication and recovery point
3. Bandwidth usage
4. Online or offline target
5. Application and data requirements
6. Damaged object replication
7. Access path recovery
8. Topology flexibility

Underlying Architecture

PowerHA

- IBM storage provides the underlying mechanism for SAN-to-SAN replication
- Sectors written to the source SAN are replicated to the tightly coupled target SAN
- Replicates applications and data in IASPs only
- Switch is done by attaching storage to secondary server and recovering to that server

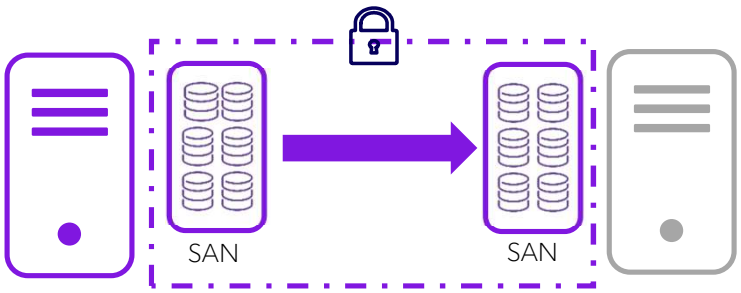

Assure MIMIX

- IBM Remote Journaling provides the underlying replication mechanism for Assure MIMIX
- Server to server replication
- Supports internal, external or mixed storage
- Replicates applications and data from *SYSBAS, ASPs, and IASPs
- Data is replicated at the end of a transaction
- Switch to a live target without recovery steps

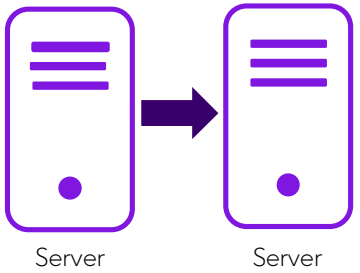

Replication and Recovery Point

PowerHA

- Data may be either in memory or on disk
- Data is not protected until written to disk
  - Unknown time after transaction completes
  - Results in non-zero, unpredictable RPO
- Journaling is required with PowerHA
  - Journaling saves data to disk before each transaction completes
  - Provides transaction integrity and near-zero RPO

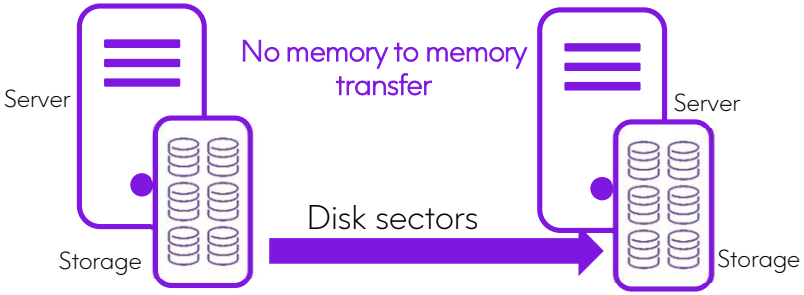

Assure MIMIX

- Journaling eliminates issues associated with delayed disk writes due to memory caching
- Journaling saves data to disk before completion of each transaction
- Remote journaling provides zero RPO (synchronous RJ) or near-zero RPO (asynchronous RJ)

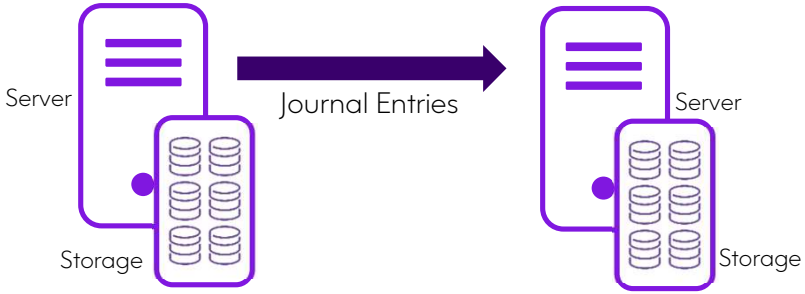

Bandwidth Usage

PowerHA

- Bandwidth needed for SAN replication is over twice that needed by logical replication
  - Journaled data is replicated
  - All sectors (files, objects and temporary files) are replicated again when written to disk
- Initial synchronization and any re-sync processes are also bandwidth intensive

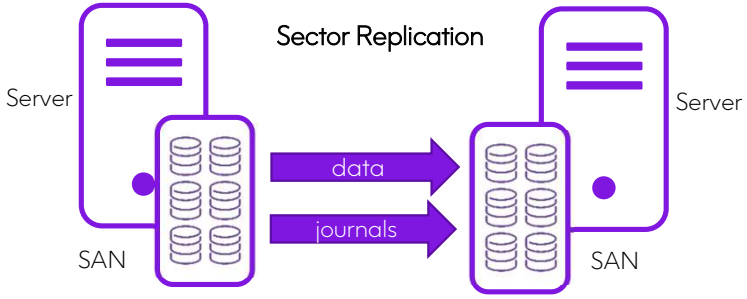

Assure MIMIX

- Journaling replicates only changed data
- Bandwidth can be further reduced through
  - Filtering non-critical data from replication
  - Journal optimizations
  - RJ filtering (Option 42 – HA Journal Performance)
  - JRNMINDTA
  - Hardware compression

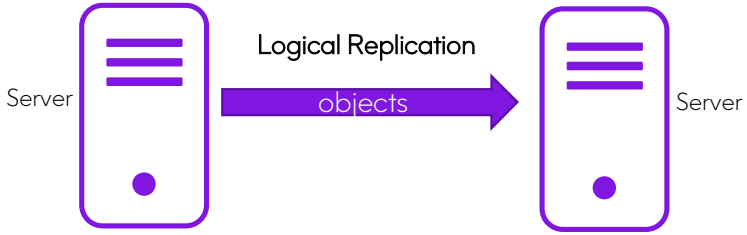

Online or Offline Target

PowerHA

- SAN to SAN replication does not support a live backup server
  - Must be varied on to a server to be used
- Some users add FlashCopy to PowerHA
  - Only adds a point-in-time copy of the data
  - Must be varied on to a server to be used
  - Not real time so quickly becomes stale

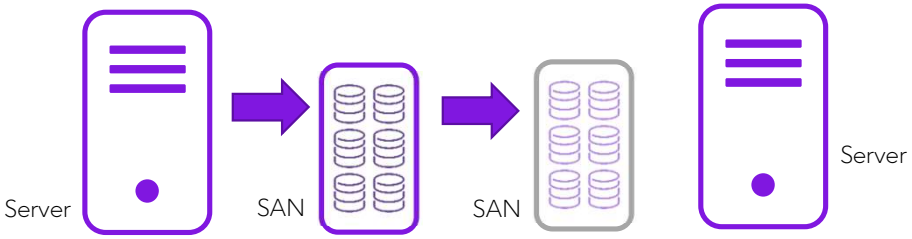

Assure MIMIX

- Backup server is always current and live
- An Assure MIMIX backup server can be used to off-load workload from the production server, including
  - Reporting
  - BI
  - Testing
  - Tape saves
  - More

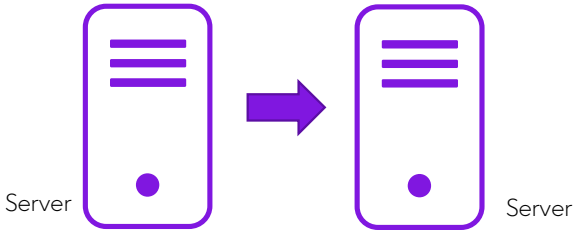

Application and Data Requirements

PowerHA

- Applications must run in an IASP to be protected by PowerHA
  - Efforts to convert and test applications can be substantial
- Application data must also reside in the IASP to be protected by PowerHA
  - Application data in *SYSBAS is not protected
  - Admin Domain does not fully protect *SYSBAS

Assure MIMIX

- No application changes required
- Assure MIMIX is independent of the underlying storage configuration
- Assure MIMIX can replicate to and from *SYSBAS, ASP, IASP or mixed environments
- Applications for every vertical are supported (e.g. banking, retail, logistics, health, etc.)

Damaged Object Replication and Repair

PowerHA

- Objects damaged on the source system will be propagated to the target
- Damaged objects detected during the IPL or vary on are marked by the OS
- Damaged objects require recovery from back-up which negatively impacts both RTO and RPO

Assure MIMIX

- Assure MIMIX cannot propagate damaged objects from the source to the target
- Assure MIMIX audits ensure objects damaged on either the source or target are detected, repaired and reported
- Live target data is always ready for switch without added object recovery steps

Access Path Recovery

PowerHA

- Recovery from unplanned failure may require access paths to be rebuilt
- Time to rebuild access paths varies based on storage, system, file and the access path itself – can be a lengthy process
- If access paths are invalidated during a switch process' vary-on, the application cannot start until keyed access paths have been rebuilt

Assure MIMIX

- There is no exposure to rebuilding access paths during an Assure MIMIX switch
- Audits verify access paths
- Cannot replicate invalidated access paths from the source to target
- Allows you to minimize downtime by switching to the target during a long access path recovery on the production server

Topology Flexibility

PowerHA

- Source and target SAN hardware must be similar
- Source and target OS and levels must be similar
- PowerHA replicates from IASP to IASP
- Distance limitations must be understood for PowerHA options

Assure MIMIX

- Storage independent to support replication in SAN, internal storage, or mixed storage environments
- OS independent to support replication between source and target servers with different OS versions, TR and PTF levels
- Supports replication of IASP, ASP and *SYSBAS
- No distance limits, enabling replication between on-premises and cloud servers
- Supports broadcast to multiple targets

Behavior in Common Scenarios



1. Initial synchronization
2. Unplanned switch
3. Planned switch
4. Reporting
5. Target integrity validation
6. Recovery from network issues
7. Operating system upgrade
8. Storage event management
9. Workload distribution

Scenario 1 – Initial Synchronization

PowerHA

- SAN solutions require a full storage sync before steady state replication can begin
- All storage sectors are sent across the network
- GeoMirror limited to 2-5 TB or it may never complete the synchronization



Assure MIMIX

- Assure MIMIX offers many options for initial sync
  - Save and restore from tape
  - Storage copy
  - Physical drive swap
  - Audit over time
  - Object sync commands
  - And more
- Reduces impact on network

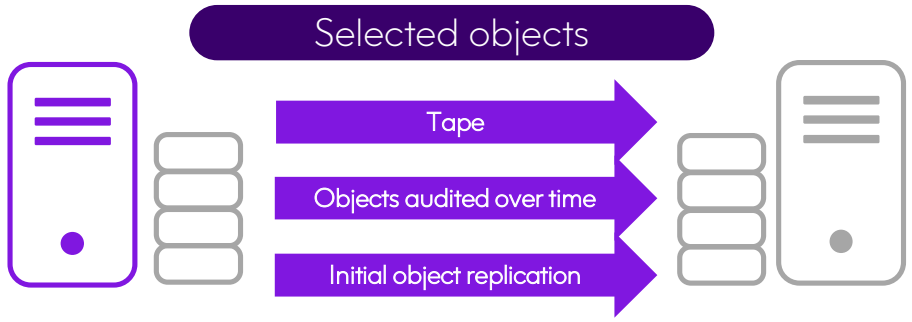

Scenario 2 - Unplanned Switch

PowerHA

- Target system must go through a vary-on process to recover transaction integrity – similar to an abnormal IPL
- Vary on has 34 steps – most of which are recovery steps of indeterminate length
- Some steps may result in invalidated access paths, damaged objects, etc.
- Damaged objects detected during vary on require manual intervention

Assure MIMIX

- Target system is a hot standby
  - No vary-on process required
  - Data has been audited for integrity
  - No need to rebuild access paths
- Single-button switch
  - Automated procedure
  - Parallelizes all possible steps
  - Provides visibility into progress of each step
- No post-switch recovery steps necessary

See diagram on following slide

Scenario 2 - Unplanned Switch Timeline

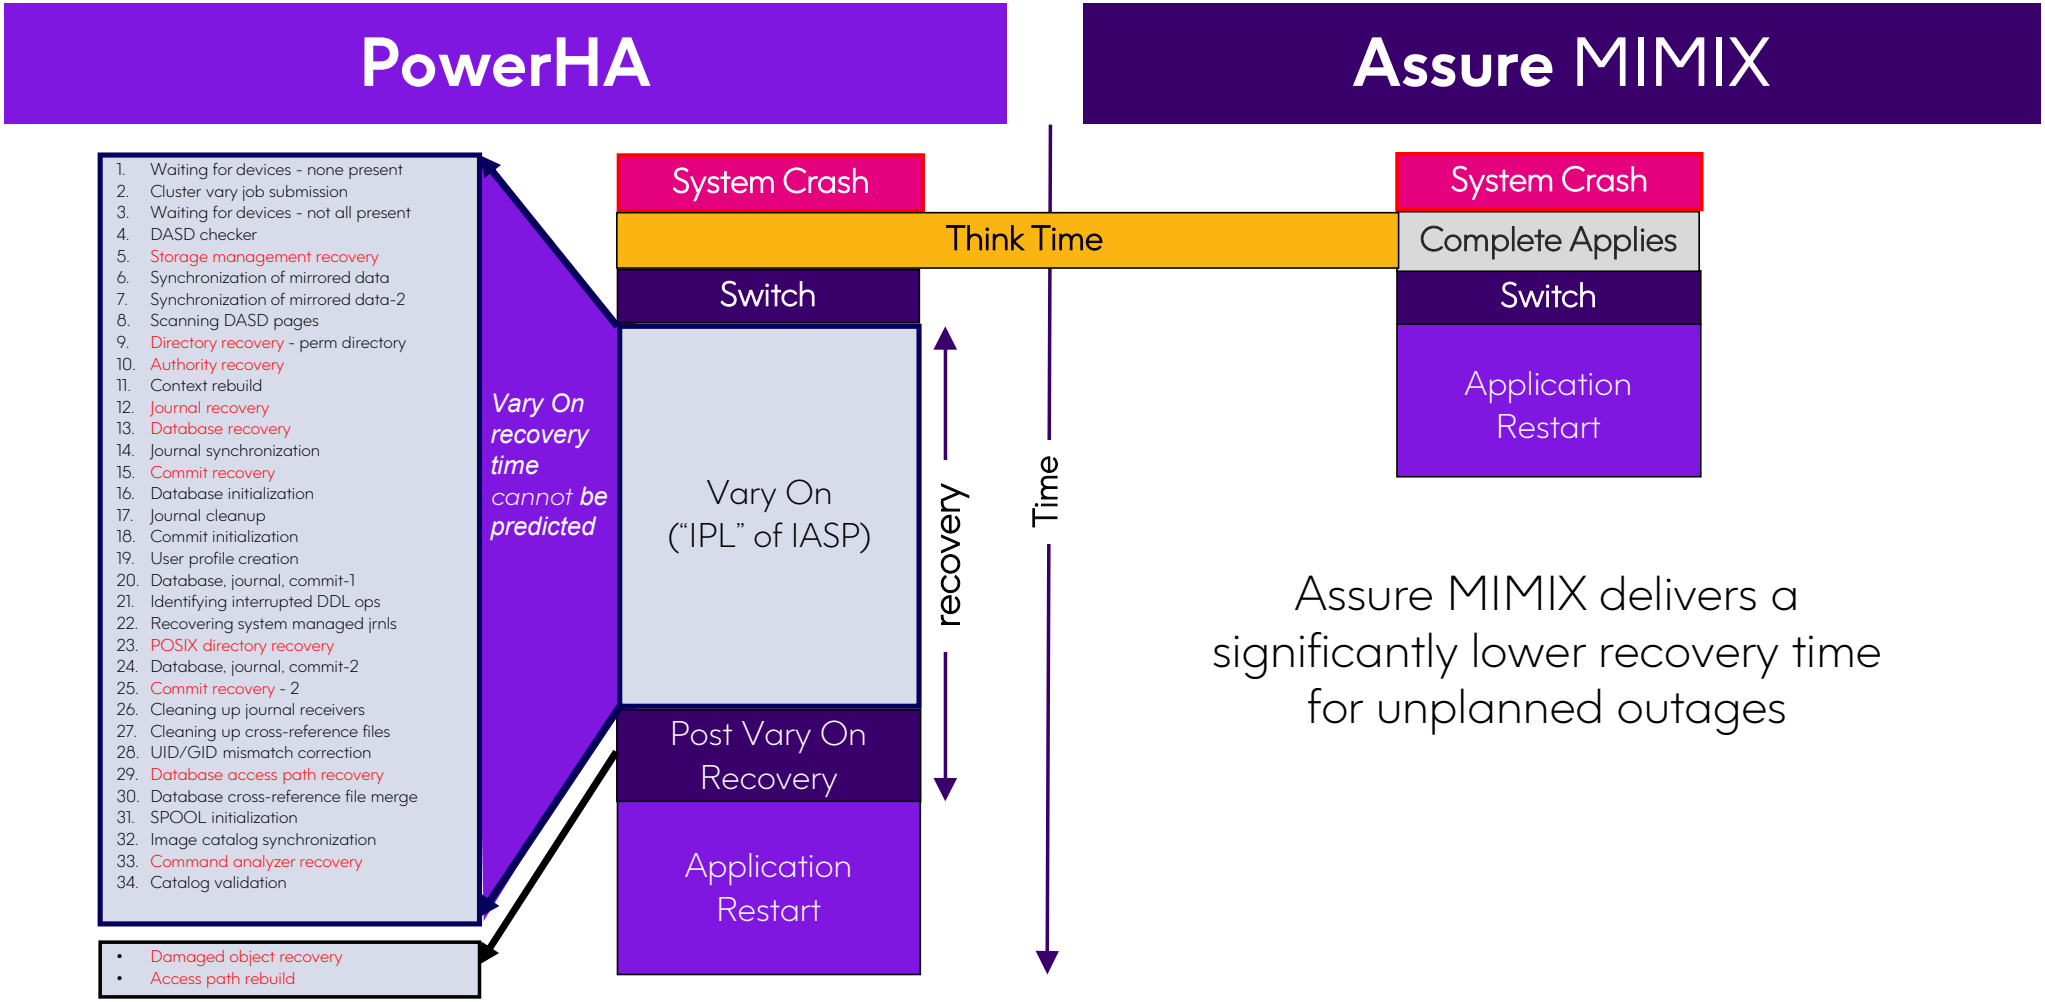

Scenario 3 - Planned Switch

PowerHA

- Requires vary off before the switch
- Requires 34-step vary on after the switch
- Vary on impacted by configuration issues, such as duplicate library names and UID/GID mismatches

Assure MIMIX

- Target system is a hot standby
  - No vary-on process required
  - Data regularly audited for integrity
- Single-button switch
  - Automated procedure
  - Parallelizes all possible steps
  - Provides visibility into progress of each step

See diagram on following slide

Scenario 3 - Planned Switch

PowerHA

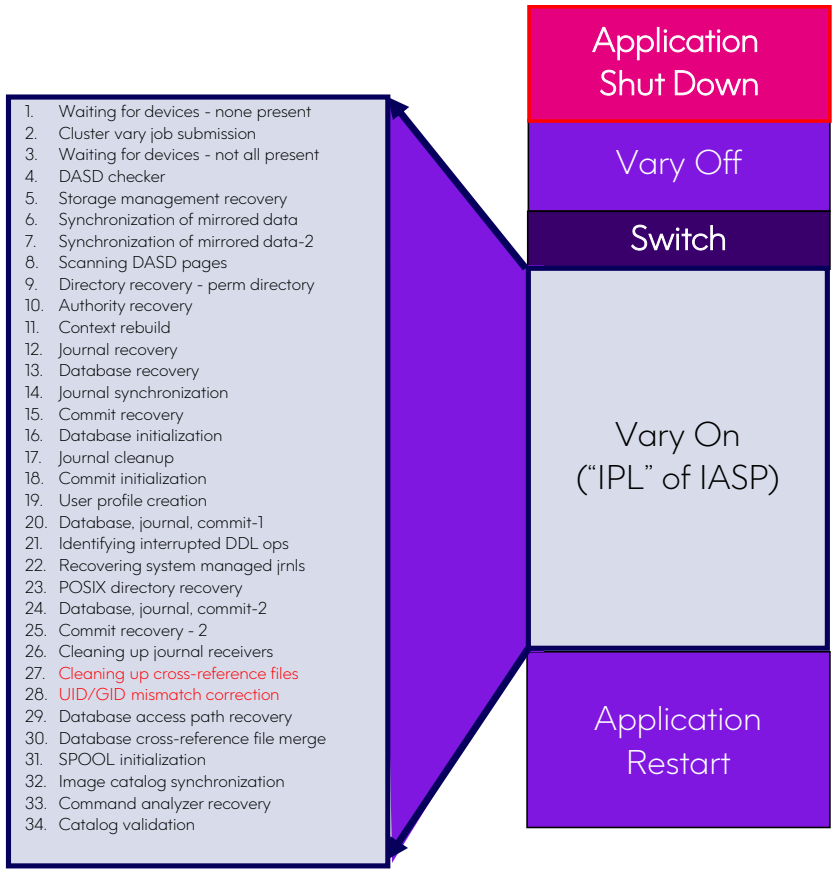

Time

Assure MIMIX

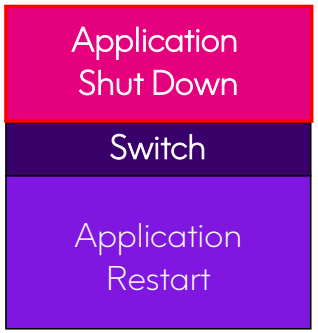

Assure MIMIX delivers a lower recovery time for planned outages

Scenario 4 – Reporting

PowerHA

- Data on the target SAN is not accessible until varied on to a server
- Other than point-in-time Flashcopy, PowerHA has no reporting option

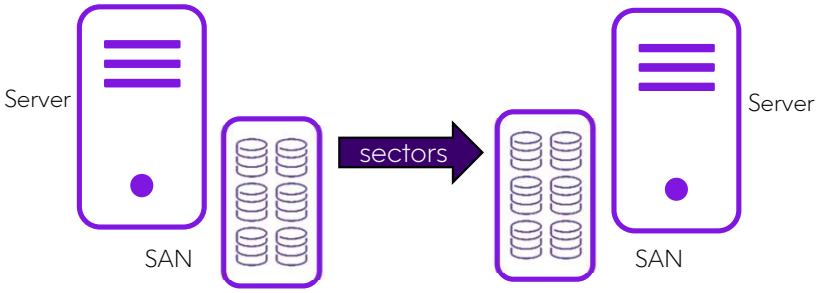

Assure MIMIX

- Reports can be run from any target
- RJ support for broadcasting enables multiple targets which can be used for queries and reports (as well as additional HA / DR nodes)
- Assure MIMIX simplifies management and switching in these multi-node environments

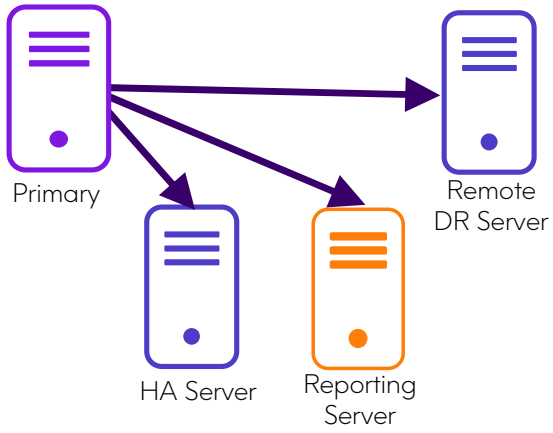

Scenario 5 – Target Integrity Validation

PowerHA

- Target SAN cannot be audited for integrity
- If target SAN becomes out of sync, partial or full resynchronization of the storage image is required
- Corruption introduced by dirty networks may be undetectable until the data is read



Assure MIMIX

- Assure MIMIX has multiple mechanisms for live target validation
  - Self-healing audits
  - Continuous monitoring of target data
  - Visibility into protection & audit status in VSP
- Virtual switch enables switch testing without disrupting production operations
- Remote Journaling offers validity check to detect network corruption



Scenario 6 – Recovery from Network Issues

PowerHA

- Depending on the extent of the outage, a partial or full synchronization of stored sectors may be required
- SAN is not protected until the full or partial resync is complete



Assure MIMIX

- After communication resumes, remote journaling repackages journal entries into large block transfers to quickly catch up
- Protection of those entries is in effect as soon as the transfers catch up



Scenario 7 – Operating System Upgrade

PowerHA

- Once you switch to the upgraded target server, you cannot switch back to the original production server until the upgrade completes successfully
- There is no option to go back to a lower OS level after switching to a higher OS level

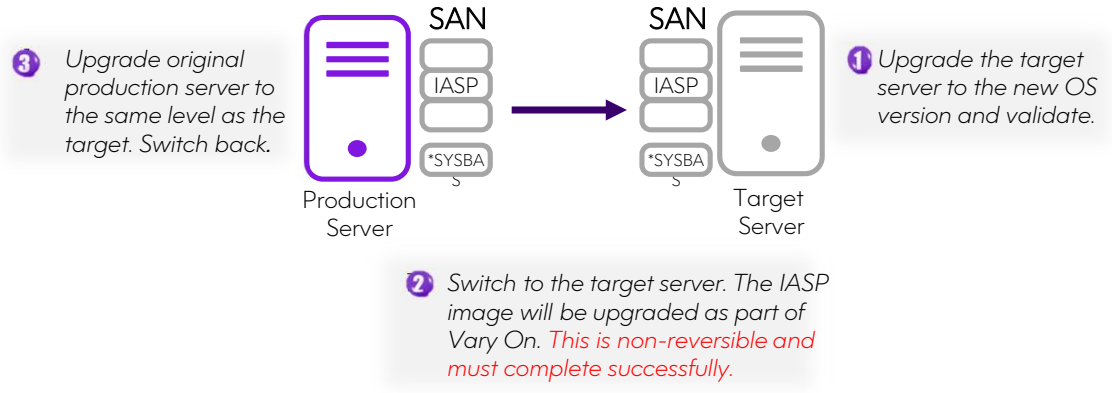

Assure MIMIX

- **Assure** MIMIX "upgrade, test and switch" process enables upgrades without risk and downtime
- Replication from higher to lower OS levels is fully supported
- 3rd node removes risk of outage during upgrade

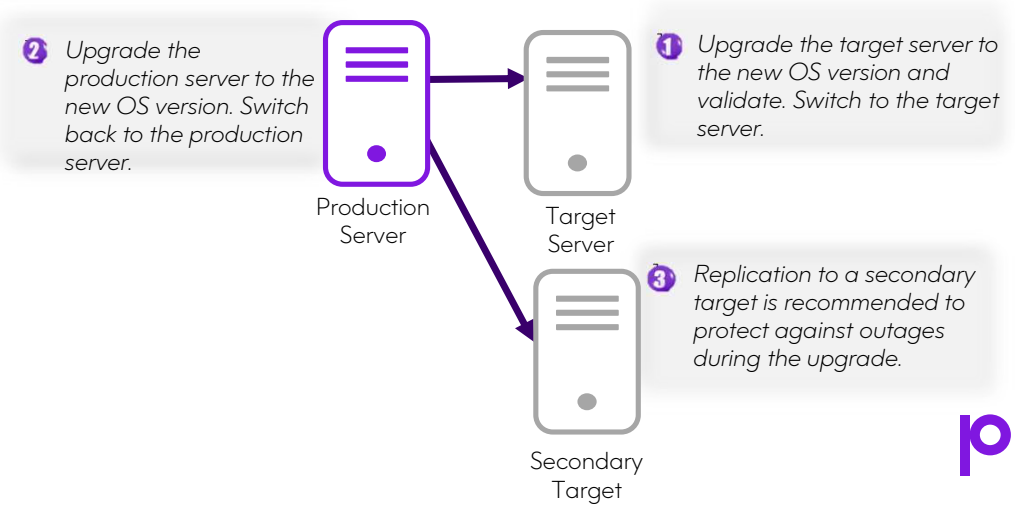

Scenario 8 – Storage Event Management

PowerHA

- Source and target storage images are tightly coupled (bit by bit copy)
- Single system storage events or storage image issues can require production downtime
- For example, Reclaim Storage (RCLSTG) events create protection downtime

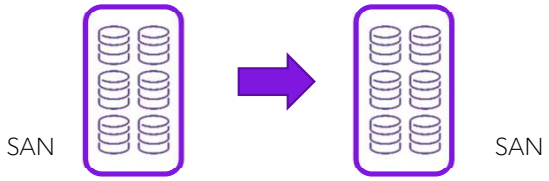

Assure MIMIX

- Assure MIMIX provides protection from storage events
  - Assure MIMIX is isolated from the storage layer
  - IBM Remote Journaling provides the underlying replication mechanism
- Storage events can be addressed with little impact on production operations
- For example, can switch to target while doing Reclaim Storage (RCLSTG) on production

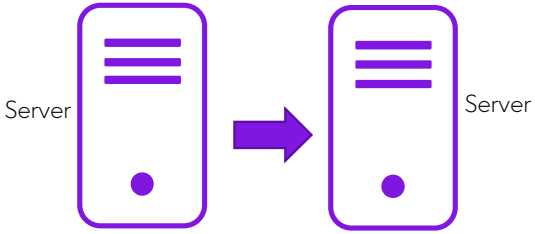

Scenario 9 – Workload Distribution

PowerHA

- Target SAN is locked and cannot be read from or written to
- Workloads run entirely on the source and cannot be balanced between the two systems
- Some PowerHA users add FlashCopy to offload work to a snapshot; however, it is point in time and requires more resources

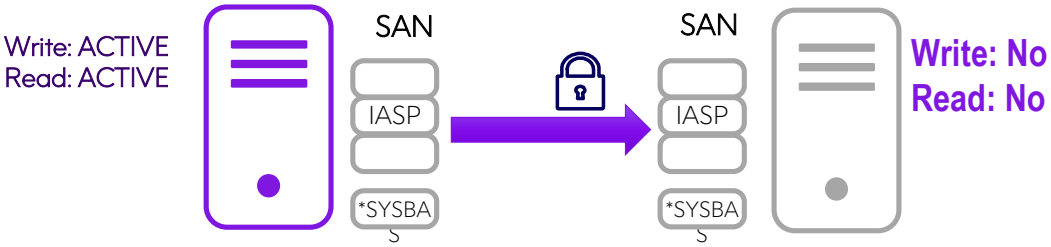

Assure MIMIX

- Assure MIMIX supports multiple replication topologies (e.g.1-to-1, broadcast, cascade, bidirectional)
- Standard 1-to-1 HA/DR would use a readable target that cannot be written
- Active-active topology enables workload distribution between two servers working on the same customer data and a near-zero RTO

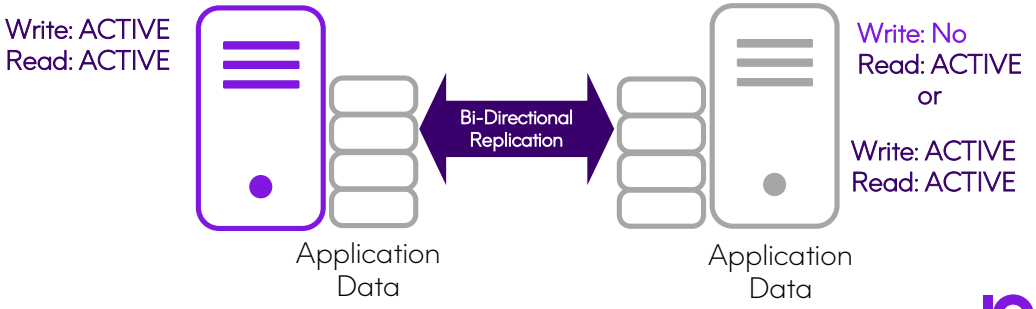


precisely

Supplemental Slides



FlashCopy MIMIX and Assure

IBM FlashCopy / Snapshot Basics

FlashCopy / snapshot basics

- Storage level technology for any OS
- Comes with PowerHA but does not require it
- Requires external storage
- Quick, easy way to create a copy of an LPAR
- Can be done in *SYSBAS and/or IASP environments

Subject to IBM i single level store challenges

- Image may be unstable based on timing of flash / snapshot
- Data in memory cache will not be copied
- Long commit cycles may be open
- Apps must pause or quiesce for a consistent snapshot

Flashed LPAR must be IPL'd and attached to a server

Safe method for stand alone FlashCopy backup

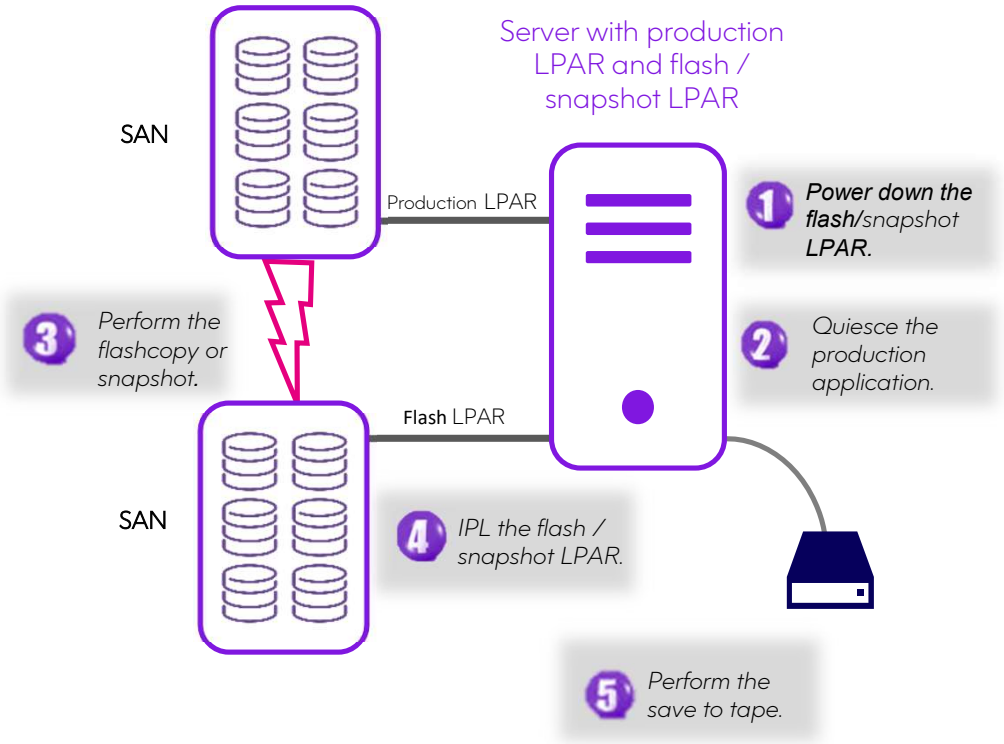

Assure MIMIX **with FlashCopy**

Assure MIMIX **with FlashCopy basics**

- Production server is protected by **Assure** MIMIX
- FlashCopy is done on the backup server without impacting the production server

Several Assure MIMIX **options eliminate challenges of single level store architecture**

- Perform flash with no pause of any kind
- End **Assure** MIMIX apply jobs, but leave RJ active to preserve RPO and transaction integrity
- Fully end **Assure** MIMIX but leave LPAR running
- Shut down LPAR or bring it to a restricted state

Assure MIMIX **with FlashCopy provides the least disruptive solution for snapshot operations**

Easy, safe backups with Assure MIMIX and FlashCopy without production server impact

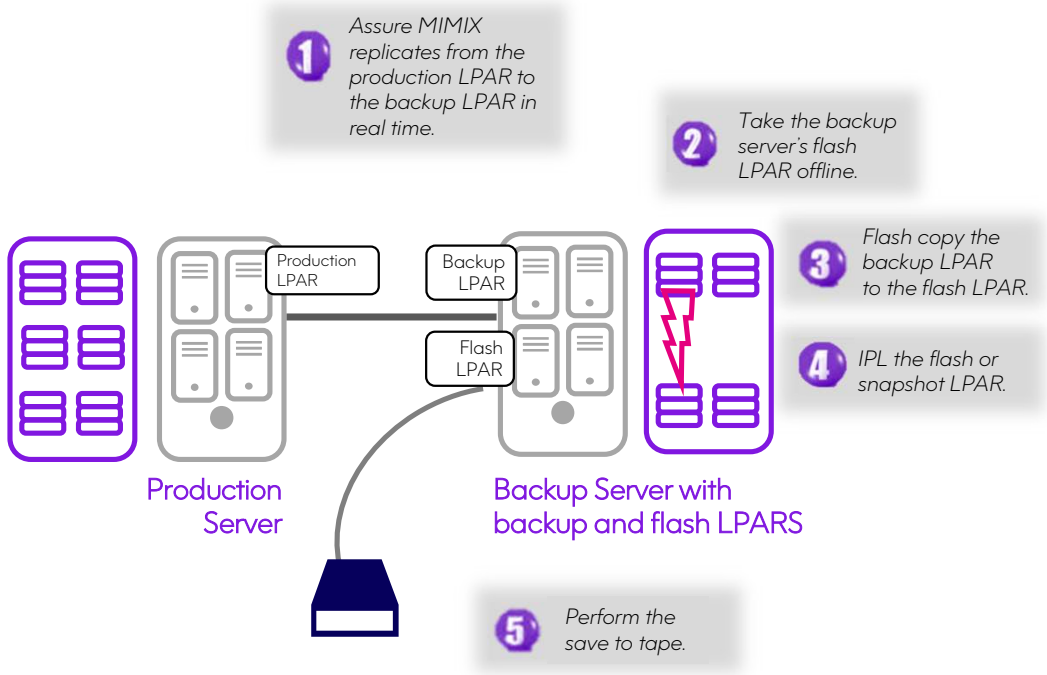

Live Partition Mobility and Assure MIMIX

What is Live Partition Mobility (LPM)?

The ability to move a running partition from one frame to another frame without disruption

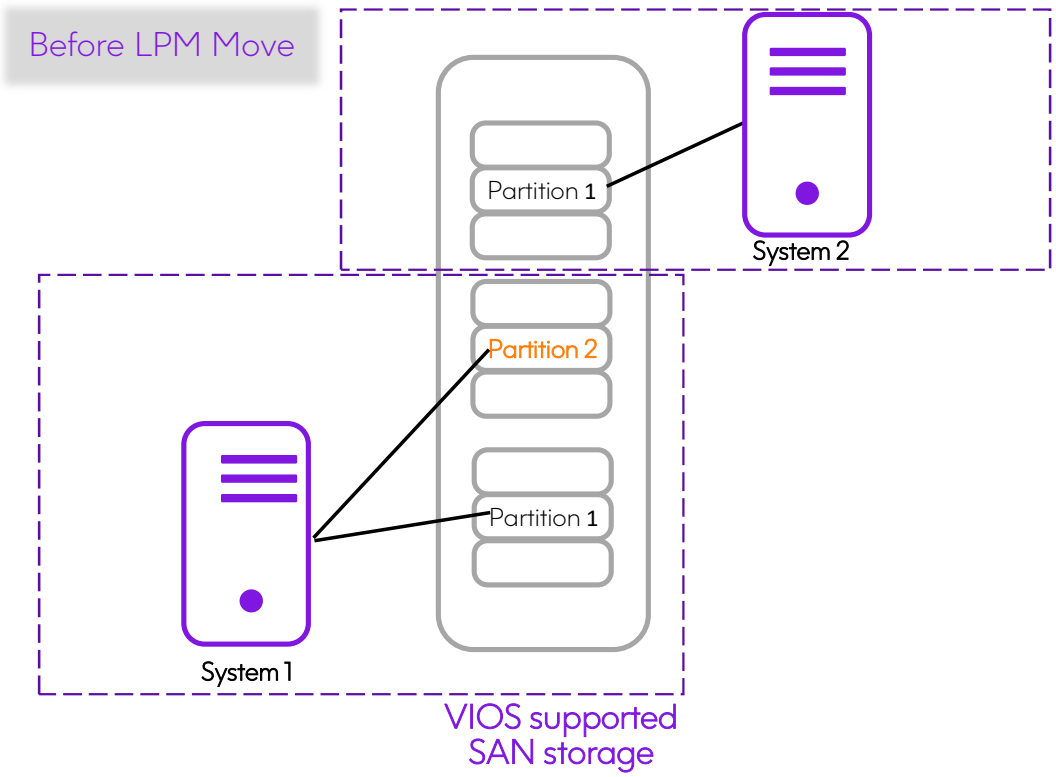

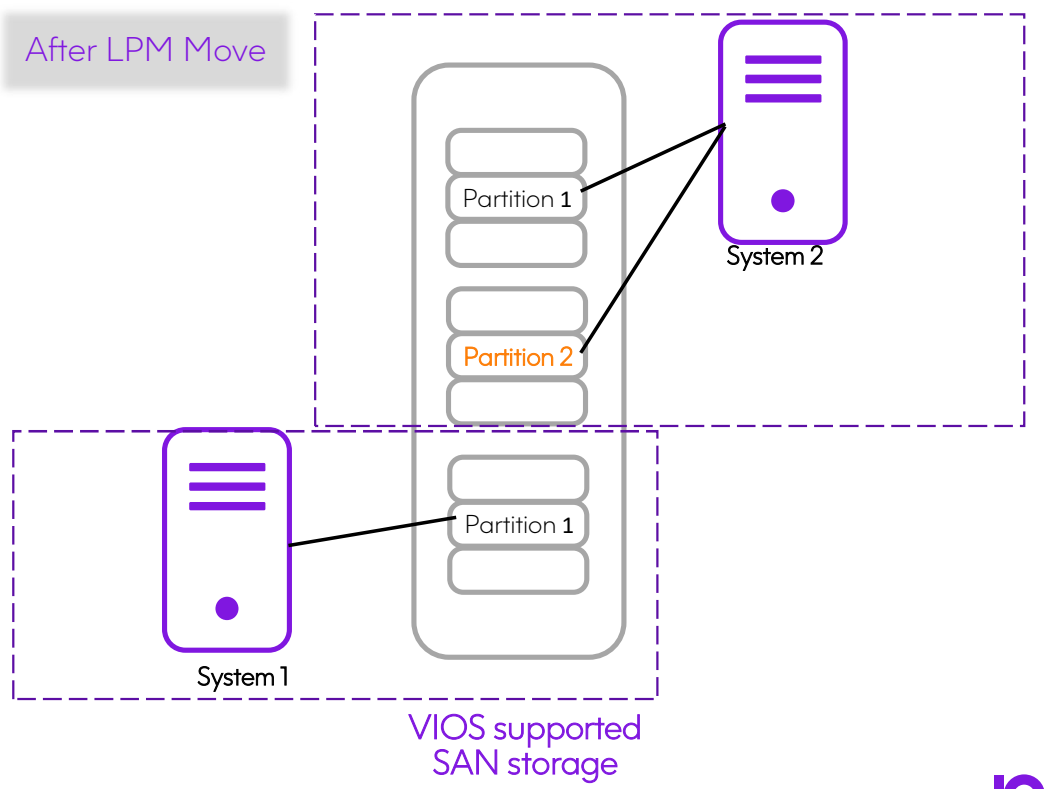

Live Partition Mobility Basics

Prerequisites for LPM

- Single SAN storage device attached to both frames with VIOS
- Both source and destination must be on same Ethernet network
- Power7 hardware and greater
- System i support was added in iOS 7.1, TR 4 (May 18, 2012)

Use Cases for Live Partition Mobility

- Workload consolidation on virtualized systems
- Workload balancing (move LPAR to an under-utilized system)
- Planned CEC outages for maintenance or upgrades
- Impending hardware outages
- Migration between Power processor generations

Live Partition Mobility Is Not an HA Replacement

- Single copy of the data
- Cannot be used for DR failover (both systems must be active)

"Live Partition Mobility increases global availability, but it is not a high availability solution. It requires both source and destination systems to be operational and that the partition is not in a failed state. In addition, it does not monitor operating system and application state and it is, by default, a user-initiated action. Unplanned outages still require specific actions that are normally executed by high availability solutions.

– IBM's PowerVM LPM Redbook

HA Licensing Considerations for LPM

- Customers should request temporary license keys for the going-to system before the move
- Apply the temporary key immediately following the LPM move
- LPAR and CPU licensing needs are handled independently

HyperSwap and Assure MIMIX

HyperSwap Basics

- Provides incremental increase in storage subsystem reliability over RAID and Mirrored DASD
- Swap of storage subsystem from one DS8000 to another in the event of a SAN failure
- Server is unaware of which SAN is the current source
- SANs must be in the same data center or on the same campus (within 25 miles)
- Coexists with server HA solutions but is not a server HA solution

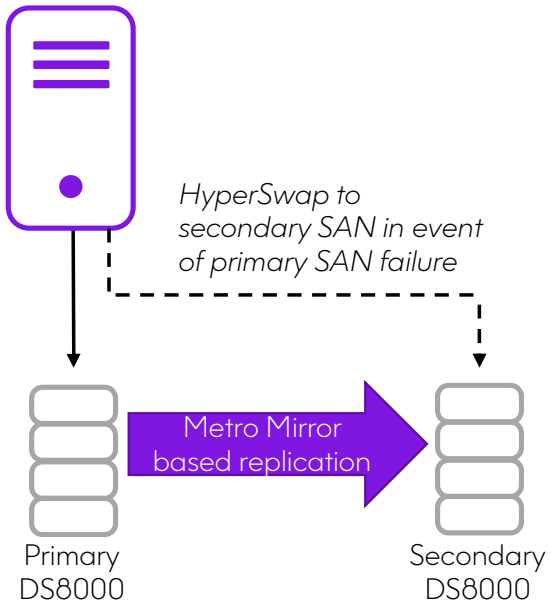

HyperSwap with PowerHA

HyperSwap can be used in a PowerHA IASP environment

- Enables an IASP or *SYSBAS to be swapped from one SAN to another
- *SYSBAS can be swapped but is not a solution for switching *SYSBAS to another server

Adds incremental protection against SAN failures but does not provide server HA

- Minimizes downtime for planned or unplanned storage outage
- PowerHA switch is required for planned or unplanned server outage with vary off / vary on of the IASP

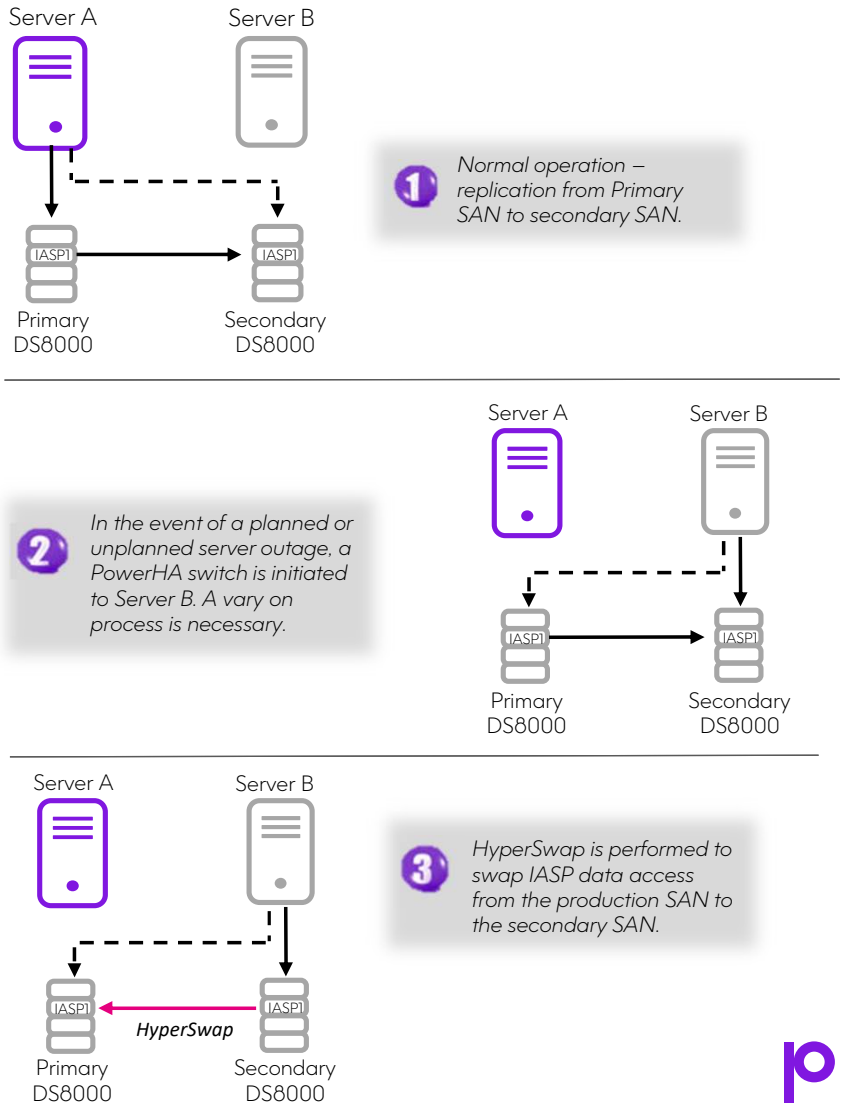

HyperSwap and Assure MIMIX

HyperSwap and Assure MIMIX can work together

- HyperSwap incrementally increases storage subsystem reliability over RAID and Mirrored DASD
- Assure MIMIX provides server HA

HyperSwap with PowerHA still needs Assure MIMIX

- HyperSwap swaps IASP or *SYSBAS to secondary SAN
- IBM PowerHA provides server switch
- Assure MIMIX for PowerHA extends and enhances HA
- Assure MIMIX is needed to add remote node

Assure MIMIX provides a solid solution

- HyperSwap requires second SAN to increase availability of already highly reliable DS8000
- HyperSwap coupled with PowerHA has all the shortcomings of PowerHA
- Assure MIMIX provides a cost-effective, complete server HA solution for SAN, internal or mixed storage

