

Session: Singapore / ASEAN Power Meetup HA/DR Solutions on IBM Power Systems

Antony Steel

Thanks to:

Steven Finnes

Ravi Shankar

Paul Moyer

Dino and the Redbook team



#include <std_disclaimer.h>

These notes have been prepared by an Australian, so beware of unusual spelling and pronunciation.

All comments regarding futures are probably nothing more than the imagination of the speaker and are IBM Confidential till after GA.



- Introduction to Availability and Disaster Recovery
- Options for IBM Power
 - Live Partition Mobility
 - Simplified Remote Restart
 - Virtual Machine (VM) Recovery Manager HA
 - VM Recovery Manager DR
 - Active / Active options
 - PowerHA System Mirror Standard Edition
 - PowerHA System Mirror Enterprise Edition
 - Geographic Logical Volume Manager
- Review of features and plans
- Comparison of options


Please

Feel free to ask questions, the chat window is being monitored
Stay on mute, unless asking a question

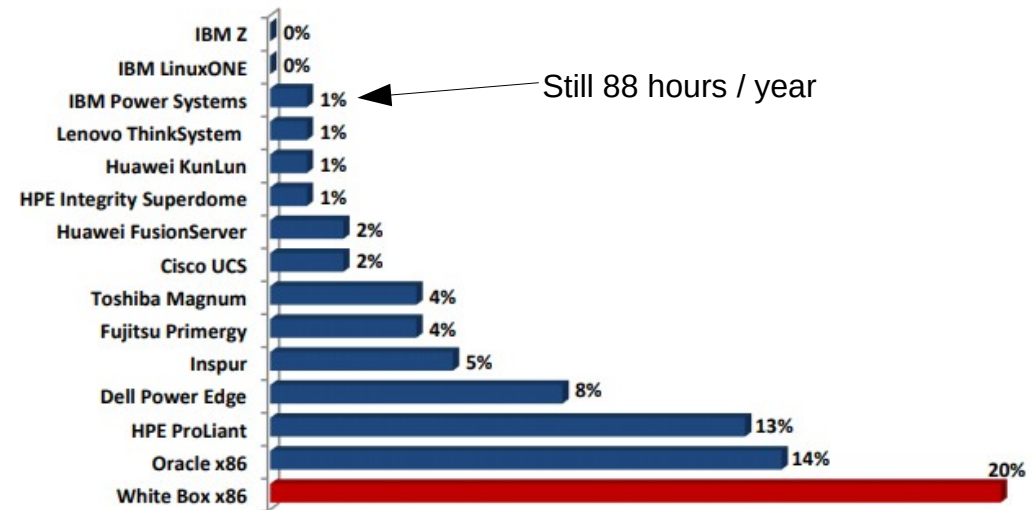
Why is HA / DR an issue

- Two-Thirds of Businesses claim that increased workloads negatively impact reliability. A 66% majority of survey respondents said that increased workloads are negatively impacting reliability compared with just 18% that said their companies hadn't experienced a decline in reliability due to an increase in workloads.
- Majority of Firms Fail to Calculate Downtime Costs. A 55% majority of companies fail to calculate the hourly cost of downtime versus 39% of firms that do track the monetary impact of server and software outages.

Factors that have the most negative impact on reliability



Unplanned downtime > 4 hours by Server platform



Source: ITIC 2019 Global Server, Hardware, OS Reliability Survey (Mid-year update)

<https://www.ibm.com/downloads/cas/DV0XZV6R>

Selecting a solution or solutions


- Review your environments and the Recovery Time Objectives (RTO) / Recovery Point Objectives (RPO) with your “Clients” (End Users, Application Owners....)
- Group your environments by platform and RTO/RPO
- Calculate your recovery time from an unplanned outage and estimate the potential data loss and associated cost
- Match the availability solutions with the requirements of each of your environments and balance cost / benefit
- You now have a budget to start to improving your HA / DR
- Don't forget about:
 - Schedule maintenance windows well into the future
 - Building a test plan and schedule testing (and follow through with it)
 - Is your documentation available; up to date; simple to follow; ever used?
 - Consider elements outside just the primary infrastructure – availability of staff; bandwidth; cross-site operations; impact of moving to DR; availability of supporting infrastructure....
 - Get agreement / buy in from the Application Owners.
 - Regularly review solutions and options (up to date? Reduced complexity?....)
 - After any outage review cause and update plan, training, infrastructure as required.
 - Keep up to date with patches and features

System maintenance

- While availability options are designed for unplanned downtime, they also assist with planned downtime
 - Maintenance
 - Testing
 - Upgrades
 - Development
 - Availability options are also a management tool that can help mask downtime
 - Understand the possible unplanned downtime events that may happen
 - User Error
 - Application Failure
 - Component Failure
 - Operating System Failure
 - Environmental Disasters
- .. and build your planning around that

When are availability solutions a good fit

- HA / DR solutions are a good fit when:
 - Aim is to reduce downtime to close to zero (not fault tolerance)
 - Solution may address planned or unplanned down time
 - Solution need not be fault tolerant but should be fault resistant
 - Solution should eliminate all single points of failure (SPOF)
- HA / DR solutions are not an answer if:
 - Cannot afford any downtime – life critical systems - Need a fault tolerant solution
 - Environment is not secure
 - Many users with root access
 - The environment is not stable
 - Change management is not respected
 - You do not have trained administrators
 - Procedures are not well documented
 - Environment is prone to user fiddle factor
 - Applications cannot be controlled
 - Scripts cannot be used to start/stop and recover applications
In particular if you cannot script the recovery of your application after an unexpected halt



HA and DR options


- Starting criteria
 - Redundancy (match depth to cost and application “value”)
 - Scripted application control and monitoring (see PowerHA Smart Assists)
- LPM
 - Not designed for unplanned outages, but to work around planned outages
- Simplified Remote Restart (SRR)
 - Correctly configured LPARs can be restarted on another server after a host “failure”
 - As just a “reboot” need application to be able to start after a hard reset
- PowerVC managed SRR
 - Automate SRR through PowerVC (includes automation and placement)
- VM Recovery Manager HA
 - Monitors host, if host fails, restarts VM on another host
 - Monitors applications / OS (registered hosts with agent installed)
 - For planned outage can use LPM to vacate host, moving VMs to other hosts in group
- VM Recovery Manager DR
 - Extends VM Recovery Manager to a DR / Backup site
 - Using scripts for network changes at DR site

HA and DR options (Cont)

- Active / Active Database options
 - PowerHA pureScale (concurrent DB2) and DB2 Mirror in IBM i
 - PowerHA “integration” with Spectrum Scale (GPFS)
 - Oracle RAC
- PowerHA System Mirror for AIX, IBMi and Linux
 - 3 separate products that monitor the infrastructure (Server; Storage; Network; Application) and move the Application to ensure availability to the End Users.
 - A large collection of configuration checking, testing and management tools
 - AIX (Linux soon) Smart Assists (sample scripts for specific applications)
 - Easy GUI for monitoring and managing all the clusters
- PowerHA System Mirror Enterprise Edition (AIX and IBMi)
 - DR and storage replication added
- Geographic Logical Volume Manager (GLVM)
 - Part of AIX, replication of storage over IP, can be managed by PowerHA EE

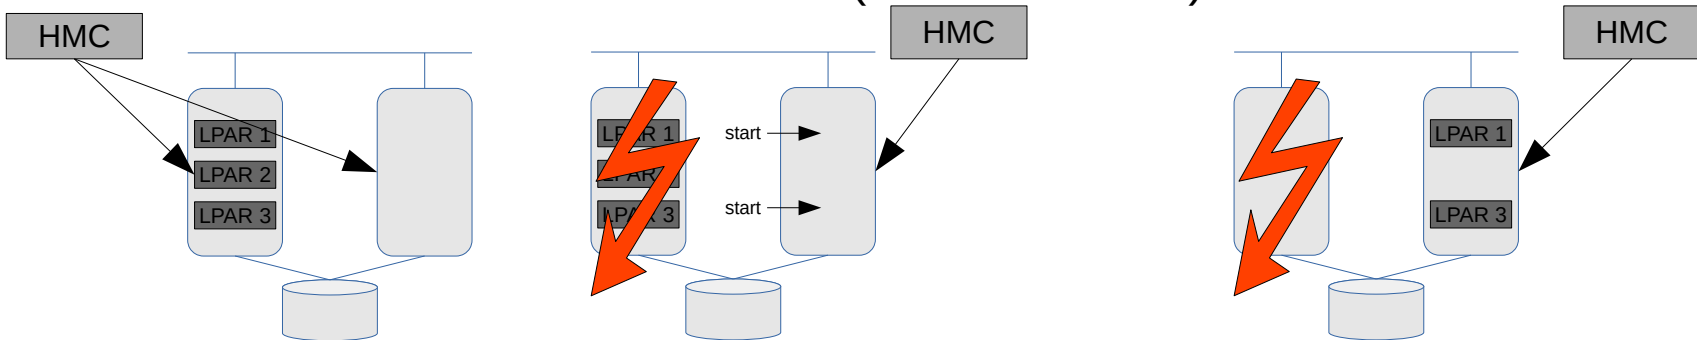
Live Partition Mobility

- Not designed for unplanned outages, but to work around planned outages
 - Host maintenance; workload balancing; some infrastructure maintenance; virtualisation layer maintenance
- Storage and network layer virtualised; consistency across hosts
- OS independent, no outage

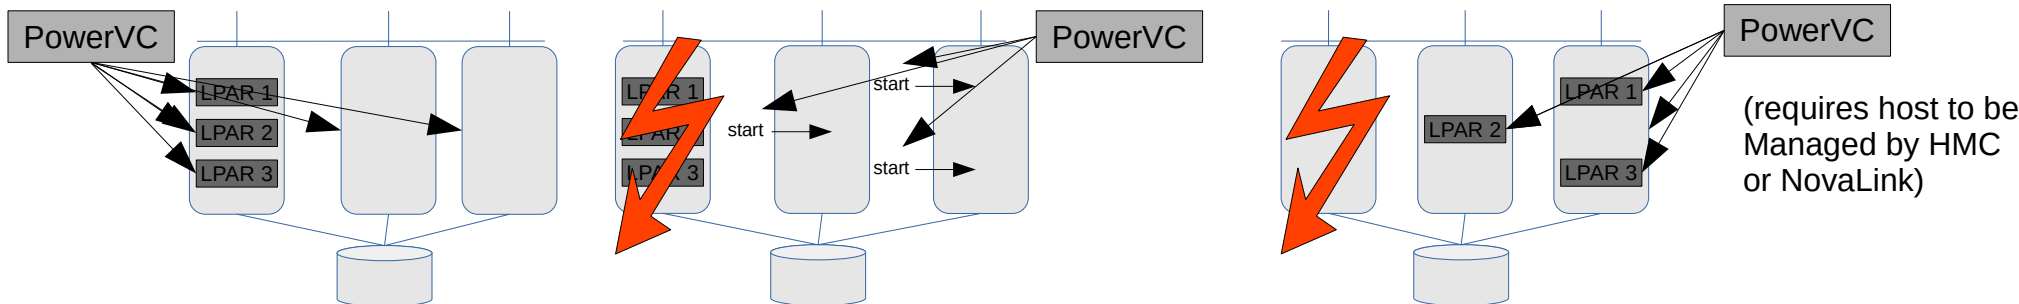


Simplified Remote Restart and PowerVC control

- If Host is in a “suitable failed” state (depending on the version)
 - Similar configuration requirements as LPM
 - OS Independent; outage (restart LPAR on new Host)
 - LPAR is restarted as a new profile, old profile must be tidied up for SRR
- HMC to restart LPAR on another host (if SRR enabled)

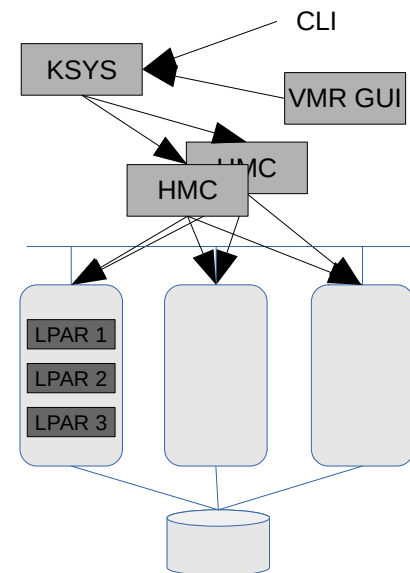


- Through PowerVC can manage which LPARs are restarted and where



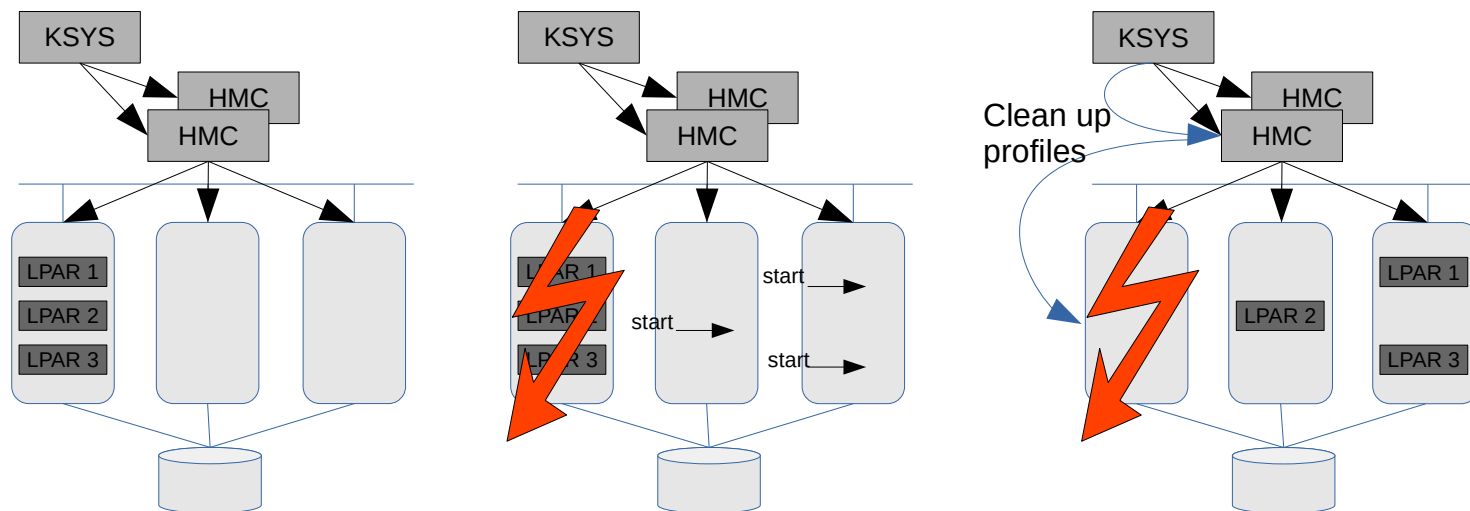
VM Recovery Manager HA

- Configuration
 - Virtualised environment with similar configuration requirements as LPM
 - OS independent; outage (restart LPAR on new host)
 - Managed by AIX LPAR (KSYS LPAR) with optional GUI, otherwise CLI
- Monitoring, KSYS LPAR monitors:
 - Advanced policies around colocation; anti-colocation; priority and server blacklist can be defined
 - Host health (are all VIO Servers healthy using Host Monitor daemon, default in VIO 3.1+) – appropriate LPARS restarted based on policies and priority
 - VM health (needs VM Agent installed) – only unhealthy LPAR restarted
 - Application monitoring and restart with VM agent installed (number of attempts locally, then move LPAR)
 - Cleans up the LPARs on failed hosts
 - Will use LPM for maintenance operations, including host vacation for maintenance



VM Recovery Manager HA (Cont)

- KSYS monitors status of Host, when failure detected
 - Specific LPARs restarted with placement following collocation/anticollocation; priority; blacklist rules
- Newer versions of KSYS also
 - monitor the OS instance and the application (User configuration required)
 - Can now get the status of the LPAR from the VIO Server if the HMC is unreachable.
 - Proactive monitoring of the VIO Servers (CPU, Memory, File system etc) and issue warning




VM Recovery Manager DR

- VM Recovery Manager DR version extends VM Recovery Manager HA to include sites and control of storage replication. It supports:
 - PowerVM for Power 7+, 8, & 9
 - Supports AIX, IBM i and Linux
 - DS8000, Storwize(FlashSystems), EMC and Hitachi & A9000 (synchronous or asynchronous)
- Same configuration of KSYS and management, except KSYS LPAR runs at DR
- One instance of KSYS can manage local VM Recovery Manager HA and VM Recovery Manager DR for a DR site
- Supports M:N configuration (ie not 1:1 hosts at Prod:DR)
- Multiple sites can have use the one DR site
- Different networks Prod and DR?
 - For AIX use *ghostdev* to clean up and script to configure in new environment, for Linux script *systemd*
- Agent now support SAP HANA and has ability to manage SAP HANA System Replication restart in primary or secondary mode
- Includes DR rehearsal (ability to create Storage snapshot in DR, start LPAR in isolated environment and clean up after testing)


VM Recovery Manager DR flow

- With local failure acts like VM Recovery Manager HA, restarting appropriate LPARs following the policy/placement rules




VM Recovery Manager DR flow (Cont)

- With any planned or unplanned outage, manually initiate recovery using *ksysmgr* on DR site.
 - With unplanned recovery, the storage is failed over to the DR site and you must manually resynchronise the storage when the Primary site recovers
 - With planned recovery, the storage replication is reversed



VM Recovery Manager DR – DR rehearsal

- Ability to start cloned LPARs in DR in an isolated network



VM Recovery Manager pre-requisites

- Hostgroup configuration: Plan and select Host/Servers such that they are cabled to:
 - Support LPM of any LPAR from any Host to any other Host within the Hostgroup
 - Would require minimally Storage and network be flat across the Hostgroup
- Two VIOS per Host is mandatory
 - If more VIOS exist on the Host, they can be included or excluded from HA configuration
 - Only 2 VIOS will be used for HA management
 - Deploy AIX rules inside VIOS to have best recommended environment
 - Have enough free space inside VIOS /, /var, /usr file systems
 - Independent redundant SAN and network fabrics across VIOS inside Hosts in the Hostgroup is mandatory
 - Need to have atleast 2 disks shared across all the VIOS that will be part of the HA management for the Hostgroup
 - Add atleast 0.5 core and 2 GB memory on top of the VIOS sizing you plan to deploy for the environment
- Two HMC per Host is strongly recommended
 - HA perspective redundancy is necessary
 - Required: Redundant connections from the KSYS to HMC and from HMC to VIOS logical partitions
 - Have enough free space inside HMC
- KSYS
 - Typical sizing: 1 core, 8GB. If you are using command line and environment managed is few LPARs, these requirements can be relaxed(User has to test over a period of time to make these adjustments to their satisfaction)
 - UI Server: Minimum: 1 core and 8 GB
 - If you are installing UI server on KSYS, you need minimum of 1 core and 8GB

Supports P7+, P8, P9

Guest OS in VMs	AIX: v6 or later IBMi: v7.2 or later (no VM Agent support) Linux: RHEL (le/be): 7.2 or later SUSE (le/be): 12.1 or later Ubuntu: 16.04 (No VM Agent support)
VIOS	VIOS 3.1
HMC	V9 R9.1.0 or later
Storage	Any supported by VIOS
KSYS LPAR	AIX 7.2 TL2 SP1 or later

VM Recovery Manager pre-requisites (Cont)

- IBM i LPARs must have “Restricted I/O Partition” selected
- LPARs must be HMC managed (if LPAR co-managed by HMC / NovaLink, then HMC must be master)
- Check the IBM Planning documentation for requirements, limitations and prerequisites
 - VM RM HA
 - https://www.ibm.com/support/knowledgecenter/SSHQN6_1.4/base/vmrecovery_ha_planning.html
 - VM RM DR
 - https://www.ibm.com/support/knowledgecenter/SSHQV4_1.4/base/vmrm_planning.html

PowerHA System Mirror Editions

- **Standard Edition**
 - High Availability within a Data Centre
 - Health management of the cluster
 - Service availability against failures in any component of the stack
 - Tools to assist with manage of clustered environment
 - Tools to assist with management and testing of PowerHA
- **Enterprise Edition**
 - All of the above
 - Adds long distance failover for Disaster Recovery
 - Low cost host based mirroring support
 - Extensive support for storage array replication
 - IBM DS8K, SVC, XIV, EMC, Hitachi, HP

	SE and EE
Centralised Management C-SPOC	✓
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	✓
	EE
Multi Site HA Management	✓
PowerHA GLVM async mode	✓
IBM Metro Mirror support	✓
IBM Global Mirror support DS8700	✓
EMC SRDF sync/async	✓
Hitachi Truecopy	✓
Stretched or linked clusters	✓
DS8000 Hyper Swap	✓

See PowerHA planning: https://www.ibm.com/support/knowledgecenter/SSPHQG_7.2/plan/ha_plan.html

PowerHA System Mirror Editions (Cont)

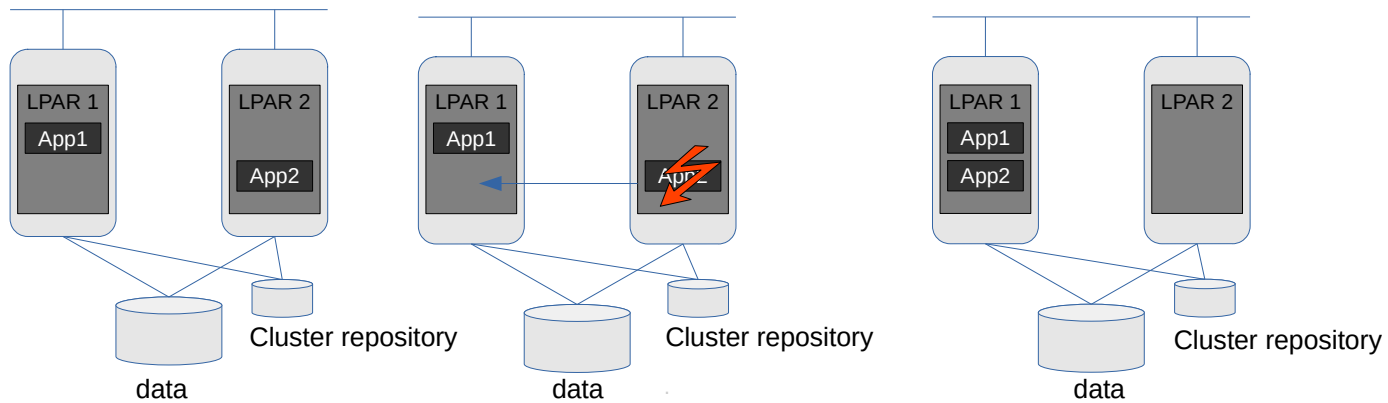
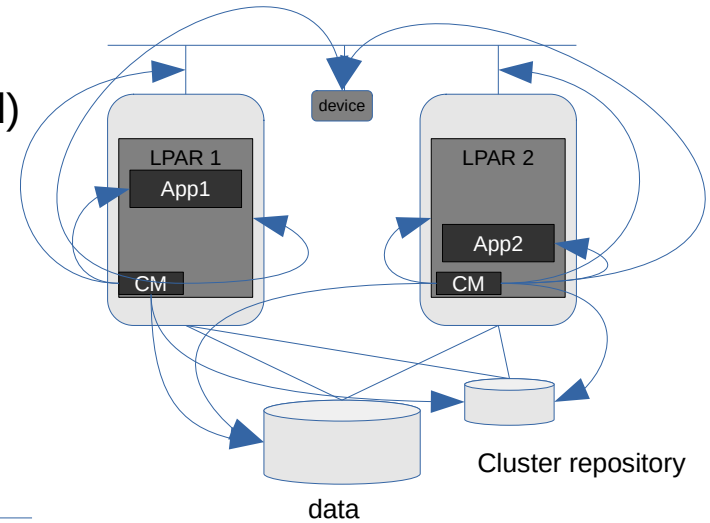
- Standard Edition
 - Supports up to 16 nodes
 - Supports Manual or Smart Assist based Deployments
 - Traditionally shares same common storage enclosure
- Supports 2 Site configurations:
 - No Copy Services Integration
 - No IP Replication Integration
 - Supports Site Specific IPs
 - Can be used with SVC Stretched Clusters
 - Used with Cross Site LVM configurations
- Enterprise Edition
 - Supports up to 16 nodes
 - Supports Stretched or Linked clusters
 - Application Smart Assistants also included for local portion of failover configuration
 - Provides local & extended cluster remote replication functions
 - Concept of tie-breaker
- Can be configured to provide local clustering capabilities at first site and automated failover to remote site
 - Automates storage level replication
 - Automates IP Replication (GLVM)
 - Integrates with DS8800 Hyperswap
 - Supports up to 2 active Sites (+ tie breaker)
 - Supports Split | Merge Policies
 - Higher Price per core

PowerHA System Mirror Standard Edition (SE)

- PowerHA SE has a long history both inside and outside IBM
- Principle is for a daemon on each LPAR to:
 - Monitor changes in Cluster Aware AIX (AIZ); RSCT (Linux); i resource monitoring (IBM I). Cluster is not formed at this level, but changes in communication paths is reported to PowerHA Cluster
 - Network connectivity
 - Storage
 - System health
 - Application
 - And respond to changes by moving application to avoid being affected by any change
- All the resources are put into a “Resource Group”, which is then managed by PowerHA, it can consist of:
 - IP address(es) – can be site specific and have location dependencies
 - Disks; File systems; Storage controllers
 - Application control scripts (start, stop, monitor)
 - CPU and memory resources
- Policies can be defined to control the behaviour of these Resource Groups
 - Startup, failover and failback policies
 - Start on same node, start on different node, start before, start after, priority and parent/child

PowerHA SE flow

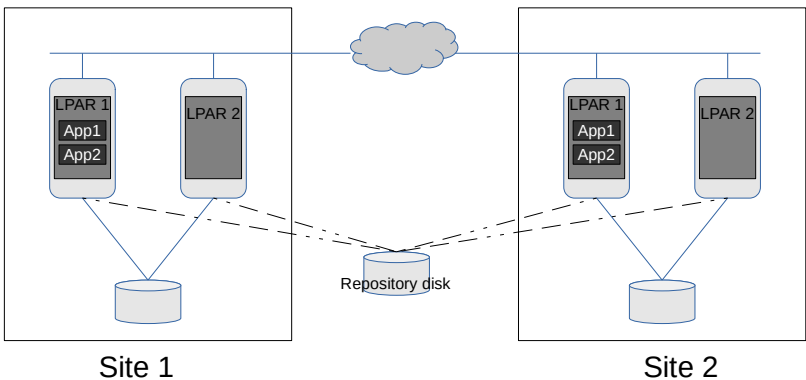
- PowerHA / CAA (or RSCT) monitoring the infrastructure
- Monitors the status from other nodes
 - Some information via network / SAN heartbeat (if supported)
 - Details in Cluster repository
- Cluster manager on all active nodes coordinate action



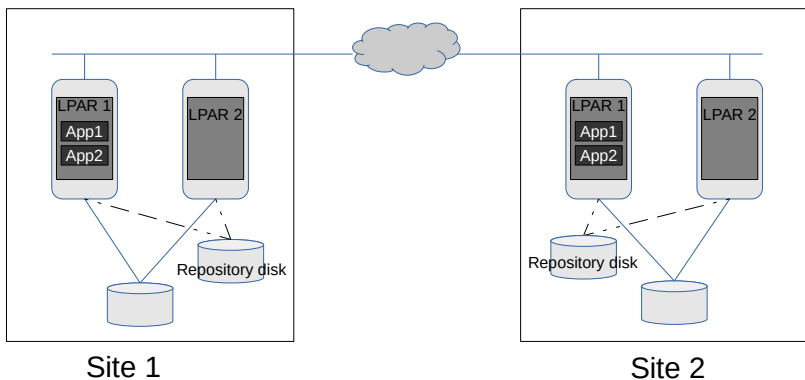
- Enterprise Edition builds on the concept of sites and adds control of storage replication between the sites.
 - Ability to manage replication via storage controller or over IP
 - Added site dependencies to Resource Group Behaviour
 - Policies around site split and merge and procedures to assist recovery of Production / DR sites
- Recent versions have added in the concept of a 3rd site or tie breaker site to avoid a split brain / site divergence (No applications can run at 3rd site for this cluster)
- The ability to have operator control over movement between sites
- Support for HyperSwap

Linked and Stretched Clusters

Stretched Cluster



Linked Cluster

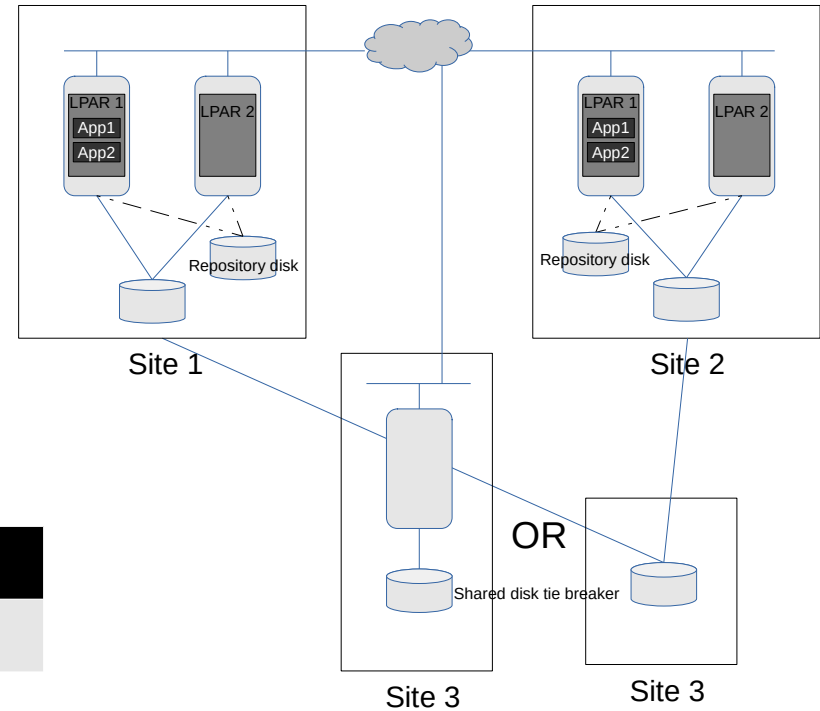


Multi Sites	Stretched Cluster	Linked Cluster
Inter site communication	Multicast	Unicast
Repository disk	Shared	Separate
Cluster Communication	Networks SAN Disk	Networks
Cross site LVM mirroring	✓	✓
HyperSwap	✓	✓
Multi site Concurrent RG with HyperSwap	✓	✓

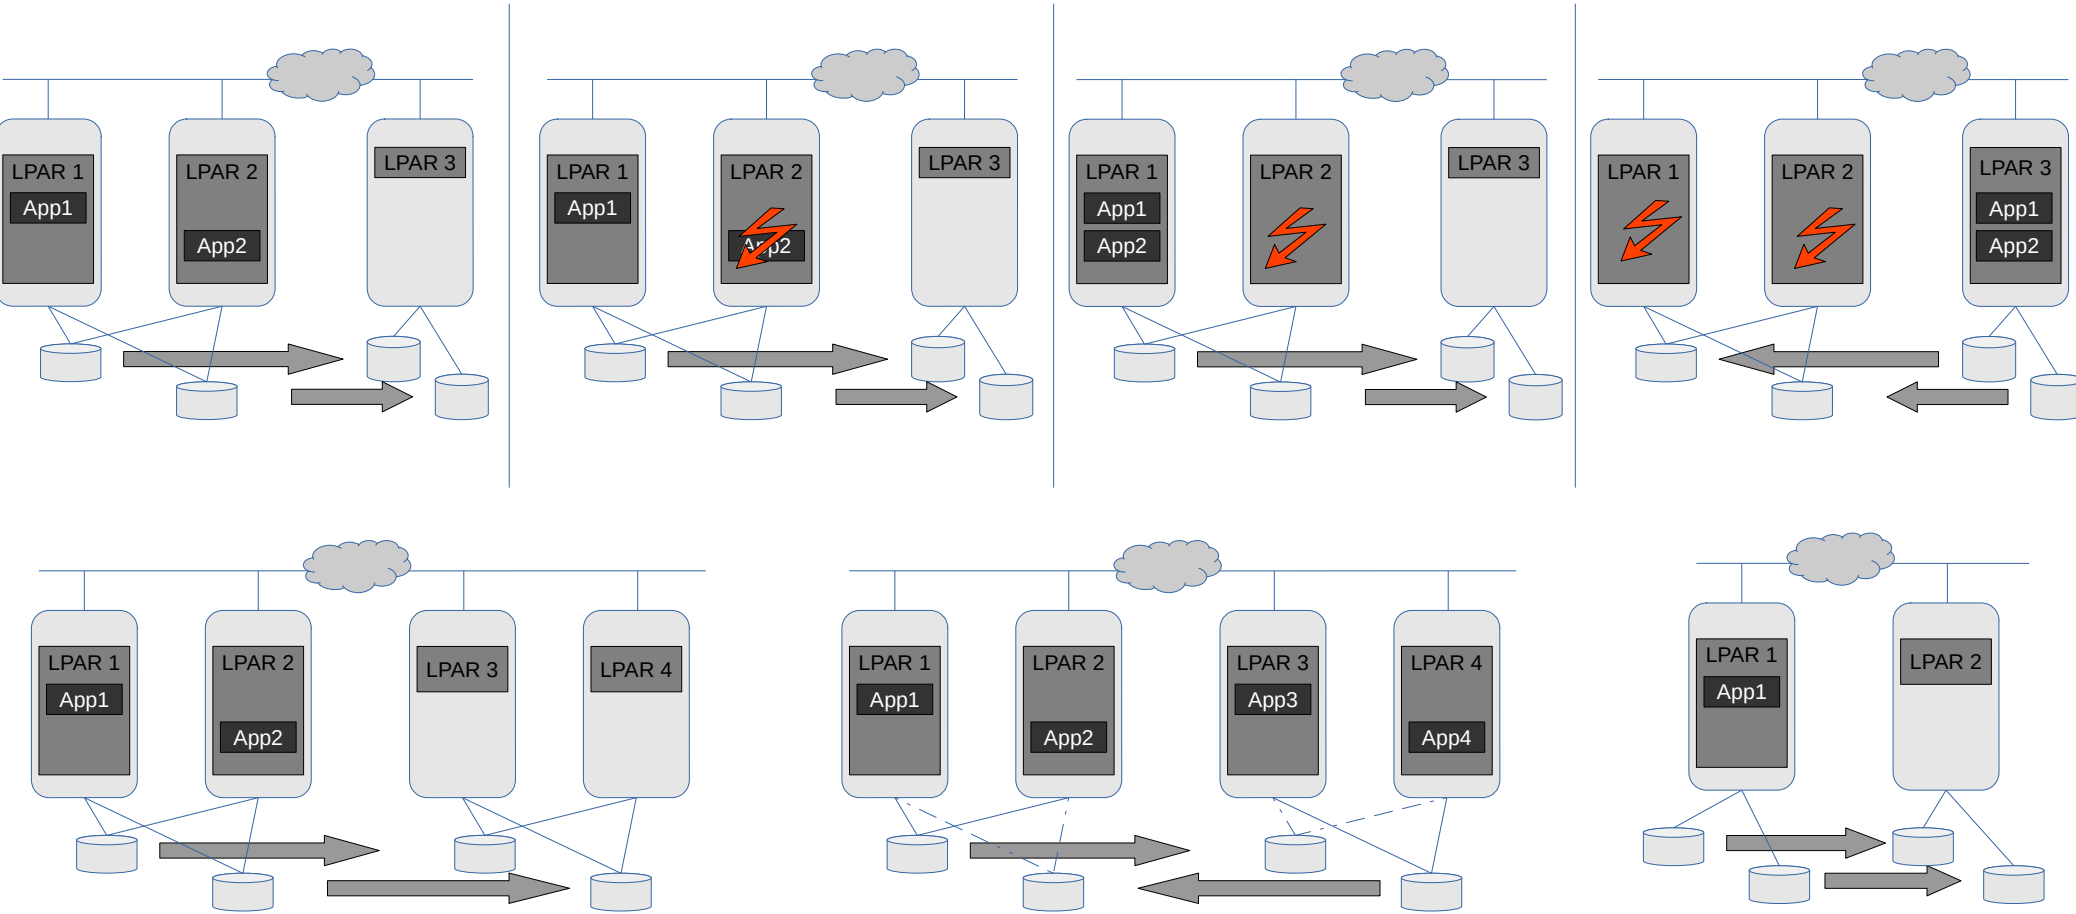
	Standard	Enterprise
Multi Site Definition <ul style="list-style-type: none"> • Site Service IP • Site Policies 	✓	✓
Stretched Clusters	✓	✓
Linked Clusters	✓	✓
HADR with storage replication management	✗	✓
Hyperswap	✗	✓

Tie breaker support

- PowerHA EE Tie Breaker Support
- More suited for Linked Clusters
 - Separate Site Split and Merge policies
 - Split/Merge: Tie Breaker policy
 - Shared (FC/iSCSI) disk or NFS file for Tie Breaker
EMC not supported with PowerPath, use AIX MPIO
 - SCSI 3 reservation disk
 - Losing side is quiesced

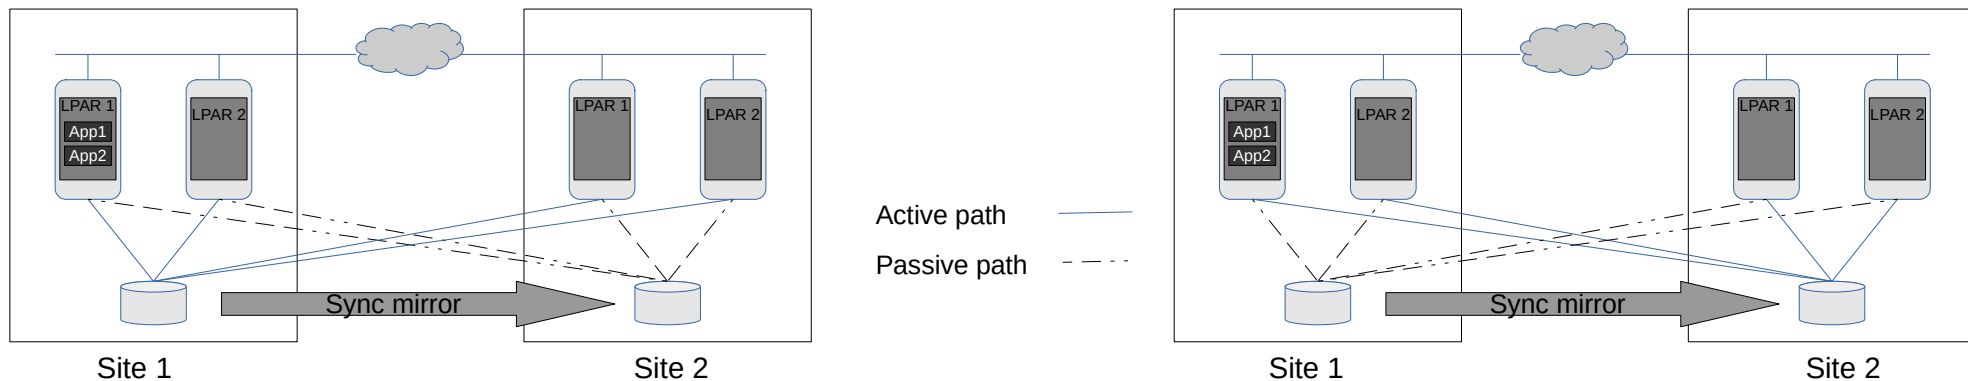


Policy Setting	Split	Merge	Comments
Tie breaker	✓	✓	Tie break Holder side wins
Majority rule	✓	✓	>N/2 side wins In case of N/2, side that includes node with the smallest node id
Manual	✓	✓	Manual steps needed for recovery to continue



PowerHA HyperSwap background

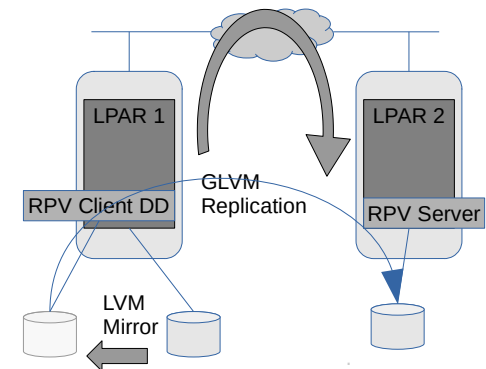
- HyperSwap originated in Mainframe Sysplex technology
 - Geographically Dispersed Parallel Sysplex (GDPS) HyperSwap is popular in z Series environments
 - Customer requests for similar capability on Power.
- PowerHA HyperSwap introduces the feature on Power with DS8800
 - The HyperSwap® function is in PowerHA® SystemMirror 7.1.2, Enterprise Edition for AIX
 - Provides for continuous availability against storage errors.
 - Based on storage-based synchronous replication
 - Enables the LPAR to transparently switch an applications I/O operation to the secondary volumes, provided physical connectivity exists between the LPAR and the secondary storage subsystem.
- Continuous availability against storage subsystem failures and supports planned HyperSwap.
- PowerHA controls recovery and auto resync.



Geographic Logical Volume Manager

- GLVM has been part of AIX since AIX 5L – part of the Logical Volume Manager
 - Rarely used without PowerHA as has no monitoring or control (easy to corrupt data)
- One of the supported methods for data replication to DR for the Enterprise Edition
- Replicates data between sites over IP network
- Supports both synchronous and asynchronous replication
- Creates a local “logical” construct of the remote disk (remote physical volume). This remote physical volume is then used to build a Volume Group with local physical disks – then normal LVM mirroring applies
- The concept of superstrict allocation policy and mirror pools were create to ensure that the correct copy of each logical partition was placed on the correct physical volume.

- ... Now offers some interesting options to replicate to the cloud



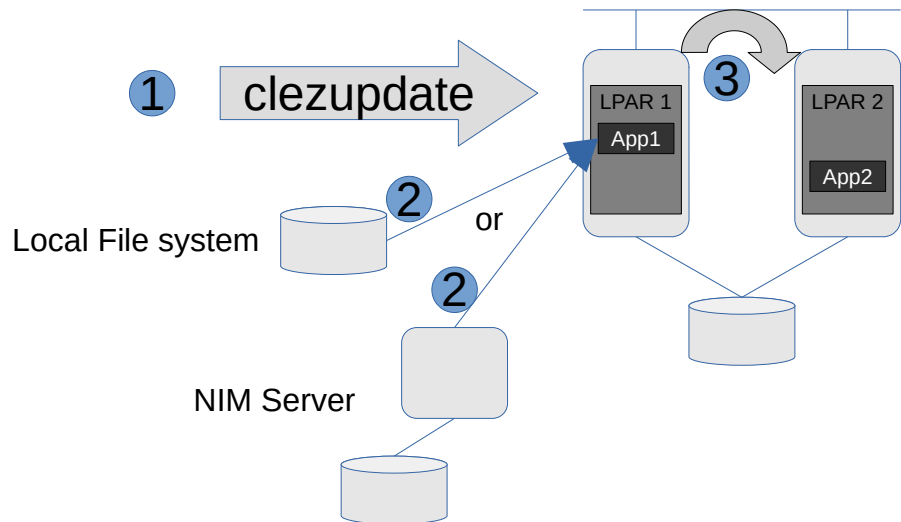
PowerHA management

- PowerHA checking / tools
 - A cluster – consistency is critical – detailed and regular verification checks
 - Feedback from support (avoid common mistakes)
 - Test tool (use automatic or customise your own test plan)
 - Customise events and responses
 - Detailed Logs (unified in GUI)
 - Deadman switch (more stable on AIX as RSCT moves to kernel based)
 - Application monitoring (using Customer or Smart Assist scripts)
 - Startup; long running; both. Include notification and can have many / application
 - PowerHA checks return code of start scripts, need application monitor to confirm
 - Label VGs as Critical VGs
 - Run a notification method; halt node processes; fence node; shutdown cluster services
 - Health dashboard
- Upgrade maintenance
 - Non-disruptive upgrade (NDU) – migrate without downtime (version dependent)*
 - Rolling upgrades (mixed levels for short periods); snapshot upgrades*

* Check readme / send me a note

PowerHA SystemMirror: Easy update tool

- Simplified IFIX and SP updates for one node or the entire cluster
 - clezupdate: Push Button, non-disruptive in production cluster update support
 - Command line tool to address various types of non-disruptive updates
 - Automate all steps of migration documentation
 - Cluster wide (one node at a time) update supported (handled internally)
 - Detailed checks, messages/guidance and error messages
 - Support for NIM Server or local file system
 - Integrated with NIM server




Update type	Reboot?	Non disruptive
PowerHA iFix, SP or TL	No	Yes
AIX (CAA/RSCT) iFixes	No	Yes
AIX SP (all AIX changes)	Mostly	No
AIX TL (coming)	Not supported	N/A


PowerHA now supports AIX 7.2 Live Kernel patching

PowerHA GUI


- Started as tool for monitoring and managing cluster
 - want to know about the health of my environment ... where the resources are, active node or standby node, etc.
 - The kind of events that get posted is overwhelming, you get lost with the important and non-important info.
 - There might be 100 lines relevant to the event how do I know which one I look at?



Browser (Chrome or Firefox)



UI Server (LPAR)
Power: AIX 7.1,7.2;
SUSE 12 le;
RHEL 7 le/be



Using the GUI to resolve PowerHA problems

- In just a few minutes we were able to:
 - determine which resource(s) had a problem

The screenshot displays the IBM PowerHA SystemMirror for AIX GUI. On the left, a sidebar lists resources with their status: demo_cluster (Warning), phac1 (Maintenance, offline), rg1 (Maintenance, offline), rg2 (Maintenance, offline), rg3 (Offline), phac2 (Normal), rg1 (Maintenance, offline), rg2 (Maintenance, offline), rg3 (Online), and cluster2 (Stable). The main panel, titled 'All Events Summaries', shows a summary for 2 Clusters, 5 Nodes, and 12 Resource groups. The summary includes counts for Critical, Warning, and Maintenance events. Below this is a filter section with icons for Critical, Warning, Maintenance, and OK, along with a search bar and a 'Reset' button. A timeline chart shows event counts from May to August 2016. The chart data is as follows:

Month	Day	Count	Status
May	11	11	OK
May	20	20	OK
Jun	72	72	OK
Jul	140	140	OK
Jul	18	18	OK
Jul	22	22	OK
Aug	297	297	Critical
Aug	687	687	Critical

The timeline also shows two event details for 'Event external_resource_state_change_complete' at 11:30:58 AM, both occurring on cluster2 -> phac3. A third event, 'Event RG_OFFLINE_STATE', is listed at 11:30:54 AM. The total number of events is 1619.

Using the GUI to resolve PowerHA problems

- In just a few minutes we were able to:
 - determine which resource(s) had a problem
 - identify the problem

The screenshot displays the IBM PowerHA SystemMirror for AIX GUI for a cluster named 'demo_cluster'. The interface shows a navigation sidebar on the left with icons for home, view, warning, and success. The main content area is divided into several sections:

- Cluster Overview:** Shows '2 Nodes' and '6 Resource groups'. For nodes, there is 1 Critical, 0 Warning, and 0 Maintenance. For resource groups, there are 4 Critical, 0 Warning, and 0 Maintenance.
- Filter and Search:** A filter bar includes icons for Critical, Warning, Maintenance, and Success, along with a search box and a 'Reset' button.
- Event Timeline:** A timeline for Thursday 18 August 2016 shows events at 9:00, 15:00, 13:00, 13:00, and 3:00 PM. The 13:00 event is marked with a green checkmark, while others are marked with red exclamation points.
- Event Details:** A detailed view of the 'Event node_down_complete' at 4:41:10 PM. The description states: 'The event script runs after the node_down event completes successfully.' The path is 'demo_cluster » phac1'.
- Event Details (Bottom):** A partial view of the 'Event RG_ONLINE_STATE' at 4:41:04 PM.
- Right Panel:** Contains detailed text for the 'Event node_down_complete', explaining that it runs when a node is not running cluster services or during a stop. It also provides instructions on how to stop cluster services manually and where to find log files (hacmp.out).
- Buttons:** At the bottom right, there are buttons for 'Open log' and 'Open terminal'.

Using the GUI to resolve PowerHA problems

- In just a few minutes we were able to:
 - determine which resource(s) had a problem
 - identify the problem
 - solve the problem

IBM PowerHA SystemMirror for AIX

demo_cluster

Events summaries | **Logs** | General | Network | Storage | Replicated storage | Applications | Snapshots

error < > | fail < > | could not < > | preamble < > X | Search logs

Open terminal

Logs

- Events log hacmp.out
- phac1 Thursday 08/18/2016 11:00:00 PM
- phac2 Thursday 08/18/2016 11:00:00 PM
- Cluster log cluster.log
- Error report
- Cluster utility log clutils.log
- Cluster manager log

phac1 Cluster verification log (clverify.log)

```
4817 BEGIN CHECK
4818
4819 END CHECK
4820 Check: PASSED
4821
4822
4823 Total elapse time: 17 second(s)
4824
4825 Verification has completed normally.
4826
4827
4828 SUMMARY REPORT
4829
4830 Results of: clver -c -V normal -O -t -r --
4831
4832 Node      Errors Logged
4833 -----
4834 All       0
4835 phac1    0
4836 phac2    0
4837
4838 Totals   0
4839
4840
```

phac1 Events log (hacmp.out)

```
1.86612 Aug 18 2016 17:39:22 GMT -04:00 EVENT STA
1.86728 Aug 18 2016 17:39:23 GMT -04:00 EVENT COM
1.86733 Aug 18 2016 17:39:23 GMT -04:00 EVENT COM
1.86736 Aug 18 2016 17:39:25 GMT -04:00 EVENT STA
1.87094 Aug 18 2016 17:39:27 GMT -04:00 EVENT COM
1.87107 HACMP Event Preamble
1.87112 Aug 18 2016 17:39:30 GMT -04:00 EVENT STA
1.87198 Aug 18 2016 17:39:31 GMT -04:00 EVENT COM
1.87209 HACMP Event Preamble
1.87210 -----
1.87211
1.87212 Stop cluster services request with 'Grace
1.87213
1.87214 Enqueued rg_move release event for resour
1.87215
1.87216 Enqueued rg_move release event for resour
1.87217
1.87218 Enqueued rg_move release event for resour
1.87219
1.87220 Node Down Completion Event has been enque
1.87221
1.87222 :check_for_site_down[+54] [[ high = high
1.87223 :check_for_site_down[+54] version=1.4
1.87224 :check_for_site_down[+55] :check_for_site
```

Using the GUI to resolve PowerHA problems

- In just a few minutes we were able to:
 - determine which resource(s) had a problem
 - identify the problem
 - solve the problem
 - and watch as PowerHA health returns.

The screenshot displays the IBM PowerHA SystemMirror for AIX GUI. The main view shows the status of the **demo_cluster**, which is currently **Stable**. The cluster consists of **2 Nodes** and **6 Resource groups**. The status indicators for both nodes and resource groups show 0 Critical, 0 Warning, and 0 Maintenance events.

The left sidebar shows a list of resources with their status:


- cluster2: Stable
- demo_cluster: Stable
- phac1: Normal
- phac2: Normal
- rg1: Online
- rg2: Online
- rg3: Online

The main panel shows the **Events summaries** section, which includes a filter bar with icons for Critical (red exclamation mark), Warning (yellow triangle), Maintenance (green wrench), and Stable (green checkmark). A search bar is available for **Search Events**. Below the filter bar is a timeline view for the year 2016, showing event counts for each month:

Month	Event Count	Status
May	352	Stable (Green Checkmark)
Jun	50	Stable (Green Checkmark)
Jun	24	Stable (Green Checkmark)
Jul	122	Stable (Green Checkmark)
Jul	22	Stable (Green Checkmark)
Aug	427	Critical (Red Exclamation Mark)

The total number of events is 997. The current event shown is **Event RG_ONLINE_STATE** at 3:55:34 PM, with a description that is not available.

PowerHA System Mirror life cycle



Extended Support



End of Life

EEOS: Estimated End of Support

PowerHA 7.2.3 SP1 (Linux)	SUSE	Red Hat
2019	SLES 12 SP1,2,3	RHEL 7.2, 7.3, 7.4, 7.5, 7.6,
2019 Q4	SLES 12 SP5 SLES 15 SP1	RHEL 8

PowerHA Licensing options

- Now a number of options to avoid paying for unused cores in a PowerHA Cluster (SE / EE)
- Resource Optimised High Availability (ROHA)
 - Standby machine with minimum resources for OS alone, extra resources added when Application(s) started, removed when Application(s) moved
 - With Applications active on all Nodes, optimise resources for each Application and add/subtract resources when Applications move
 - Supports
 - Enterprise pools
 - CoD
 - Enterprise Capacity Backup (ECBU)*
 - Discounted processor nodes matching the installed production server processor nodes
 - No charge, annually renewable active standby memory = $365 \times N \times 32\text{GB}$ where N is the number of active mobile cores on the production
 - Mobile processor activations are transferred from production to ECBU via Enterprise Pool transfers
 - Registration of primary system and ECBU required. Primary and ECBU must be within the same enterprise

* Talk to your IBM Rep

PowerHA new features summary

- New features for PowerHA 7.2.3
 - Cloud Backup Management
 - Customers can now use IBM Cloud or Amazon AWS, as a data backup option. Collects the flash copy of the application data at the defined backup schedule, and copies to the IBM Cloud Object Storage S3 or Amazon AWS S3 storage using S3 boto interface.
 - Availability metrics
 - Log analyser and miscellaneous features
 - GUI enhancements (Power UI Server can be AIX 7.1,7.2; SUSE 12 le; RHEL 7 le/be)
 - Miscellaneous tidy of inconsistencies and features:
 - Some customers requested use of rsyslog (AIX uses syslog by default)
 - CSPOC support for LVM preferred read
 - Smart assist for Oracle now has options for shutting down oracle (immediate, abort, normal, transactional)
 - Administration events, so that administrator initiated tasks are logged and handled the same way.
 - Event serial number to assist tracking through the various PowerHA logs
 - Improvements in recovering from script failure and option to cancel remaining event processing
 - Failures in pre/post/notify event processing can stop the processing of the event
 - Improvements in the log analyser (added options and strings not case sensitive)
 - Network flapping events – define a threshold

Planned improvements

- Cross cluster verification (Done – Dec 19)
 - Compare test and development clusters
- Availability metrics
 - Add failover times
 - Provide historical averages
- Smart Assist updates
 - TSM
 - Federated Security
 - DB2 multiple instance support
- Customer requests
 - Support for SVC role based access
 - Encrypt all passwords
 - Increase maximum number of resource groups to 256
 - Support for Websphere MQ listener
 - Support for SVC v8.1

- After 7.2.2 was released, a review conducted:
 - Now that I know where the problem is, I want to be able to fix it.”
 - The ability to take actions on the managed clusters was asked for, which implied a higher level of security and organization
 - User management
 - Role-based access control
 - Cluster zones
 - Added support for Power Linux
 - Linux clusters can be managed on AIX server
 - AIX clusters can be managed on Linux server
 - One server can manage both Linux and AIX clusters
- For 7.2.3
 - 5 customers were involved in the research
 - Adds more administrative capabilities
 - Modify existing cluster configuration objects
 - Delete existing cluster configuration objects
 - Installation improved
 - Highly available

PowerHA Linux new features

- New features for PowerHA linux 7.2.2
 - Filesystem support on shared disk
 - Support for ext3, ext4 and XFS
 - FQDN support
 - LDAP support
 - Support of SAP HANA and NW on same cluster
 - SAP HANA support on RHEL
 - Globalization support - translation to Japanese
 - PowerHA Migration (offline and rolling)

PowerHA for IBM i new features

- 7.4 release has improvements around integration, enhanced usability and automation of the administrative domain
 - Integration - Integration with IBM Copy Services Manager and DB2 Mirror for i
 - Enhanced Usability - Grouping of Cluster Resource Groups (CRGs) with CRG containers and additional control over CRG failure and IASP vary on behaviour
 - Administrative Domain - Automated management of monitored resource entries and improved control over synchronisation options
- Integration
 - PowerHA SystemMirror for i continues to integrate with additional products and technologies. The 7.4 release brings the following new integrations:
 - IBM Copy Services Manager for DS8000 Replication
 - PowerHA contains support for Copy Services Manager (CSM) Metro Mirror, Global Mirror, and HyperSwap with Global Mirror.
 - IBM Copy Services Manager (CSM) is a strategic technology used to manage Metro Mirror or Global Mirror replication as well as other advanced storage capabilities. The integration of CSM with PowerHA for i enables PowerHA to support advanced functionality such as DS8000 HyperSwap with Global Mirror
 - IBM DB2 Mirror for i
 - PowerHA SystemMirror for i standard edition integrates with Db2 Mirror for i to provide support for near-continuous availability of IFS data within a Db2 Mirror for i environment. In addition, PowerHA SystemMirror for i enterprise edition enables Db2 Mirror disaster recovery capability through the use of new CRG containers.

PowerHA for IBM i new features (Cont)

- Enhanced Usability
 - The usability of the product continues to be enhanced by providing additional integrated options for configuring and controlling a PowerHA environment
 - Support for customising and controlling behaviours in a PowerHA and cluster environment via PowerHA Policies.
 - A new PowerHA policy has been introduced to allow integrated control over canceling automatic Cluster Resource Group (CRG) failovers for certain events with the new QCST_CRG_CANCEL_FAILOVER policy. For additional information see: PowerHA Policies.
 - Additional customization on device CRGs with a new vary online setting of *CURRENT. When a device entry in a CRG has a vary online setting of *CURRENT the device will be varied online after a switchover or failover on the new primary node if the device was varied on prior to the switchover or failover.
 - New CRG containers that serve as a control object for a collection of CRGs to be managed as a single entity. Additional information on CRG containers is coming soon.

PowerHA for IBM i new features (Cont)

- Administrative Domain
 - Automated Management
 - Three new PowerHA policies: QCST_AD_CREATE, QCST_AD_DELETE, and QCST_AD_RESTORE enable the following functionality:
 - Automate the addition of Monitored Resource Entries (MREs) across nodes in the administrative domain upon object creation.
 - Automate deletion of MREs across all nodes in the administrative domain upon object deletion.
 - Honour the restore of attributes on MREs in the administrative domain upon system save and restore operations.
 - Improved control over Synchronization Options
 - When starting the cluster administrative domain with the Start Cluster Administrative Domain (STRCAD) command, it is now possible to choose which node is used as the source of updates to push to other nodes in the cluster. Previously, when resource changes were made to nodes in an inactive Admin Domain, multiple nodes could change the same resource and there was no coordination for multiple different changes to the same resource.
- Other Improvements
 - New DSPASINF command functions.
 - Continued improvements to reduce IASP vary on times.

	Storage Requirements	Automated Failover	Source Server Status	Source VIO Server Status	VM/Application Outage	RTO	RPO	License	Cost
Live Partition Mobility	Shared	No	Active	Active	No (if VM is active)	N/A	N/A	N/A	N/A
Simplified Remote Restart	Shared	No	Inactive	Inactive	Yes	Operator + IPL + App Start	0	N + 0	\$
VM RM HA	Shared	Yes	Active or Inactive	Active or Inactive	Only if a server/ VM outage occurs	IPL + App Start	0	N + 0	\$\$
PowerHA	Remote Copy	Yes	Active or Inactive	Active or Inactive	Yes	App Start	0	N + 1	\$\$\$
PowerHA EE	Remote Copy	Yes	Active or Inactive	Active or Inactive	Yes	App Start	Sync or Async	N + 1	\$\$\$\$
VM RM DR	Remote Copy	No	Active or Inactive	Active or Inactive	Yes	VMRM HA local, DR +	Sync or Async	N + 0	\$\$

Useful links

- V6 to v7.1.3 SP3 Migration guide
 - http://www-01.ibm.com/support/knowledgecenter/SSPHQG_7.1.0/com.ibm.powerha.insgd/ha_install_mig61.htm
 - youtube video: <https://www.youtube.com/watch?v=MaPxuK4poUw>
- PowerHA Forum. QA forum, & Linked in
 - Linked In: <https://www.linkedin.com/grp/home?gid=8413388>
- PowerHA & AIX IFIX Bundles
 - Fixes: <https://ibm.biz/PowerHAFixes>
- PowerHA SystemMirror & Power: Capacity Backup
 - CBU: <http://www-03.ibm.com/systems/power/hardware/cbu>
- PowerHA Fix Level Rec Tool links
 - Release Recommendation: <https://www-304.ibm.com/webapp/set2/flrt/liteTable?prodKey=hacmp>
 - PowerHA SystemMirror Notification Subscription (My Notifications): <http://www-01.ibm.com/software/support/einfo.html>
- PowerHA 7.2.3 Redbook
 - IBM PowerHA SystemMirror V7.2.3 for IBM AIX and V7.2.2 for Linux July 2019 SG24-8434-00
- VM Recovery Manager
 - Implementing IBM VM Recovery Manager for IBM Power Systems July 2019 SG24-8426-00

Useful links (older)

- Developerworks has been sunset
 - Developerworks Wiki: (now moved to: <https://www.cleartechnologies.net/phawiki/>)
 - QA Forum: <http://ibm.biz/developerworks-PowerHA-Forum>
- PowerHA videos
 - UI/Dashboard 2 min video demo
 - https://www.youtube.com/watch?v=d_QVvh2dcCM&feature=youtu.be
 - Non-disruptive upgrade from 713 to 72
 - <https://www.youtube.com/watch?v=1Kzm7I2mRyE>
 - ROHA: Flexible Capacity Management
 - <https://www.youtube.com/watch?v=G-zefev-XYU>
 - Automated Repository Replacement
 - <https://www.youtube.com/watch?v=HJZZDCXLwTk>
 - AIX Live update (PowerHA node)
 - <https://www.youtube.com/watch?v=BJAnpN-6Sno>
 - rootvg Failure monitoring
 - <https://www.youtube.com/watch?v=OZcrhVGvkBg>

Session: HA/DR Solutions on IBM Power Systems

¿ Questions ?

Thanks!

For further information....

Contact:

Antony (Red) Steel

antony.steel@belisama.com.sg

+65 9789 6663



Backup slides