



PowerHA SystemMirror

Cesar Diniz Maciel
Executive IT Specialist
IBM Global Techline
cmaciel@us.ibm.com



15 de Mayo, 2013. IBM Argentina

Agenda

- § High availability and DR concepts
- § PowerHA product offering and versions
- § New version 7 Enterprise Edition





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



Business Objectives for Disaster Recovery

Recovery Time Objective

§How long can you afford to be without your systems?

Recovery Point Objective

§When it is recovered, how much data can you afford to recreate?

§No data loss requires a **synchronous** solution

§**Asynchronous** solutions usually lose some data

Network Recovery Objective

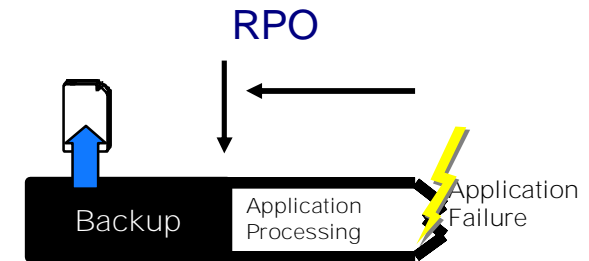
§How long to switch over network?

§Determine telecom line costs at various bandwidths

What is your cost / recovery time curve:

§If I spend a little more, how much *faster* is Disaster Recovery?

§If I spend a little less, how much *slower* is Disaster Recovery?



Determining the cost vs. RTO recovery curve is the key to selecting proper solution(s)


High Availability Clustering Solutions

§ High Availability (HA) Management


- Business continuity by workload failovers
- Provides for redundant operating environment
- Automated workload bring ups
- Environment specific outage actions
- Planned and unplanned outage management

§ Disaster Recovery (DR) Management


- Geographically dispersed locations
- Host based or Storage based replication
- Simplified compliance testing



Data Center Clustering



Multi-Site Clustering





PowerHA SystemMirror Value Proposition

Reducing both Planned and Unplanned downtime

<p><u>Unplanned Outages:</u></p> <p>§ System Failures</p> <ul style="list-style-type: none">– Hardware– Operating System Crash– Power Loss– User Error <p>§ Component Failures</p> <ul style="list-style-type: none">– NIC– SCSI/SAN Adapter– Network Hub/Switch– SAN Switch– Disk Failure (both O/S and application data)– PowerHA SystemMirror 7.1 is specially resilient to loss of storage	<p><u>Planned Outages:</u></p> <p>§ Maintenance</p> <ul style="list-style-type: none">– System Hardware Change/Upgrade– OS & Application Upgrades & Fixes– Non-disruptive cluster SW updates <p>§ Non-Disruptive System Maintenance</p> <ul style="list-style-type: none">– When Combined with Live Partition Mobility <p>§ Testing</p> <ul style="list-style-type: none">– Applied Fixes– Failure scenarios for HA & DR
--	---

Edison Group Whitepaper on deep value of PowerHA integration with AIX



PowerHA SystemMirror Editions

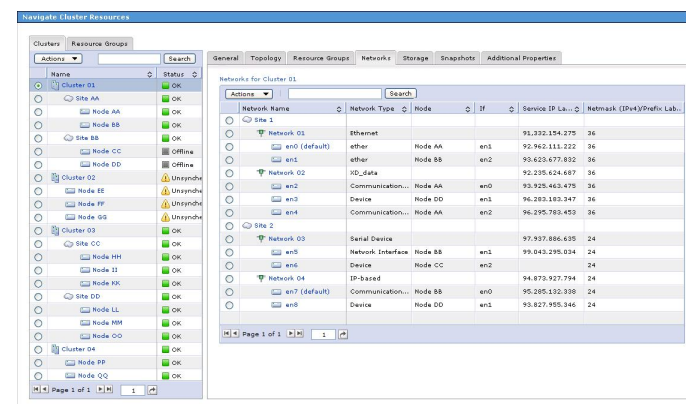
PowerHA SystemMirror Standard Edition

- § 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



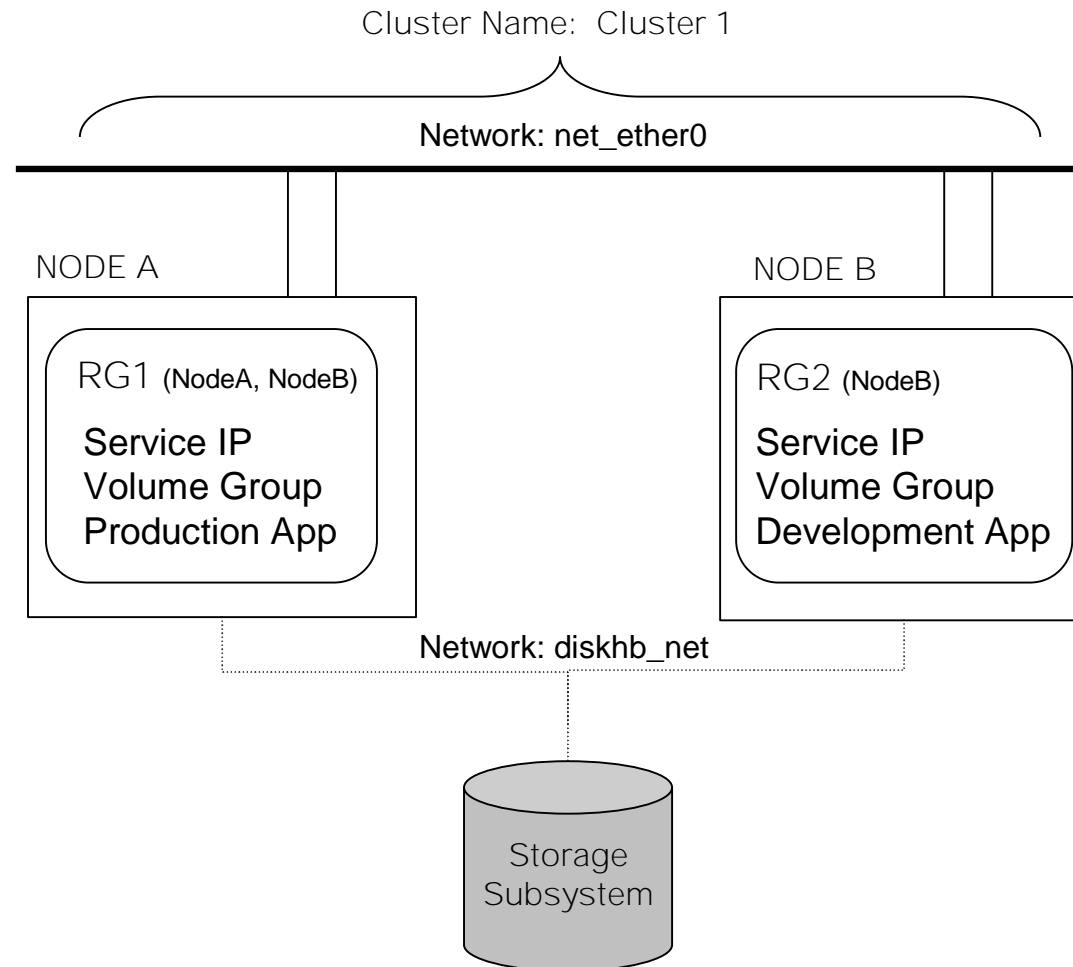
PowerHA SystemMirror Enterprise Edition

- § Cluster management for the Enterprise
 - Multi-site cluster management
 - Includes the Standard Edition function





IBM PowerHA SystemMirror for AIX Sample Cluster Configuration



Topology Components:

- Cluster name
- Node Names
- IP Network
 - Interfaces
- Serial Network

Types of Resources

- Service IP
- Volume Group/s
- Application Server

Resource Components:

- Resource Group/s
 - Policies
 - startup
 - fallover
 - fallback
- Dependencies
 - Parent / Child
 - Location



PowerHA SystemMirror 7.1 Standard Edition

Simpler to deploy and easier to manage with IBM Systems Director, intuitive interfaces, cluster and resource group wizards, management dashboards and Smart Assists for SAP and other popular applications

Minimize IT operations with cluster aware AIX; cluster wide AIX commands, kernel based event management, device naming, central repository and multi-channel communications

Robust cluster integrity with disk fencing and multi-channel heart beat which automatically uses available I/O including SAN.

Complete end to end failover automation with policy driven resource group relationship sequencing





PowerHA SystemMirror for AIX 6.1 Versus 7.1

PowerHA SystemMirror	6.1	7.1	PowerHA 7.1 Benefit
IBM Director based graphical user interface	NA	ü	Ease of Use
Cluster Aware AIX (CAA)	RSCT	ü	Reliability
Triple redundant heartbeat (multicast)	NA	ü	Effectively eliminates partitioning
SAN based communications	NA	ü	Additional cluster communication path
Stretched cluster (shared repository)	NA	ü	Two-Site multicast HA/DR shared network
Cross Site Mirroring (single site stretch cluster)	NA	ü	LVM mirroring with CAA
Linked clusters (separate repositories)	NA	ü	Two-Site HA/DR separate networks
HyperSwap with DS8800	NA	ü	Two-Site continuously available storage
Multi-Site set up wizard	NA	ü	Speeds up implementation
Tie Breaker	NA	ü	Can eliminate split-site scenarios
Federated Security	NA	ü	Cluster wide security management
Live Cache SAP hot standby	NA	ü	Fast failover for APO SCM



PowerHA SystemMirror 7.1 Enterprise Edition

Simpler to deploy and easier to manage **multi-site configurations** with IBM Systems Director, intuitive interfaces, multi-site install wizard

Stretched Cluster; **Cluster wide AIX commands**, kernel based event management
single repository multicast communications

Linked Clustering; **cluster wide AIX commands**, kernel based event management,
linked clusters with unicast communications & dual repositories

HyperSwap for continuously available storage in two-site topologies

Cluster Split/Merge **technology** for managing split-site policy scenarios




Announce Date: Oct 3 2012

GA Date: Nov 09 2012



© 2013 IBM Corporation

PowerHA 7.1 Two-Site Solutions



Two-Site Stretched Cluster



Two-Site Linked Cluster

Two Sites	Stretched Cluster	Linked Clusters
Inter Site Communication	Multicast	Unicast
Repository Disk	Shared	Separate
Cluster Communication	<ul style="list-style-type: none"> ü Networks ü SAN ü Disk 	<ul style="list-style-type: none"> ü Networks ü SAN *
Cross Site LVM Mirroring	ü	ü
Storage based replication	ü	ü
HyperSwap	ü	ü
Multi Site Concurrent RG with / HyperSwap (active-active) *	ü	NA


PowerHA V7.1	Standard	Enterprise
Multi Site Definition	ü	ü
§ Site Service IP		
§ Site Policies		
Stretched Cluster	ü	ü
Linked Clusters	ü	ü
HADR with Storage Replication Management	NA	ü
HyperSwap	NA	ü

* Future capability




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

PowerHA for AIX Approach For Data Synchronization


- § PowerHA SystemMirror Enterprise Edition
- Options for managed synchronous replication
 - Metro Mirror - **SVC and DS8000 based synchronous replication solution**
 - GLVM - **AIX based mirroring solution over IP**
 - EMC SRDF – **EMC's data replication**
 - Hitachi Remote Copy – **Hitachi's data replication**
 - SVC Global Mirror – **Integration with IBM Storage Virtualization Controller**






PowerHA Enterprise Edition HyperSwap

- § Multi-Site PowerHA cluster with continuous storage availability
- § Non-disruptive - applications keep running in the event of a storage outage
- § Storage maintenance without downtime
- § Storage migration without downtime



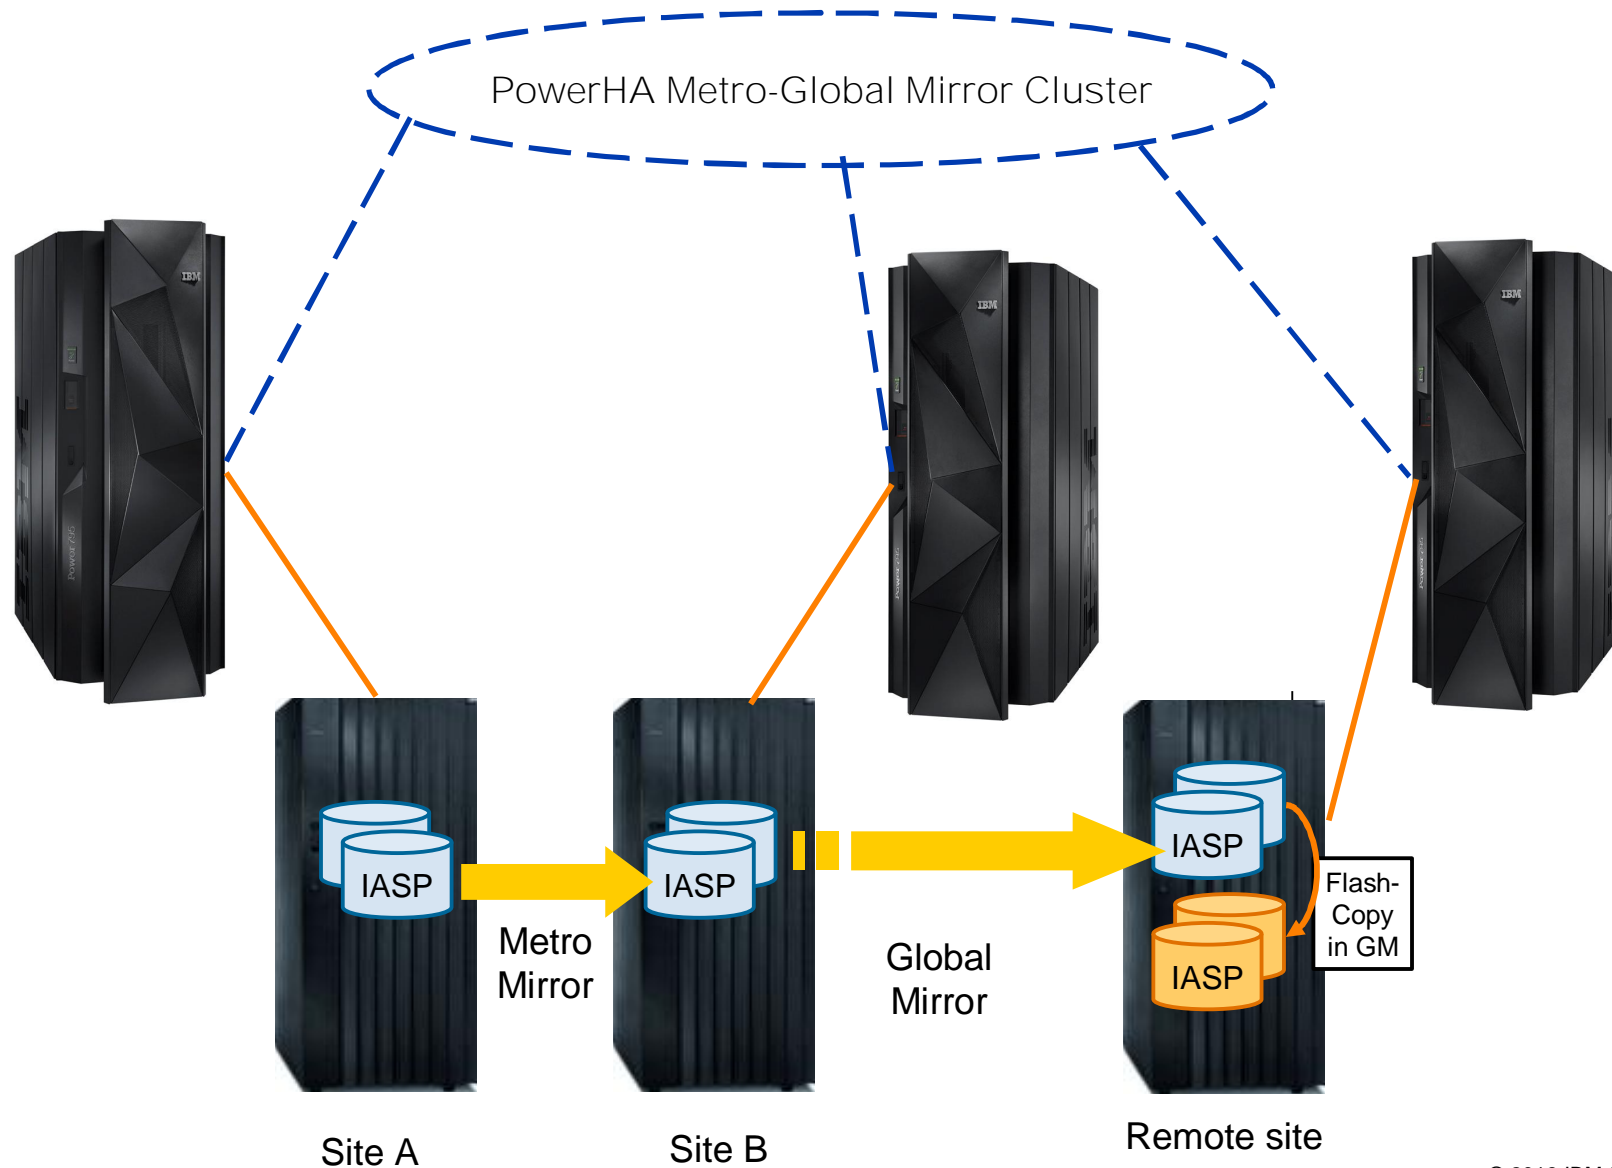
PowerHA for IBM i - Geomirror



- § Internal disk is not switchable (LUNs required), you use geomirror with internal disk configurations (or for that matter any disk configurations)
- § Use Geomirror for configs with IASPs under 5 Tbytes (a suggested guideline ... key point is proper sizing for bandwidth when considering a complete re-sync after an outage)
 - Sync mode included with Standard Edition (~ 40 KM or less)
 - Async mode requires Enterprise Edition
 - Storage agnostic (typically internal disk)




IBM I (MGM) Metro Mirror Global Mirror Three Site Cluster





PowerHA: Resource Optimized High Availability

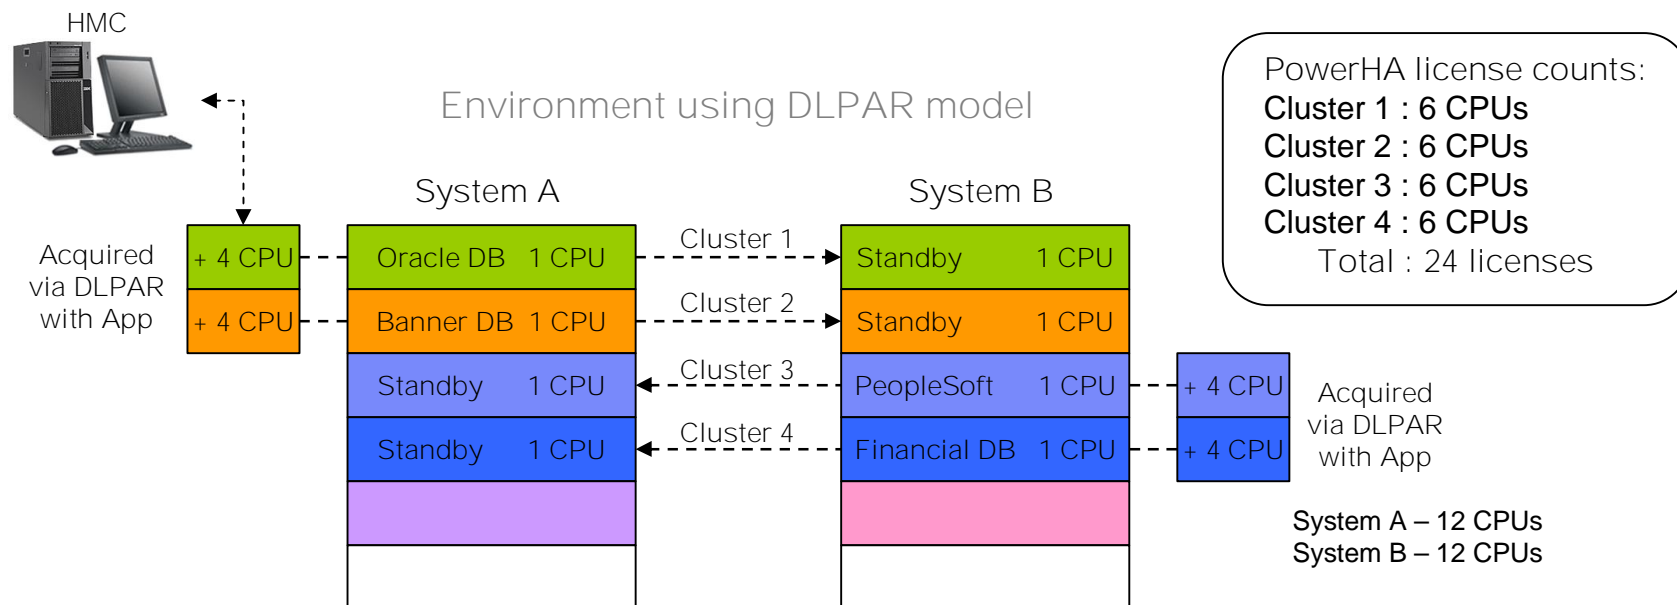
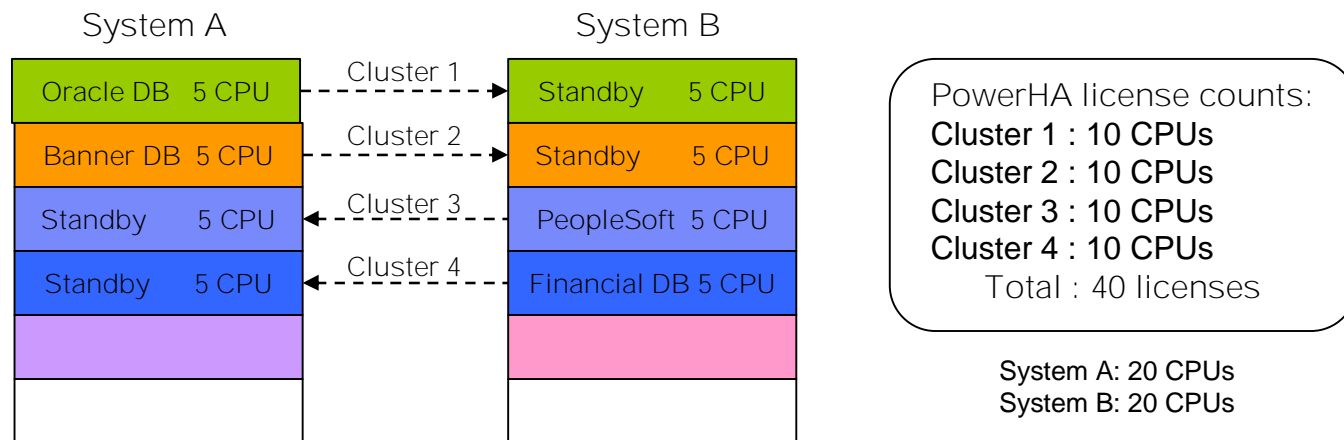


- § PowerHA provides for LPAR resource management during failovers
- § Can readjust resources from the free pool
- § Can enable CPU and memory as needed using CoD



PowerHA SystemMirror: DLPAR Value Proposition

Environment using dedicated CPU model (No DLPAR)





PowerHA SystemMirror 7.1 Enterprise Edition

Storage Options for your Enterprise

	Sync mode	Async mode	Active Active	Consistency groups
GLVM	Yes	Yes	ü	ü
XIV	Metro Mirror	Global Mirror	ü	ü
Storwize® V7000	Metro Mirror	Global Mirror	ü	ü
SVC	Metro Mirror	Global Mirror	ü	ü
DS8800	Metro Mirror	Global Mirror	ü	ü
EMC	SRDF	SRDF	ü	ü
Hitachi	TrueCopy Universal Replicator	TrueCopy Universal Replicator	ü	ü
HP Storage Works P9500	Truecopy Universal Replicator	Truecopy Universal Replicator	ü	ü
HP XP12000, XP24000	Continuous Access Software replication	Continuous Access Software replication	ü	ü

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





PowerHA SystemMirror 7.1 for AIX

- § AIX embedded clustering technology
- § *AIX Cluster Aware, Kernel based event management*
- § PowerHA skills are an extension AIX skills
- § WW service and support

PowerHA SystemMirror V7	Standard Edition	Enterprise Edition
AIX IT operations skill base	ü	ü
AIX Technical skill support base	ü	ü
AIX Shared Storage management	ü	ü
AIX Volume manager/Mirroring	ü	ü
AIX File System	ü	ü
AIX SMIT Management Tools	ü	ü
AIX DLPAR HA/DR management	ü	ü
IBM Director based graphical interface	ü	ü
AIX Cluster Aware	ü	ü
Kernel based event management	ü	ü
Federated Security	ü	ü
Stretch cluster (shared repository)	ü	ü
Cross Site Mirroring	ü	ü
Triple redundant multi-cast heartbeat	ü	
SAN based communications	ü	
Live Cache SAP hot standby	ü	
Linked clusters (separate repositories)	ü	ü
HyperSwap		ü
Multi-Site set up wizard		ü
Split/Merge management options		ü



PowerHA 7.1 Standard Edition Smart Assists

- § Simplified application availability
 - Included in the Standard Edition

- § Advanced PowerHA agent
 - *Discover, Configure, and Manage*

- § Support to define and deploy HA policies
 - *Exploit advanced relationships*
 - *Flexible failover policies and resource order policies*
 - *Out of the box start, stop and health monitoring for common workloads*

1	Websphere
2	IBM DB2
3	Oracle Database
4	Oracle Application Server
5	IBM HTTP Server
6	Tivoli Directory Server
7	SAP
8	FileNet
9	Tivoli Storage Manager (3 Smart Assists)
10	Lotus Domino Server
11	MAXDB
12	Websphere MQ
13	AIX Print Server
14	AIX DHCP
15	AIX DNS



Resources & Assets: Disaster Recovery with PowerHA, Storage + SSW

- § High Availability with IBM PowerHA Marketing Page
- § <http://www-03.ibm.com/systems/power/software/availability/index.html>
- § IBM PowerHA SystemMirror for AIX
- § <http://www-03.ibm.com/systems/power/software/availability/aix/index.html>
- § HA & HAMCP Wiki – PowerHA for AIX
- § <http://w3.tap.ibm.com/w3ki03/display/hacmp/PowerHA+for+AIX>
- § IBM PowerHA SystemMirror 7.1 for AIX Redbook
- § Sale eLearning:
- § [Unified HA and Disaster Recovery Solution - IBM Power, Storage and Systems Software A Sellers Guide to Opportunity](#)
- § Disaster Recovery Whiteboard
- § <http://swglearning.raleigh.ibm.com/whiteboard/>
- § Whitepapers and Red papers
- § [New Clipper Group White Paper - A Win-Win Strategy for Mitigating Risks and Being Prepared](#)
- § [Disasters Remind Us Why We Need a High-Availability Data Center](#)
- § [Edison Group - IBM PowerHA SystemMirror 7.1 — The Value of Deep Integration](#)
- § [IBM Storage Infrastructure for Business Continuity](#)
- § Risk Self Assessment:
- § www.ibm.com/smarterplanet/us/en/business_resilience_management/overview/index.html?re=2brf24
- § Business Resilience QuickChecks:
- § <https://www-309.ibm.com/technologyconnect/tgcm/TGCMServlet.wss?alias=quickcheck&linkid=1Q0000>
- § References
- § [Disaster Recovery References](#) *(Search PowerHA + Disaster Recovery)*
- § HA CoC Services
- § [HA CoC Offerings \(hacoc@us.ibm.com\)](#) <http://www-03.ibm.com/systems/services/labservices/solutions/hacoc.html>



Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

© 2013 IBM Corporation



Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, ChipHopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

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.

AltiVec is a trademark of Freescale Semiconductor, Inc.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010

© 2013 IBM Corporation