



Technical Forum & Executive Briefing

17 al 21
Octubre
2011

Imagine **PODER** Imagine **CAPACIDAD**

PowerHA SystemsMirror for AIX Offerings and Strategy

Cesar Diniz Maciel
Executive IT Specialist – IBM Power Systems
Global Techline – Latin America
cmaciel@us.ibm.com



Strategy

- World domination!!!! 😊
- Seriously, to provide customers with mission-critical infrastructure in order to support the business environment

Agenda

- Offering strategy and capabilities/functionality
- Offering value to customers
- Future offering capabilities
- Customer success stories
- Resources for additional information

Cost of downtime

• Most users experience 66% more downtime than anticipated

• Average cost of mission critical application downtime \$10k/min

Application	Cost/Minute
Call Location	\$27,000
Number Portability	\$14,400
ERP	\$13,000
Supply Chain Management	\$11,000
Electronic Commerce	\$10,000
Internet Banking	\$7,000
Universal Personal Services	\$6,000
Customer Service Center	\$3,700
ATM/POS/EFT	\$3,500
Messaging	\$1,000

• 72% of mission critical applications experience nine hours of downtime per year: Most companies are under-investing in high availability technology"

Source: Standish Group Research Note 1998
Penny Wise & Pound Foolish

Resiliency Without Downtime -- PowerHA

High availability solutions for UNIX, and i clients



- Cluster management solutions for high availability
 - Can enable near continuous application service
 - Minimize impact of planned and unplanned outages
- PowerHA™ SystemMirror for AIX®
- PowerHA SystemMirror for i
- PowerHA pureScale

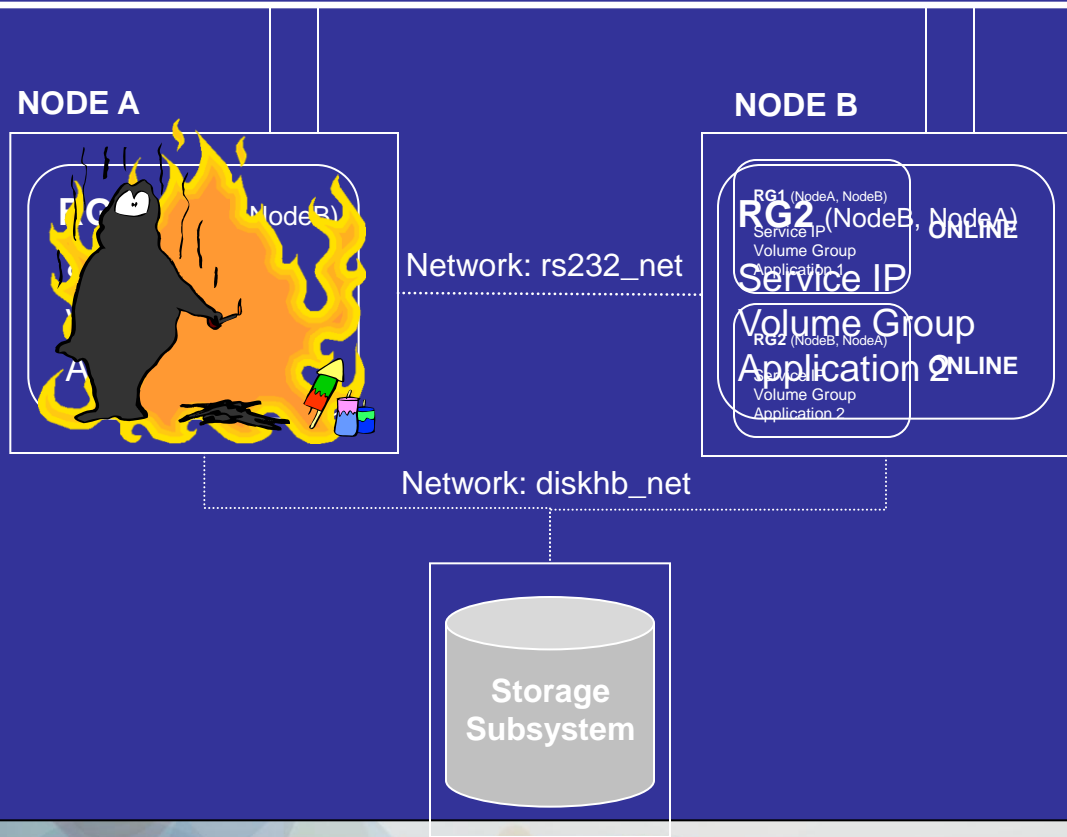
PowerHA SystemMirror Solution

- PowerHA SystemMirror
 - Active/Standby integrated AIX clustering
 - Storage based data resiliency (host based or storage server based)
 - Full system resiliency (application, data, networks etc)
 - Data center and multi-site resiliency

- Pre-reqs AIX
 - PowerHA SystemMirror 6.1, AIX 6.1 (AIX 5.3 currently supported)
 - PowerHA SystemMirror 7.1, AIX 6.1 and AIX 7.1

Cluster Name: Cluster 1

Network: net_ether0



Environment:

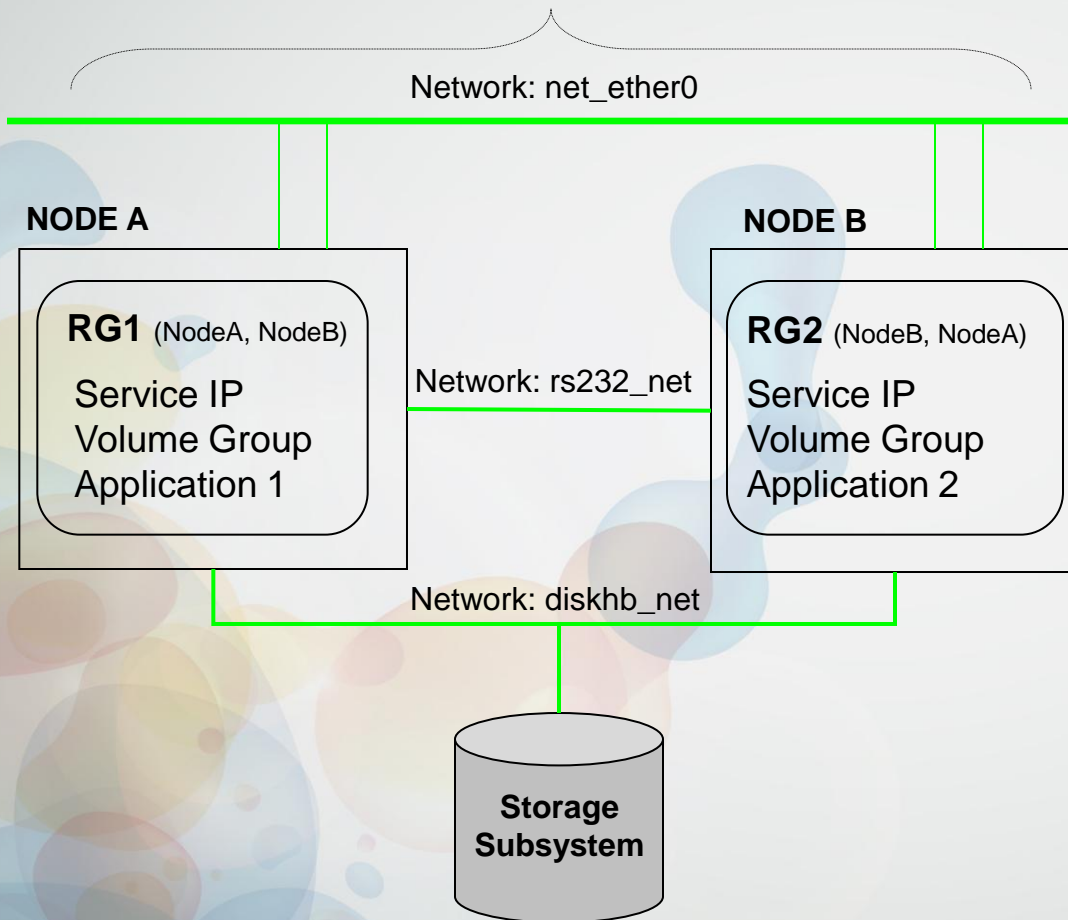
- Each Machine is running its own Production application
 - Node A fails to Node B
 - Node B fails to Node A

Fallover Behavior:

- On fallover the target machine will need enough CPU & memory resources in order to handle the load of both applications

PowerHA – How it protects you

Cluster Name: Cluster 1



Heartbeating Components

- IP Networks
- Serial Networks

Cluster Behavior

- Local Adapter Swaps
- Selective Fallover on IP Loss
- Selective Fallover on VG Loss

Monitoring

- Application Process Monitoring
- Application Custom Monitoring
- AIX Error Notification
- SMS Paging Service

PowerHA SystemMirror



- Complete IBM Power Systems integrated end to end solutions for HA DR
 - Focus: 24x 7 Application availability through planned or unplanned outage events
 - Deeply integrated extension of host operating system (AIX, IBM i)
 - Developed and supported by the Power Systems engineering development team
- HA Clustering technology
 - Clustering provides the applications a complete resiliency infrastructure
 - Cluster monitors and manages primary and secondary resources for HADR operations
- Storage based data resiliency
 - Data resiliency is an extension of the host system storage management architecture
 - Storage volumes are either switchable or mirrored between nodes in the cluster
 - Hardware based replication services for Multi-Site Operations
 - Host Based Replication (GLVM for AIX)
 - Storage Base Replication (Metro Mirror or Global Mirror)
- Overall solution characteristic
 - Automation, minimal IT operations involvement
 - Data between primary and secondary nodes always in sync always ready for a failover event

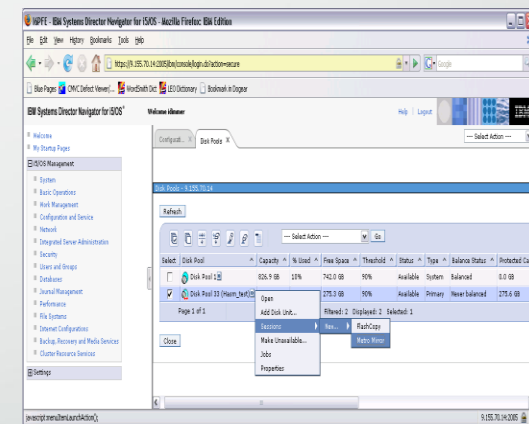
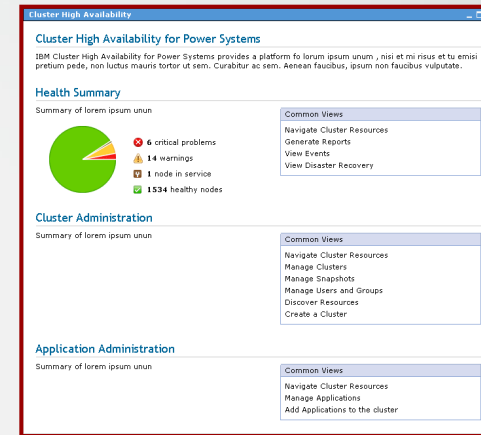
PowerHA SystemMirror Value Proposition

For:	<i>Operations executives such as CTO or IT Director</i>
Who Need:	<p><i>Near continuous application availability</i></p> <p><i>With minimal dependencies on IT administration</i></p> <p><i>With multi-site solutions for disaster recovery</i></p> <p><i>A host system OS technology based HA solution</i></p> <p><i>An IBM Power Systems supported solution</i></p>
IBM provides:	<p>HA Solutions for the data center:</p> <p><i>PowerHA SystemMirror Standard Edition</i></p> <p>HADR Solutions for multi-site configurations:</p> <p><i>PowerHA SystemMirror Enterprise Edition</i></p>
Unlike:	<p><i>Cross platform HADR offerings:</i></p> <p><i>PowerHA SystemMirror is an integrated extension of AIX</i></p> <ul style="list-style-type: none"> <i>•AIX kernel, LVM & storage management</i> <i>•Director based user interface</i> <i>•IBM Power Systems development, integration and support</i>
IBM's offering has/does:	<p><i>Automated HADR operations</i></p> <p><i>Options for both host based and storage based replication</i></p> <p><i>Integrated Cluster Aware host operating systems; AIX</i></p>

PowerHA SystemMirror Strategy

Simplification – Usability- Reliability

- **Solution Package Optimization**
 - Standard and Enterprise Editions
- **Deep Integration with AIX**
 - Cluster Aware AIX (7.1)
- **Ease of use**
 - Smart Defaults, Graphical Interfaces
- **Storage Integration & Disaster Recovery**
 - Extend PowerHA Storage Support
- **Middleware Integration**
 - Enhanced Application HA Management
- **Virtualization**
 - Exploit and align with Power Systems platform virtualization
- **Cluster wide resource management**
 - Single point of control




Client Need: High Availability – Data Center

- Requirements
 - Near continuous application service
 - Recovery Time Objective (RTO) measured in minutes
 - Recovery Point Objective (RPO) is Zero (no data loss)
 - Eliminate affects of planned outages
- Solution Strategy:
 - PowerHA SystemMirror **Standard Edition**



Simplifies server maintenance enabling Svenska Handelsbanken to switch between production and backup systems in 14 minutes. Provides ease of administration without requiring downtime from the bank's production server.



Client Need: Disaster Recovery at Remote Site

- Requirements
 - Recovery of IT operations at a remote location
 - Recovery Point Objective (RPO) near zero
 - Recovery Time Objective measured in hours (2 or 3 hours)
- Solution Strategy:
 - Data replication services (eg. Metro Mirror, Global Mirror, logical replication, srdf)
 - **Manually** recover data & operations from replicated data



Data Center



Remote Site

Regulated by the US Federal Reserve, CLS Bank as a designated critical component of the US Financial Industry, is required to withstand any regional disaster and be able to resume its service from another location in another region, within 2 hours.

Today's Client Need: Unified HA/DR Solution

- Requirements
 - HA/DR clustering for data center & multi-site operations
 - Automated role-swap operations
 - Readily demonstrate remote recoverability compliance
- Solution Strategy:
 - **unified clustering solution** for data center & multi-site resiliency
 - **PowerHA SystemMirror Enterprise Edition**



Remote Site



Data Center

SK Energy is South Korea's leading oil producer. To improve system resilience and introduce a disaster recovery plan, IBM recommended and implemented a twin data center approach. The Power Systems servers are configured as a cluster using IBM PowerHA, which provides high availability and rapid failover.

Business Value

- Keeping business applications available 24X7 with no data loss and minimal time lost in the event of an unplanned outage
- The IBM PowerHA™ solution enables IBM Power Systems™ clients to deploy solutions for high availability and disaster recovery.
- **Pain Points**
 - Need for continuous application and data availability in the data center
 - Need for a multi-site HA/DR solution supporting disaster recovery and compliance regulations
 - Need for solutions that are robust and yet require minimal IT operational involvement
 - Need an HA/DR solution that is offered by IBM with IBM service and support
- **Benefit:** PowerHA SystemMirror support for HA clusters based on hardware-based data replication ensures no data loss with synchronous replication and minimal data loss with asynchronous replication; easy to manage operations and easy on-demand failover to a backup facility. The PowerHA SystemMirror Enterprise Edition is the best price/value solution because it supports both multi-site and data center solutions with one package.
- **Benefit:** PowerHA SystemMirror is an extension of the AIX or IBM i operating system therefore the operating system specialists readily expand their skills to HA/DR operations.
- **Benefit:** PowerHA SystemMirror technology is developed by IBM Power Systems, is integrated and tested by Power Systems at the factory and is supported by IBM and IBM Power Systems experts

IBM Systems Storage™ & PowerHA SystemMirror

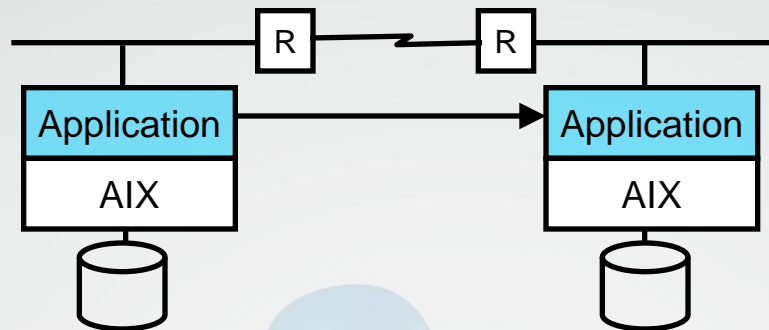


- Integrated HA/DR PowerHA cluster
- Data Center High Availability via shared/switchable storage volumes
- LVM mirroring for AIX environments
- Global Mirror provides Multi-Site Disaster Recovery
- FlashCopy provides fast, off-line backups of production applications and DB
- Reliable on demand role swap operations



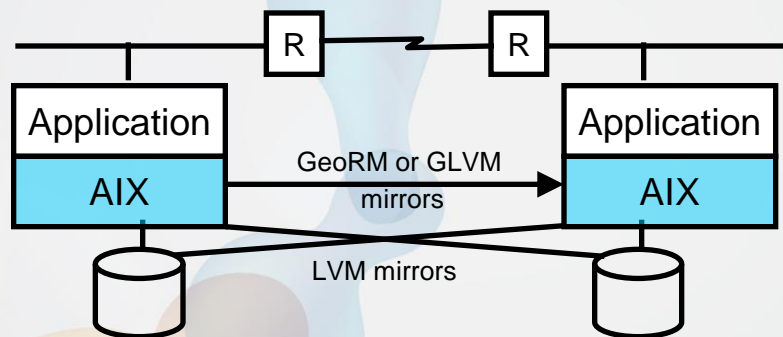
Power Systems site disaster mirroring options

Application replication




E.G., Lotus Notes replication, DB2 HADR or warm standby/log shipping

OS layer mirroring



E.G., LVM mirroring over FC, and GLVM over IP

Disk subsystem mirroring



E.G., MetroMirror

Differences between remote mirroring options

- Application data replication
 - ▶ Requires less bandwidth
 - ▶ Usually asynchronous mirroring only
 - ▶ More complex implementation
 - Implement it for each application
 - Only mirror *application data*
- AIX layer remote data mirroring
 - ▶ Relatively simple implementation
 - ▶ Mirror LVM *logical volumes*
 - Application independent
- Disk subsystem layer remote data mirroring
 - ▶ Simpler implementation
 - Mirror *logical disks*
 - Mirror disks for multiple heterogeneous systems
 - Application independent
 - Implement with application online!
 - ▶ Requires more bandwidth

Implementation
Complexity



Network
Bandwidth

Synchronous versus asynchronous solutions

■ Synchronous mirroring

- Writes complete at both sites prior to returning acknowledgement to application
- Only in-flight transactions lost due to a site disaster
- **Potentially significant performance impact**

■ Asynchronous mirroring

- Acknowledgement returned to application when local write completes and **before remote write completes**
 - In-flight transactions, **plus additional transactions will be lost**
 - Amount lost depends on how far behind the secondary site is
 - Limited performance impact: no additional write latency, but additional work to keep track of inter-site inconsistencies
- Some method must be used to keep track of what remote writes have not completed (always for asynchronous, and for synchronous when we can't talk to the secondary site)

Remote mirroring performance effects

- Local writes to disk cache take about 2 ms
- Remote writes often take 10X or longer
 - ▶ The issue is write latency, not necessarily bandwidth
 - ▶ Inadequate bandwidth means writes will queue increasing write times beyond the latency
 - ▶ Light thru fiber travels 1 km in 5 micro-seconds
 - ▶ Additional latencies for code stacks (routers, TCP/IP, etc.)
- Actual performance impact depends upon the application and hardware
 - ▶ Larger impact for transactional applications

PowerHA SystemMirror Editions

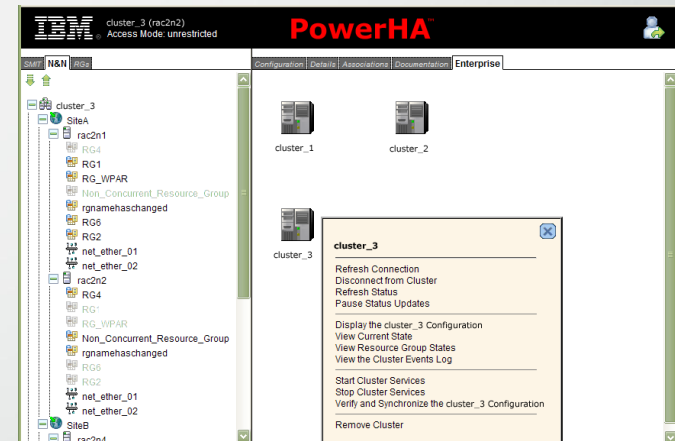
PowerHA SystemMirror for AIX Standard Edition

- Cluster management for the **data center**
 - Monitors, detects and reacts to events
 - Establishes a heartbeat between the systems
 - Enables automatic switch-over
- IBM shared storage clustering
 - Can enable near-continuous application service
 - Minimize impact of planned & unplanned outages
 - Ease of use for HA operations



PowerHA SystemMirror for AIX Enterprise Edition

- Cluster management for the **Enterprise**
 - Multi-site cluster management
 - Includes the Standard Edition function
 - Optimized for IBM storage
 - Supports EMC SRDF
 - [Hitachi Truecopy](#)



PowerHA SystemMirror for AIX Editions

- Editions optimized for software value capture
- Standard Edition targeted at datacenter HA
- Enterprise Edition targeted at multi-site HA/DR
- Tiered pricing structure Small/Med/Large
- Expanded storage support

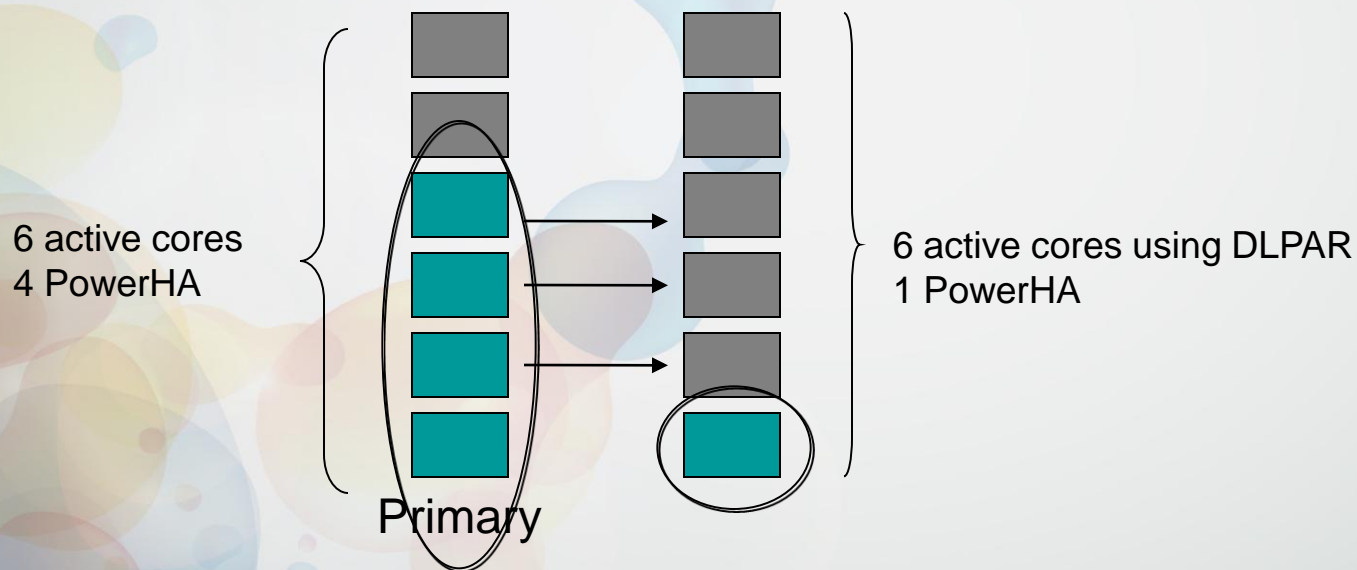
	Standard Edition	Enterprise Edition
Centralized Management CSPOC	✓	✓
Cluster resource management	✓	✓
Shared Storage management	✓	✓
Cluster verification framework	✓	✓
Integrated disk heartbeat	✓	✓
SMIT management interfaces	✓	✓
AIX event/error management	✓	✓
Integrated heartbeat	✓	✓
PowerHA DLPAR HA management	✓	✓
Smart Assists	✓	✓
Multi Site HA Management		✓
PowerHA GLVM async mode		✓
IBM Metro Mirror support		✓
IBM Global Mirror support DS8700		✓
EMC SRDF sync/async		✓
Hitachi Truecopy		✓

PowerHA SystemMirror Licensing Considerations

- Licensed/priced per processor core configured in an HA partition
- Sub capacity: license only the processor cores involved in the HA configuration
 - Example: 16 cores active while 8 are doing HA for SAP; you purchase 8 PowerHA entitlements (not 16 !)
 - Isolate HA application to a partition and license all of the cores in that partition
- Target System have one core isolated to a partition and upon failover use DLPAR to collapse primary partition and expand target partition thus shifting entitlements to target cores

PowerHA SystemMirror on AIX in a HA/DR Configuration

- PowerHA SystemMirror on AIX
 - One PowerHA LPP per core being used for HA on the primary
 - You don't need to license all of the cores, only the cores being used for HA/DR support
 - Only one PowerHA LPP on the back up server (no CBU required)
 - Isolate HA environment to its on LPAR and license all the processor cores in the LPAR
 - Use DLPAR to contract primary partition and expand target partition



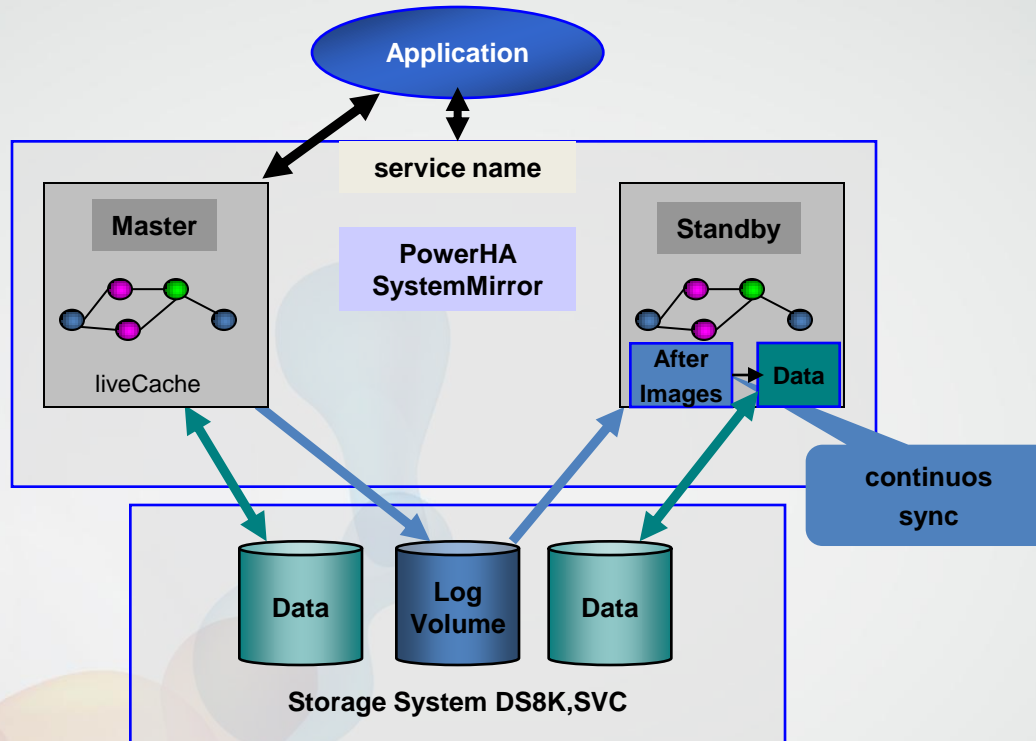
Resiliency without Downtime

PowerHA SystemMirror 7.1 Standard Edition

- **PowerHA SystemMirror 7.1 integration with Cluster Aware AIX**
 - Kernel based monitoring
 - Cluster wide event management
 - Cluster wide AIX command interface
 - Cluster wide device naming
 - Cluster wide storage device management
- **New Director-based GUI interface**
 - Intuitive interface for PowerHA SystemMirror deployment and management
- **New Smart Assist – Discover, Configure and Manage HA for**
 - SAP, DB2, Oracle, WebSphere
 - PowerHA SAP LiveCache Hot Standby Solution
 - MQSeries Smart Assist Agent
 - LDAP Smart Assist Enhancement
- **Cluster wide Security Administration**
 - PowerHA Federated security,
- **Complexity reduction**
 - Director Module Enhancements for PowerHA Systems Mirror
- **Exploitation of AIX Clustering facility**
 - Enterprise content manager,
 - Tivoli storage manager, LDAP, IBM HTTP, printers
- **PowerHA SystemMirror 6.1 EE SP7**
 - Expand Storage Integration & Disaster Recovery
 - Exploit IBM XIV Disaster recovery
 - Kernel level event management



SAP LiveCache Hot Standby Solution



Eliminates Downtime - **Failover in minutes (MaxDB)**

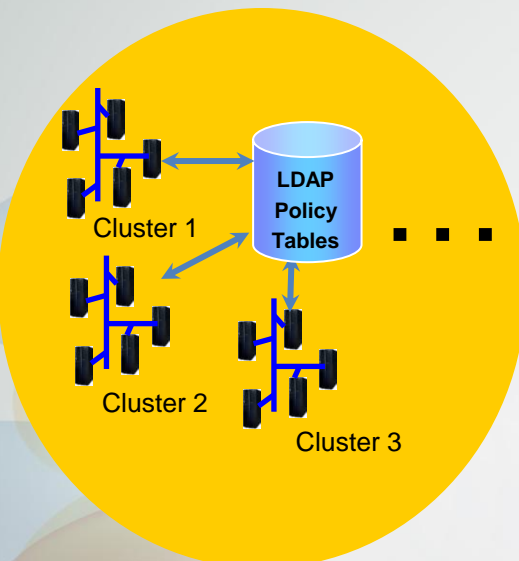
Easy to Manage - **Director or SMIT based interface**

Easy to Deploy - **SAP LiveCache Hot Standby Setup Wizard**

Included in PowerHA 7.1.1 Standard Edition



PowerHA Federated Security



The screenshot shows the IBM Systems Director web interface. The main window displays 'Cluster and Resource Group Management' for a cluster named 'SAP_Cluster'. The interface is divided into several sections:

- Left Panel:** A navigation menu with categories like Automation, Availability, Inventory, Release Manager, Security, System Configural, System Status and Task Management.
- Resource Groups Table:** A table listing resource groups and their HA Status.

Select	Name	HA Status
<input checked="" type="radio"/>	SAP_Cluster	OK
<input type="radio"/>	SAP_active	OK
<input type="radio"/>	DB_active	OK
<input type="radio"/>	SAP_DB_standby	OK
<input type="radio"/>	WAS_Cluster	OK
<input type="radio"/>	WAS_Node1	OK
<input type="radio"/>	WAS_Node2	OK
<input type="radio"/>	WAS_Node3	OK
<input type="radio"/>	WAS_Node4	OK
- Federated Security for SAP_Cluster Table:** A table showing user groups and their attributes.

Select	Name	ID#	User Role	Registry	Keystore	Home Directory
<input type="checkbox"/>	SAP_Group1					
<input type="checkbox"/>	susantj_1010	10	Admin	LDAP	LDAP	/var/spool/lp
<input type="checkbox"/>	grehint_1014	14	User	LDAP	LDAP	/var/spool/uicppublic
<input type="checkbox"/>	mark_1006	6	Monitor	LDAP	LDAP	/usr/sbin/snapp
<input type="checkbox"/>	SAP_Group2					
<input type="checkbox"/>	gabialex_1013	13		Local	LDAP	/var/adm/invscout
<input type="checkbox"/>	todde_1017	17		Local	LDAP	/var/adm/pconsole
<input type="checkbox"/>	poolr_1002	12		Local	LDAP	/home/guest
<input type="checkbox"/>	SAP_Group3					
<input type="checkbox"/>	susantj_1010	10	Admin	LDAP		/var/spool/lp
<input type="checkbox"/>	grehint_1014	14	User	LDAP		/var/spool/uicppublic
<input type="checkbox"/>	mark_1006	6	Monitor	LDAP		/usr/sbin/snapp
<input type="checkbox"/>	rajibadh_1026	26	Monitor	LDAP		/var/spool/lp
<input type="checkbox"/>	kenhigh_1023	23	Log Viewer	LDAP		/usr/bin/uucp

Centralized Administration,
System security, encrypted file system, role based administration..

LDAP Servers

IBM Tivoli Directory Services

Windows Active directory

Included in the 7.1.1 Standard Edition



PowerHA SystemMirror 6.1 Enterprise Edition

Storage Options for your Enterprise

	Sync mode	Async mode	Active Active	Consistency groups
GLVM	Yes	Yes	✓	✓
XIV	Metro Mirror	Global Mirror	✓	✓
Storwize® V7000	Metro Mirror	Global Mirror	✓	✓
SVC	Metro Mirror	Global Mirror	✓	✓
DS8800	Metro Mirror	Global Mirror	✓	✓
EMC	SRDF	SRDF	✓	✓
Hitachi	TrueCopy Universal Replicator	TrueCopy Universal Replicator	✓	✓
Hitachi VSP	Truecopy Universal Replicator	Truecopy Universal Replicator	✓	✓
HP Storage Works P9500	Continuous Access Software replication	Continuous Access Software replication	✓	✓



Solutions for disaster recovery
 Major vendors, popular options
 IBM storage options that span the market



Additional Resources PowerHA for AIX

- PowerHA Website
 - www.ibm.com/systems/power/software/availability/
- Availability Factory
 - Contact your IBM representative or an IBM Business Partner and they will contact us via e-mail (hacoc@us.ibm.com) to learn more.
- IBM Technology Service Offering for PowerHA SystemMirror XD deployment
 - <http://www-935.ibm.com/services/us/index.wss/offering/its/a1000032>
- **Redbooks**
 - PowerHA on AIX (SG24-7739-00)
 - SG24-7841 IBM PowerHA Enterprise Edition Implementation Guide
 - SG24- 7845 IBM PowerHA v.71 Education: PowerHA for AIX Implementation, Config and Administration AN410
 - Go to IBM.com/services/learning (search for AN410 or PowerHA)
- Education: Lab Services AN44 Extended Distance and Disaster Recovery
 - http://www-304.ibm.com/jct03001c/services/learning/ites.wss/us/en?pageType=course_list&subChapter=194&subChapterInd=S®ion=us&subChapterName=AIX+high+availability&country=us
 - <http://lpar.co.uk/HA/home.html>
- GLVM white paper
 - www.ibm.com/systems/resources/systems_p_os_aix_whitepapers_pdf_aix_glvm.pdf
- IBM storage virtualization offerings
 - www.ibm.com/systems/storage/virtualization
- SAP consulting services for PowerHA and PowerVM
 - gehenni@us.ibm.com
 - sbranden@us.ibm.com
- Wiki
 - <http://www.ibm.com/developerworks/wikis/display/WikiPtype/High%20Availability>