

# PowerHA SystemMirror good practice (a203866)

—  
Antony (Red) Steel  
Belisama Pte. Ltd.



#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.

Thanks to:  
Steven Finnes  
Ravi Shankar  
Paul Moyer  
Dino and the Redbook team



# Agenda

- Introduction to PowerHA – Standard and Enterprise
- Cluster aware AIX
- PowerHA maintenance and features
- Recommended practice
- Problem determination
- PowerHA what is new and directions



# Introduction to PowerHA SystemMirror for AIX

- PowerHA SystemMirror for AIX Standard Edition
- Cluster management for the data centre
  - Monitors, detects and reacts to events
  - Establishes a heartbeat between the systems
  - Enables automatic switch-over
  - Closely integrated into Power and AIX
    - LVM mirroring; capacity upgrade, Live kernel patching
  - Multi-cluster federated security, user management, encryption
- IBM shared storage clustering
  - Can enable near-continuous application service
  - Minimise impact of planned & unplanned outages
  - Ease of use for HA operations
- Smart Assists – application agents
  - Out of the box deployment for SAP and other popular applications
- Mature Product
  - 26 Major releases (averaging one a year)
  - Over 18,000 customers worldwide

87% of respondents consider 99.99% (52.56 minutes) of unplanned per server/per annum downtime as the minimum acceptable level of reliability for mission critical servers and applications.

Also reported issues that cause downtime:

Human Error	59%
Security	56%
Software flaws	38%

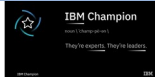
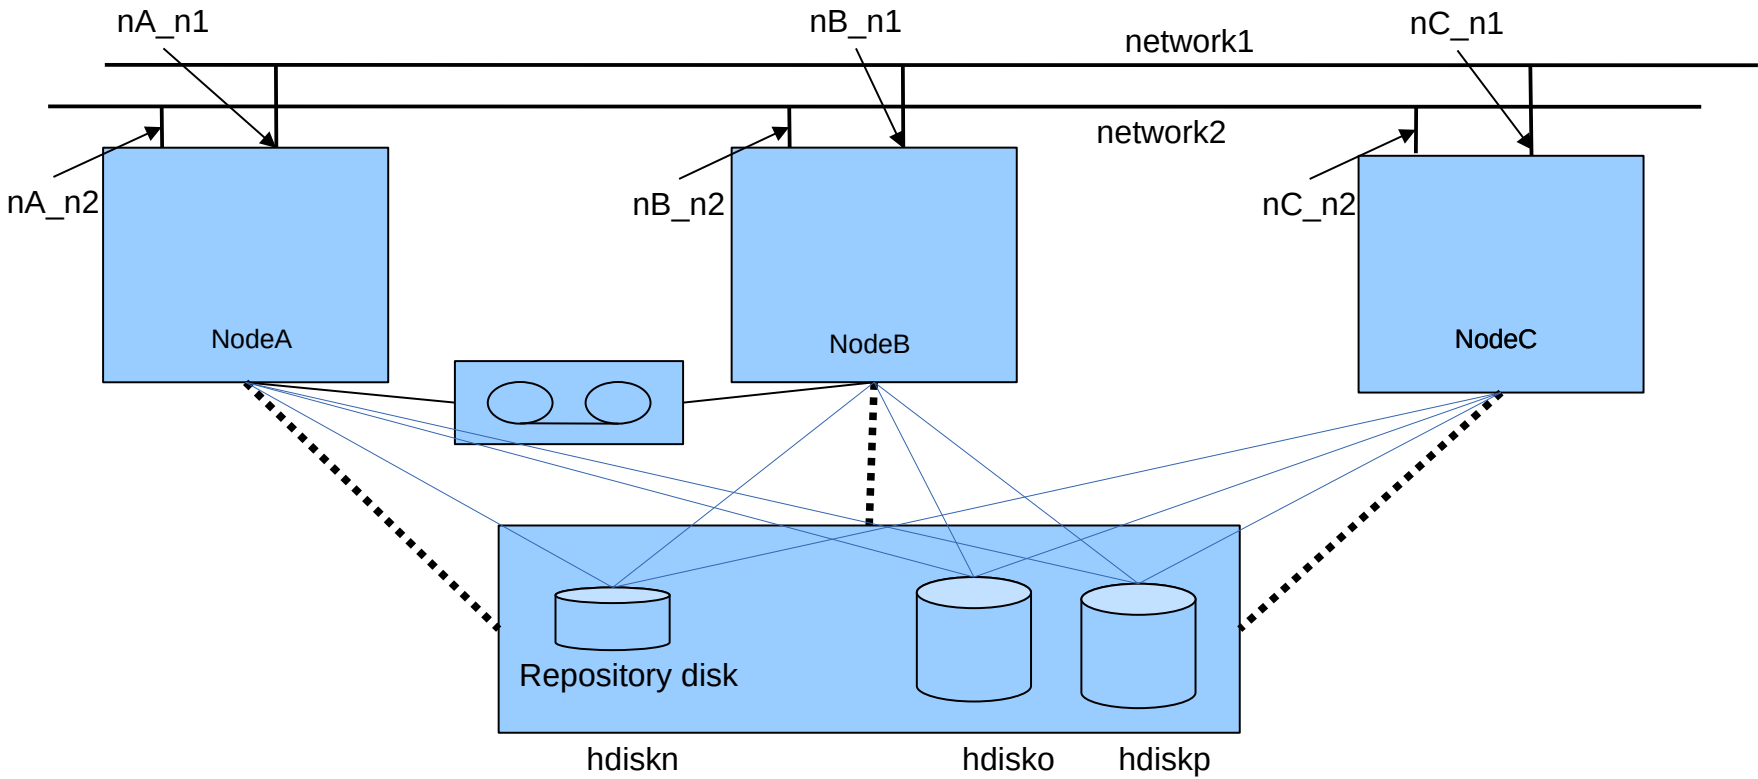
ITIC's 2020 Reliability poll

- PowerHA SystemMirror for AIX Enterprise Edition
  - Cluster management for the Enterprise
  - Multi-site cluster management
  - Includes the Standard Edition function
  - Low cost mirroring option support

# PowerHA SystemMirror editions

- 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 (GLVM) Integration
  - Supports Site Specific IPs
  - Can be used with SVC Stretched Clusters
  - Used with Cross Site LVM configurations
  - Supports Split | Merge Policies when configured as a Linked Cluster
- Enterprise Edition
  - Supports up to 16 nodes
  - Supports Stretched or Linked clusters
  - Application Smart Assistants also included for local failover configuration
  - Provides local & extended cluster remote replication functions
  - MetroMirror; GlobalMirror EMC SRDF; Hitachi TrueCopy/HUR
  - Can be configured to provide local clustering capabilities at one site and automated failover to remote site
    - Automates storage level replication
    - Automates IP Replication (GLVM)
    - Integrates with DS8800 Hyperswap
    - Supports up to 2 Sites
    - Supports Split | Merge Policies
    - Higher Price per core


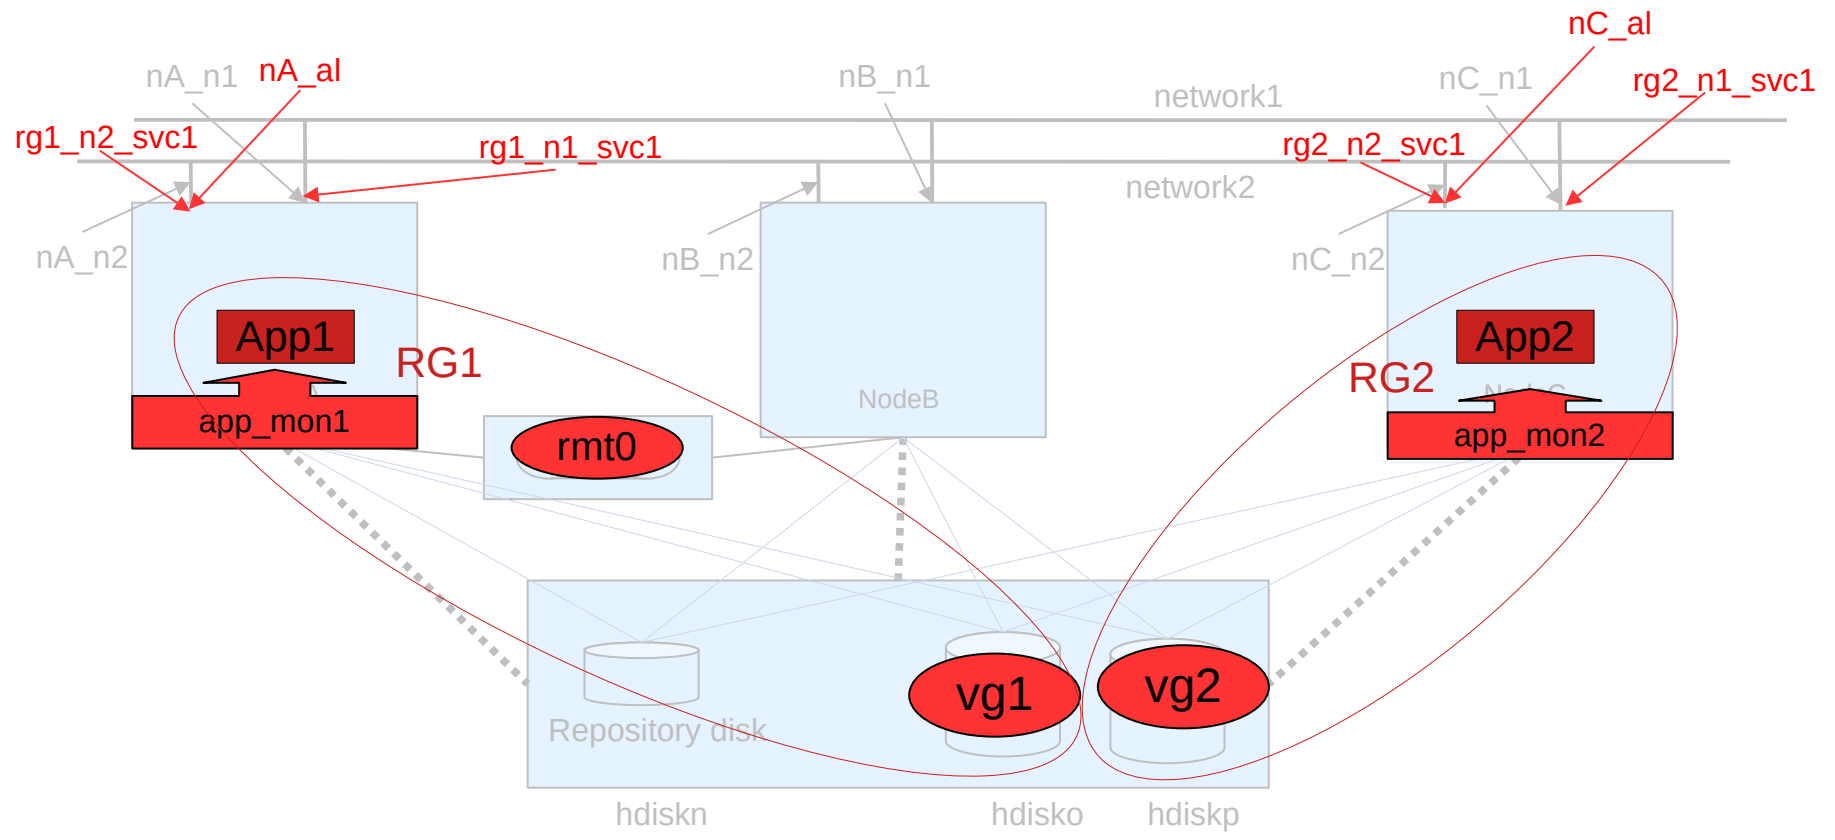
# Cluster topology



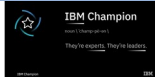
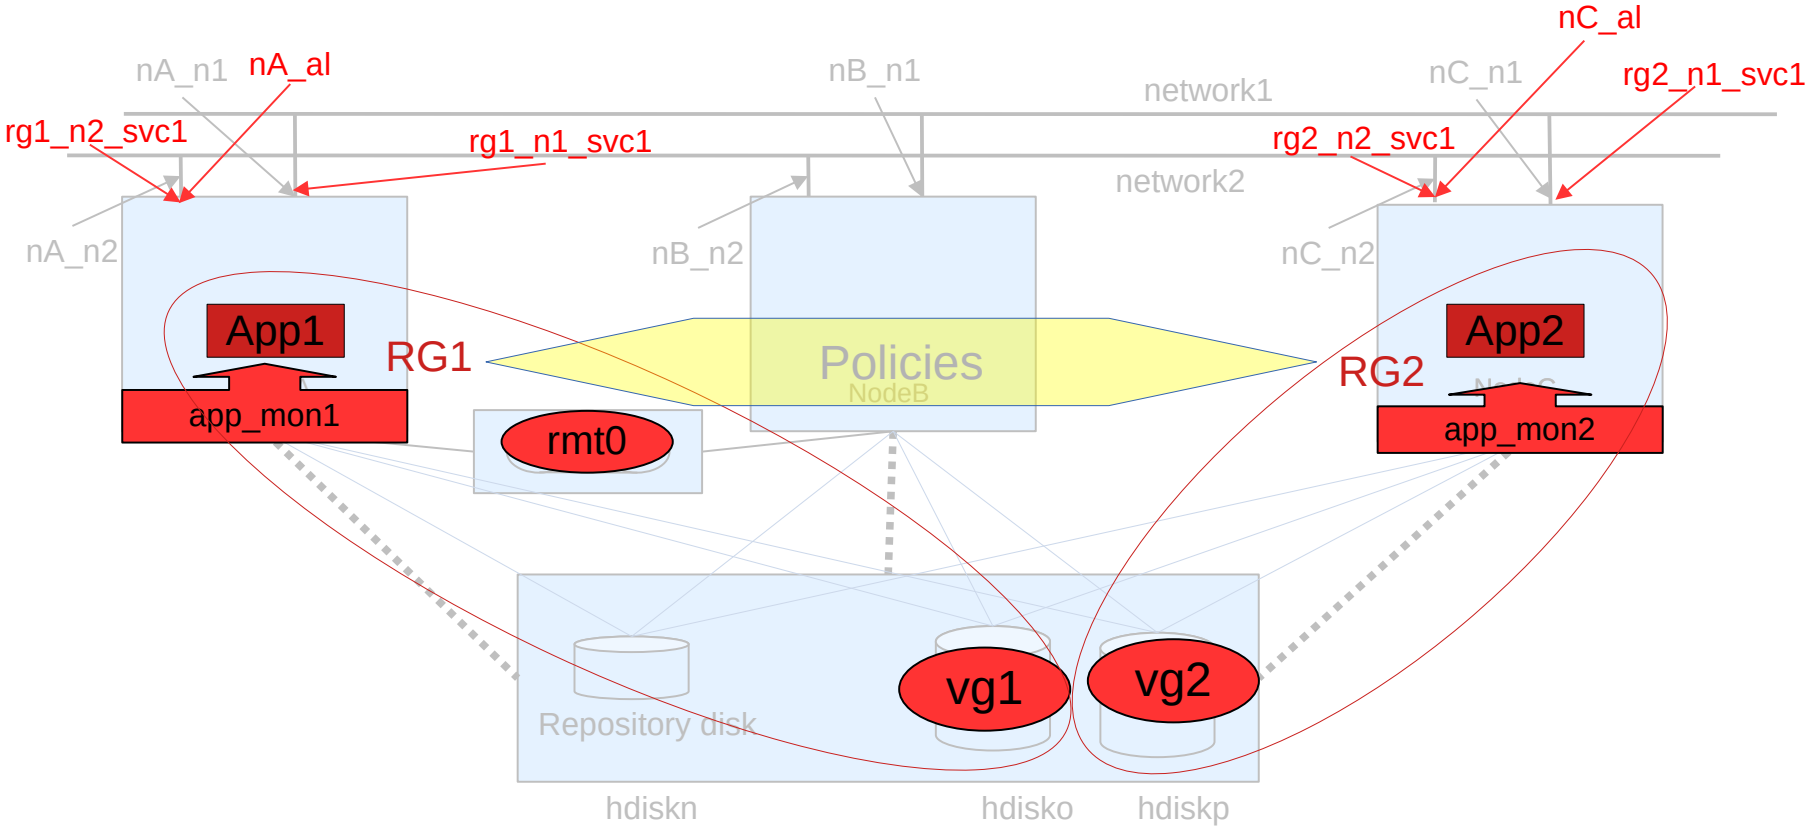
## Resource groups

- A collection of all the resources required to run the Application. Depending on the application may contain:
  - List of nodes on which to run (“home” node is the first node in the list)
  - IP Address
  - Volume Group / File systems
  - Application start script
  - Application stop script
  - Application monitor script(s)
  - Policies governing
    - Startup
    - Fallover
    - IP Address placement
- Dependencies can be configured to control the behaviour of the resource group(s)

# Resource groups

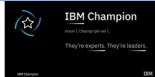
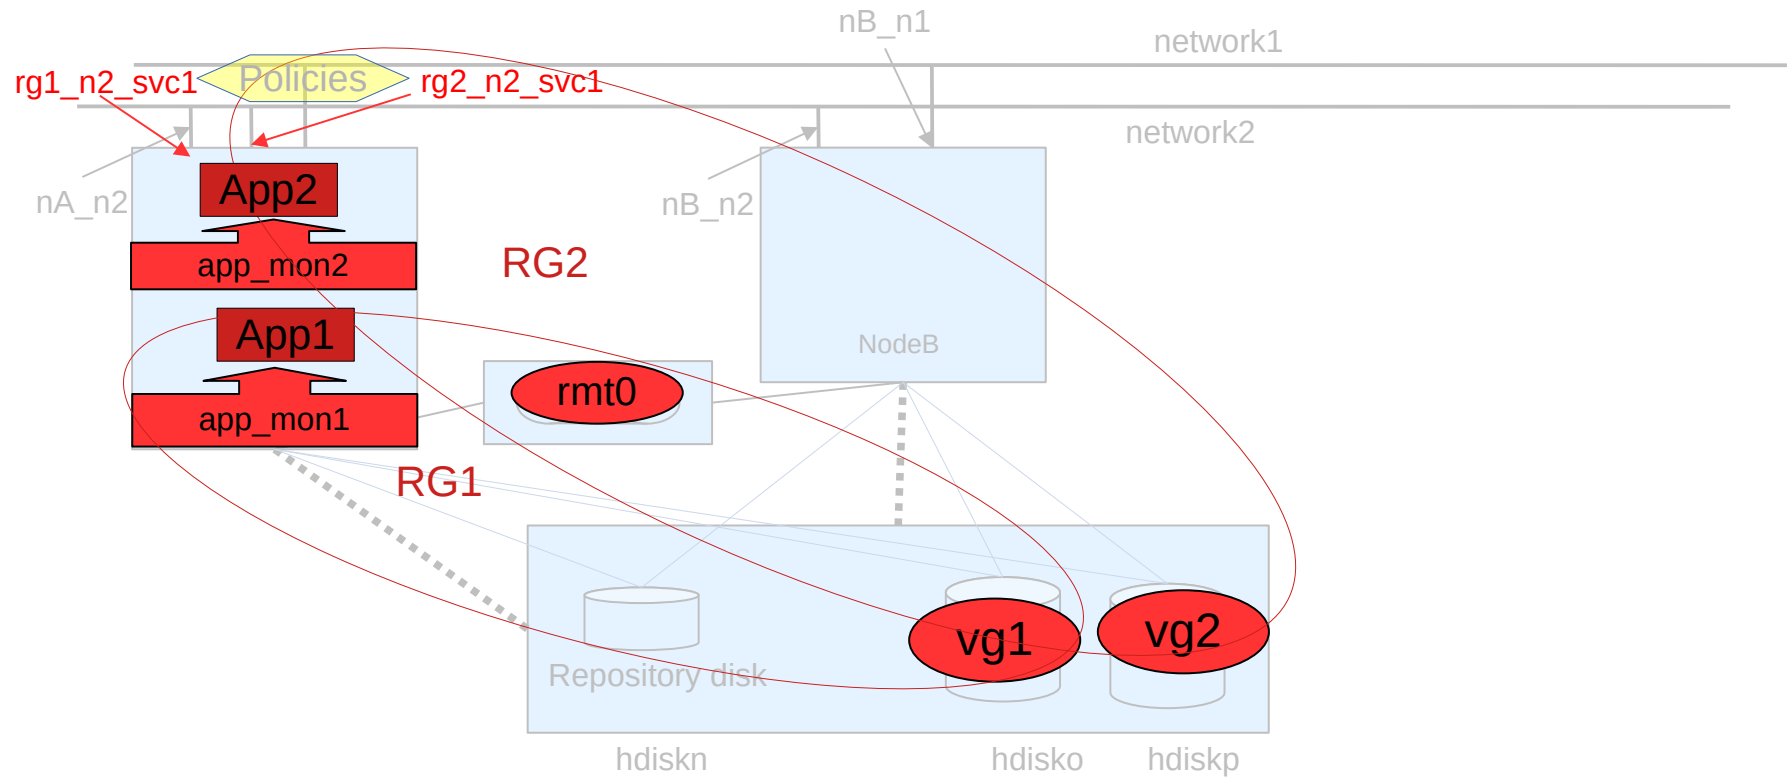


# Resource groups





# Resource groups

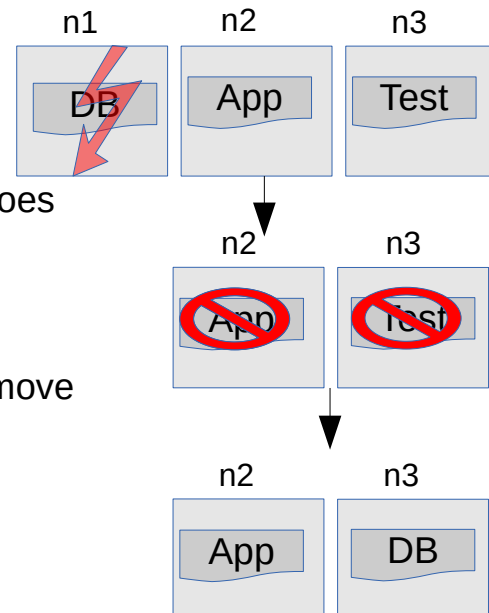


# Cluster behaviour

- Resource group policies
  - Startup
    - Online on home node only
    - Online on first available
    - Online on all available
    - Ability to control start up distribution
  - Fallover
    - Fallover to next node in the list
    - Fallover using Dynamic node priority (CPU, Paging space, Disk IO, Adaptive (user defined))
    - Bring offline
  - Fallback
    - Fallback to higher priority node
    - Never fallback
    - Delayed fallback timer
- Resource group dependencies
  - covered in detail below
- IP distribution preferences
  - covered in detail below
- Inter site management policies
  - Online on Both Sites
  - Online on Either Site
  - Prefer Primary Site
  - Ignore

# Resource Group dependencies

- Online on same node dependency
  - Resource groups come online on the same node
  - Start after and stop after
- Parent child dependency
  - Child will come online after the parent is stable, will go offline if the parent goes offline. Can have up to 3 levels
- Location (Online on different node) dependency
  - High, intermediate and low
  - High will force intermediate and low to move, intermediate will force low to move
  - Same priority cannot come online on same node
  - Same priority will not cause a movement
- Dynamic Node Priority
  - Processor Utilisation (cl\_highest\_idle\_cpu)
  - Memory Utilisation (cl\_highest\_free\_mem)
  - Disk I/O Utilisation (cl\_lowest\_disk\_busy)
- Dynamic Node Adaptive Fallover – user defined
  - cl\_lowest\_nonzero\_udscript\_rc
  - cl\_highest\_udscript\_rc



Online on different node dependency

DB: n1,n3,n2	High
App: n2,n3,n1	Intermediate
Test: n3,n2,n1	Low

Parent / Child  
DB – parent; App - Child

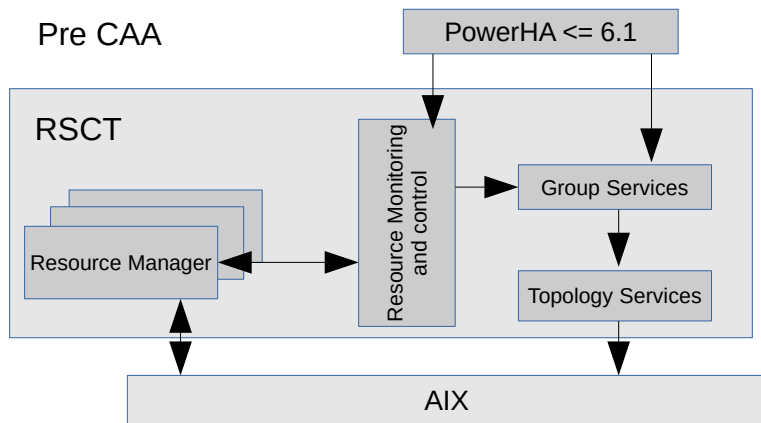
# IP distribution preferences

- Collocation
  - All Service labels will be on the same “adapter”
- Collocation with persistent
  - all service labels will be on the same “adapter” as the persistent IP.
- Collocation with Source
  - all service labels will be on the same “adapter” and the customer can choose the source IP of the outgoing packets
- Anti-collocation
  - all resources of this type will be allocated on the first “adapter” which is not already serving (or serving the least number of) addresses
- Anti-collocation with 1<sup>st</sup> Source
  - Same as above with the service IP being the source address of all outgoing packets.
- Anti-collocation with Persistent Labels
  - service labels will almost never be on the same “adapter” as the persistent IP, that is, service will occupy a different interface as long as one is available, but if no other is available then they will occupy the same interface.
- Anti-collocation with Persistent Labels and Source
  - Same as above with all outgoing packets having the service IP as the source address.

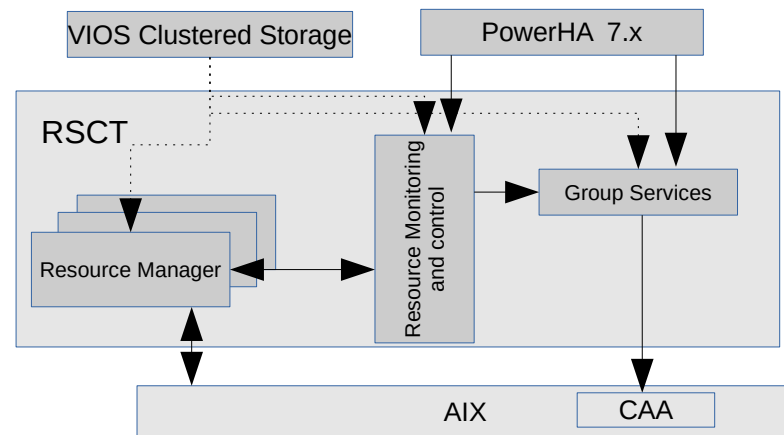
# Cluster Aware AIX

- Kernel based
- A set of services/tools embedded in AIX to help manage a cluster of AIX nodes and/or help run cluster software on AIX
  - IBM cluster products (including RSCT, PowerHA, and the VIOS) will use and/or call CAA services/tools.
  - CAA services can assist in the management and monitoring of an arbitrary set of nodes and/or running a third-party cluster.
- CAA does not form a cluster by itself. It is a tool set.
  - There is no notion of quorum. (If 10 nodes of a 11 node cluster are down, CAA still runs on the remaining node).
  - CAA does not eject nodes from a cluster. CAA provides tools to fence a node but by itself never fences a node and will continue to run on a fenced node
- Requires a repository disk (protected at the storage level)
- By default all interfaces monitored, not just those defined in PowerHA (lscluster -i)
- Debugging – use “snap caa” to collect PD data
- Can secure communications between nodes.

## Cluster Aware AIX (cont)




- Debugging
  - `snap caa`
- Logging via syslog
  - `lscluster -s` for stats
- List node and repository disk UUID
  - `lsattr -El cluster0`



- Lists valid cluster repository disks
  - `/usr/lib/cluster/clras lsrepos`
- Displays storage framework UUID for disks
  - `/usr/lib/cluster/clras sfinfo -d hdisk2`
- Displays contents of cluster repository disk
  - `/usr/lib/cluster/clras dumprepos`

# Default multi channel health management

- Simple setup
- Multiple channels of communication
  - Network
  - SAN
    - Supports virtual fibre
      - LPAR uses sfwcomm VLAN device which uses VLAN 3358 to communicate with VIO Server, Server passes packets out fibre to other nodes \*
    - Not all current fibre adapters support SAN heartbeat
  - Central repository disk
    - Contains matrix with state of all nodes in cluster
  - 3 independent / redundant modes of communication
- Critical to avoid spit brain / split cluster



\* IBM PowerHA SystemMirror V7.2 for IBM AIX Updates (Redbook)


# PowerHA DR options






# PowerHA SystemMirror Enterprise Edition

- Enterprise (DR), introduces storage replication management / hyperswap
  - Stretched
    - Essentially one cluster operating across and aware of 2 sites
  - Linked
    - Two clusters coordinating activity and responses to events




Multi Sites with Stretched Cluster



Multi Sites with Linked Clusters

# Managing Sites

- Site split and merge policies
  - Majority rule
    - Site with > half nodes or site with node with lowest ID wins
  - Manual
    - Administrator controlled
  - Tie breaker
    - Site holding tie breaker wins
- Tie breaker support introduced (third site SCSI reserver (FC; iSCSI); or NFS)
  - SCSI reservation FC/iSCSI disk or NFS file as the tie breaker
  - Losing side is quiesced
  - More suited for Linked Clusters



## LVM Split Site or Cross Site LVM

- Assumes SAN connected disks and nodes at two locations
- Define shared volume group with super strict mirror pools
  - Mirror pool for each location
  - Disks must be manually assigned to each mirror pool
  - Knowing which disks are where is a user responsibility
  - LVM mirrors logical volume between two locations
  - Resource group definition should allow forced varyon
- In the event of node and disk loss at one location
  - Volume group forced on line at other location by PowerHA
  - Mirror pool set up guarantees a local copy of the data
  - Manual recovery of repository using Repository Resiliency


- Risk analysis
  - Study the current environment. An example would be that the server room is on a properly sized UPS but there is no disk mirroring today.
  - Perform requirements analysis. How much availability is required and what is the acceptable likelihood of a long outage.
  - Hypothesise all possible vulnerabilities. What could go wrong?
  - Identify and quantify risks. Estimate the cost of a failure versus the probability that it occurs.
  - Evaluate counter measures. What does it take to reduce the risk or consequence to an acceptable level?
  - Finally, make decisions, create a budget and design the cluster.
- Include into your backup planning and backup infrastructure
- Plan for a test cluster
- Applications must be able to be started after crash
- Plan
  - Application data location can be shared
  - Location of application executables and configuration files

- Infrastructure planning
  - Power Redundancy; I/O Drawers; SCSI Backplane; SAN HBAs / Multipathing; Virtualised or Dedicated Deployments; Backup Strategies; Application Fallover Protection
- LPM
  - Live move of OS/Application between frames; Workload management; Energy management; Hardware management
- Partition Suspend/Resume
  - Resume where stopped; suspend low priority workloads; Firmware updates without stopping / restarting the application
  - VMRM HA and DR
- IBM Power RAS
  - Take advantage of Power Hot Swap features, built in redundancy and concurrent maintenance features.

- Designing High Availability
  - A spare should be available for every single hardware and software component that is required to keep application running
    - No 'Single Point of Failure'
    - Whilst a generally accepted principle, not always adhered to (cut to reduce cost)
    - 2 ports on one adapter is not removing a single point of failure
  - Perform a risk analysis, think through the effects of the failure of a each component, each possible combination and the desired response.
  - Monitor environment (have found unnoticed previous failures not acted on that reduced PowerHA's ability to respond to next failure)
- Use PowerHA features to assist
  - Automatic (regular) cluster verification – add your own customisation!
  - File collections
    - Keep your configuration files and AIX important files
  - PowerHA cluster password utility
    - maintain passwords across the cluster
  - Monitor
    - Use PowerHA logs to alert you of failures that may not have been visible at the time (as above)

## More PowerHA features that assist

- Are you aware of / using
  - Fast failure detection
  - Application monitoring
    - Startup, long running or both
    - Process or custom
    - Initiate fallover
  - Disk handling changes
    - Enhanced concurrent mode volume groups – now the default. Fast disk takeover
    - No longer required to manage SCSI reserves
  - C-SPOC
    - Manage disks, volume groups ..
  - Cluster Test tool
  - Exploitation of mount guard
    - JFS2 facility to prevent accidental 2<sup>nd</sup> mount
  - Nightly Verification of configuration



- Documenting the environment
  - Cluster Snapshots
  - GUI
- Ease of Deployment
  - Smart Assistants (SAP, DB2, Oracle...)
- Federated Security Integration
  - Encrypted File systems
  - RBAC
    - ha\_admin; ha\_op; ha\_mon; ha\_view

Remember that in the new versions of PowerHA, the developers used feedback from the field/PMRs to stop users making common mistakes

## Designing High Availability (cont)

- Remove single points of failure
  - Node
    - Using multiple nodes
    - Check PowerHA support matrix
      - <https://www.ibm.com/support/pages/powerha-aix-version-compatibility-matrix>
    - Many clusters are mutual takeover, few hot standby
    - Check common environmental factors (power, consistency)
    - Mirror rootvg
    - LPAR good sizing (ie sufficient cpu, memory and I/O to handle PowerHA while application at peak)
      - Planning must allow for worst case, include sufficient resources for CoD or DLPAR operations
  - Power source
    - Using multiple circuits or uninterruptible power supplies
    - Check the rack cabling
  - Virtualisation layer
    - Multiple VIO Servers
    - Multiple network and storage adapters in each VIO Server



## Designing High Availability (cont)

- Remove single points of failure (cont)
  - Network
    - Design for client access and application communication requirements
    - Plan multiple pathways / networks to connect nodes / clients
    - Be aware of environments subnet rules
    - Use /etc/netsvc.conf to protect against DNS problems
    - Persistent alias for node monitoring / administration
    - Single adapter networks – populate netmon.cf with “reliable” addresses
      - PowerHA needs to be able to determine adapter has network connectivity
      - Note it is recommended not to use gateways as they may drop ICMP packets if under load
    - Understand that network packet from AIX will use first address on the adapter
      - PowerHA can assist with controlling this – collocation with source
  - Network adapter
    - Using redundant network adapters and bonding (ether channel etc)
  - TCP/IP subsystem
    - Using non-IP networks to connect nodes (heartbeat)
      - Repository disk
        - Have a backup repository disk
      - Use SAN heartbeat if supported by your adapter

## Designing High Availability (cont)

- Remove single points of failure (cont)
  - Storage
    - Multiple Storage systems / built in redundancy
    - Accessible by all nodes
    - Monitor for internal errors, including AIX errors, check for failed paths/components
    - Avoid nested file systems
    - Use Enhanced Concurrent Volume Groups
    - Enable fast failure detection New disk heartbeat allows fast failure detection (dump process writes to disk)
    - LUNs built on multiple physical disks with mirroring RAID
  - Disk adapter
    - Using redundant disk adapter, multi-path hardware
  - Application
    - Adding node for takeover
    - Use application monitor(s)
    - Good scripting (more later)


## Designing High Availability (cont)

- Remove single points of failure (cont)
  - Site
    - Replicate infrastructure at 2nd site
    - Distance > range of your likely threats
    - Sufficient resources to keep your system running for production rebuild
      - particularly data centre connectivity
    - All required infrastructure available at DR site to maintain application
    - Run applications at both sites / run ½ year at each site
  - Administrators
    - good documentation 'clear' design
    - training
      - PowerHA training
      - AIX Administration, LVM and Storage environment
      - TCP/IP
      - Able to give cluster perspective to application owner
      - Able to give cluster perspective to change management process
    - Access to site(s)



## Designing High Availability (cont)

- Remove single points of failure (cont)
  - Cluster testing plan
    - actually use it
    - Post incident review to feed back into design and plan

Check IBM PowerHA support matrix and FLRT



# Notes if using HBA heartbeat



- Application controllers are scripts
  - Can be started in two modes:
    - Foreground (HA event processing will wait for the completion of the application start script)
    - Background
  - Poorly designed scripts may cause hangs (config\_too\_long in HA events)
  - With earlier versions of PowerHA, the return codes were not checked, now will cause “EVENT ERROR” if return code is not 0
  - Use good script writing
    - Log all details, use time/date stamps, trap errors (to hacmp.out?)
    - Don't grep for process name and kill
    - Adequate alerting, add notifications (email, page, SMTP traps, integrate with other monitoring tools)
    - If nesting, keep logging and check return codes
    - Keep scripts consistent across the cluster, use hostname / code path for local variations
    - Use change control and keep consistent across the cluster
  - Must be able to recover/start application after a crash without operator intervention
  - Check the application is not already running (will be run when starting from unmanaged state if no monitors configured)

## Application controllers (cont)


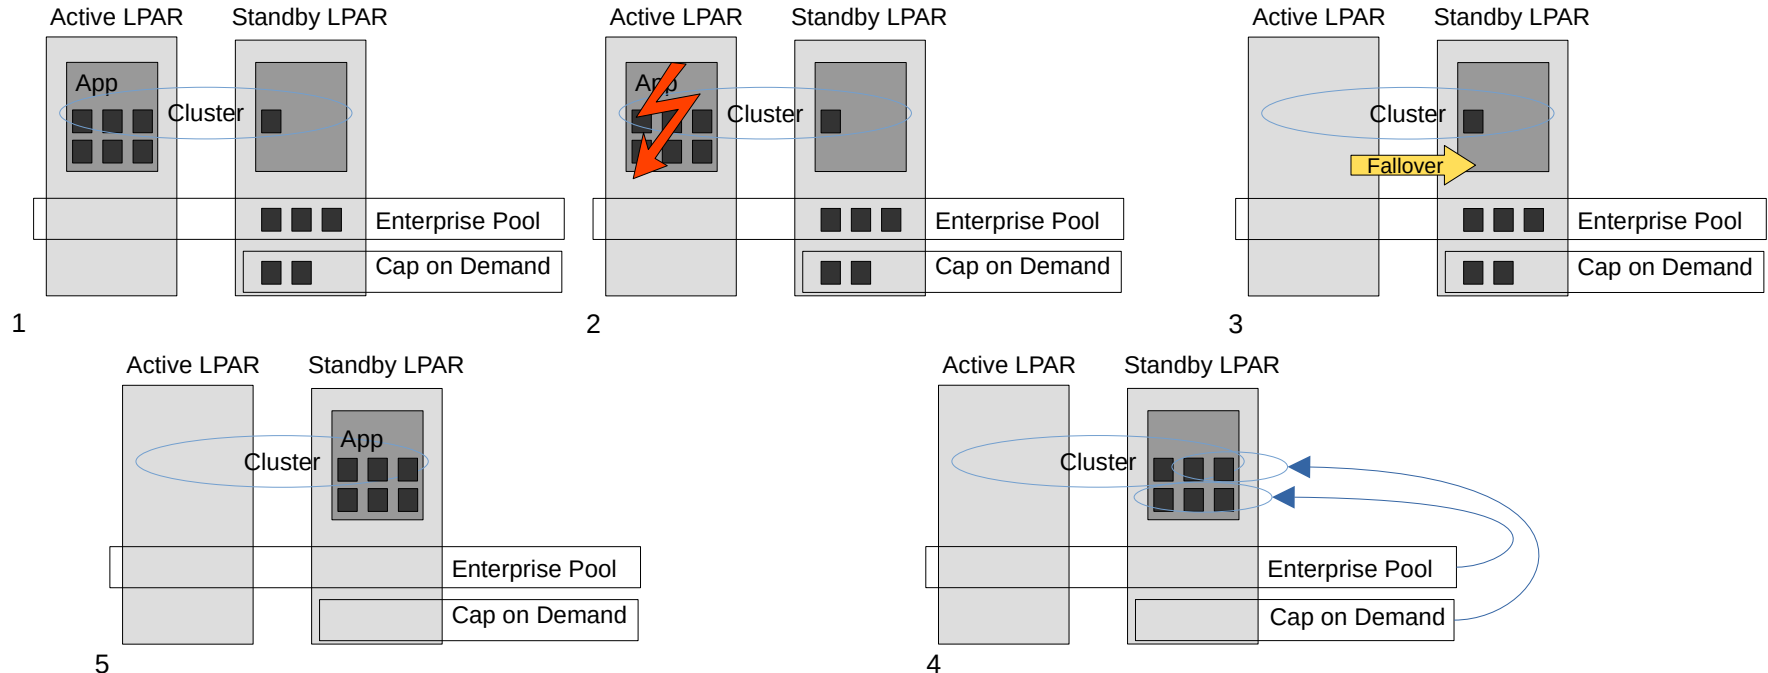
- Application controllers (cont)
  - Should check the environment before starting the application (all resources available – disks, file systems, IP addresses, prerequisite services and the state of the data)
  - Must be able to successfully stop the application
  - On stopping the application ensure that the application is stopped and all associated process cleaned up (child processes may be using disk resources)
    - Check to the point that the file system can be unmounted (fuser)
  - Exiting with a non-zero exit code will stop processing (config\_too\_long)
- Understand dependencies (remember you are in a clustered environment)
  - Those that PowerHA can control
  - Configuration files etc (HA can assist with consistency)
  - User information, passwords(HA can assist)
- Application resourcing, managing licenses
  - Move resources with the application
    - PowerHA can use CoD / DLPAR operations to ensure sufficient resources available for application.
    - Large allocations of memory may be slow, breakdown into smaller chunks
  - Check licensing requirements for multiple node operation

- PowerHA uses application monitors to check on the application health
  - Each Application controller can have 0 or more monitors (at least 1 is strongly recommended)
  - Application monitor can trigger a response (goal is for self healing clusters)
    - Application restart
    - Notification
    - Resource group fallover
  - See the monitors in the Smart Assists for examples
  - The application monitor is bound to the application controller and there are two types:
    - Process Monitor
      - Checks the process table (entry type and number)
    - Custom Monitor
      - Your own script to check on the application, for example, read for a record in a table



- Application monitors (cont)
  - Three modes:
    - Startup Monitor
      - Invoked when application started and used to confirm application running
        - Required by other dependencies, for example when to start the child resource group
      - Keeps checking until return code 0
    - Long running Monitors
      - Run at regular intervals to confirm application is still running, else trigger a response
    - Both
      - same monitor for checking startup and long term running

- Application provisioning
  - Ability to move CPU and memory resources with the application
  - PowerHA Enterprise Pool and CoD integrated management
    - Significant cost savings in both hardware and software



# Maintenance, testing and monitoring



## Maintenance - you cannot let sleeping clusters lie

- Why touch the system ?? – has been working now for over 3 years...
  - Hardware may need to be upgraded (6 monthly f/w update – 1/year may not be concurrent).
  - Replacement hardware may be at unrecognisable firmware levels..
  - Application may need to be upgraded, which may require new software levels or fixes
  - Either the OS and/or application out of support
  - Business expands
  - Lots of customers have helped IBM improve the code
  - Possibly longer outages and more complex recovery paths.
  - PowerHA designed to manage/support upgrade process (Shawn has a number of excellent videos covering the process)
    - Rolling upgrades
    - Snapshot conversions
    - PowerHA supports AIX Live Update
  - cl\_ezupdate tool designed to help administrator with software upgrade management tasks.
    - Automatic update of the entire cluster or a subset of its nodes in most cases without disruption of access to the application.
    - cl\_ezupdate has a rollback option

Shawn's videos: <https://www.youtube.com/channel/UCK-QtfzCBTQRihxQLjfDVHg>



- Upgrade planning
  - Check that cluster and application are stable and that the cluster can synchronise cleanly
  - Take a cluster snapshot and save it to a temporary non cluster directory (export `SNAPSHOTPATH=<some other directory>`)
  - Save event script customisation files / User Supplied scripts to a temporary non cluster directory. If you are unsure that any custom scripts are included, check with `odmget HACMPcustom`.
  - Check that the same level of cluster software (including PTFs) are on all nodes before beginning a migration
  - Ensure that the cluster software is committed (and not just applied).
- Plan upgrade and maintenance of rest of infra
  - HMC, PowerVC
  - Firmware
  - Virtualisation layer (VIO Servers; NovaLink)
  - AIX
  - PowerHA

## Performing upgrades

- For all AIX, PowerHA or application changes
  - Test in test cluster (test the whole process as would be done in production)
  - take cluster snapshot
  - change standby node
    - PowerHA will run with mixed levels, but there can be no changes to configuration
    - If no standby node, fail resources to another node
  - Move resource group over to upgraded node and test
    - Any problems, fail back to previous node
    - No problems, test for a reasonable period of time
  - Roll through the cluster
  - or depending on code levels use the easy update tool
- Have a regular maintenance schedule, test plan and build faith in the cluster

## Upgrade considerations

- Check the PowerHA documentation and readme
- Non-Disruptive Upgrade functionality is NOT available to get to 7.X
  - Can perform non-disruptive upgrade to load patches, for example:
    - From 7.1.0.1 to 7.1.0.4
- Migration to 7.X releases is quite different
  - Migration is disruptive
  - Migrate snapshot and apply
  - Requires some reconfiguration of cluster topology
- Options for migrating from older releases to recent levels
  - May require a couple of steps
  - May be easier to rebuild test cluster and switch.

- Change control planning
  - Is the change necessary?
  - How urgent is the change?
  - How important is the change? (not the same as urgent)
  - What impact does the change have on other aspects of the cluster?
  - What is the impact if the change is not allowed to occur?
  - Are all of the steps required to implement the change clearly understood and documented?
  - Was the change well tested and documented?
  - How is the change to be tested and approved?
  - Often many parties involved, have they all been notified and approve?
  - Have the users been notified?
  - What is the plan for backing out the change if necessary?
  - Is the appropriate expertise be available should problems develop?
  - Does the maintenance period include sufficient time for a full set of backups prior to the change and sufficient time to back out and or a full restore should the change fail testing?
  - Has the system documentation been updated?
  - Collected system configuration (snaps) and PowerHA configuration (snapshots)



# Maintenance - AIX

- AIX environment – keep your firmware, PowerHA and OS current
  - Maintain the system
    - Microcode and Software (OS and application)
    - configuration (tuning)
      - don't wait till performance becomes a problem
      - understand your application vendor's recommendations
    - Tools
      - Subscription services, SUMA (Service Update Management Assistant)
        - Analyse, filter and download maintenance files.
      - compare\_report
        - Compare filesets on target system
      - NIM, Pre-requisite tool, power matrix
      - Microcode discovery, FLRT and inventory scout
  - For this to be effective - develop a routine, plan ahead, schedule regular outage periods throughout the year for maintenance
    - If no maintenance scheduled, run the cluster test plan

- Changes to storage – C-SPOC vs lazy update
  - When changes to disks, volume groups, file systems are made through C-SPOC, all nodes that share the resource group will be updated
  - If the change is made outside C-SPOC, PowerHA will detect that a change has been made, so will export and re-import the volume group.

- Test cluster prior to going live
  - Use cluster test tool or document your own plan
- Run cluster verification regularly
- Plan a regular test – book 6 monthly outages for tests/upgrades. Set the expectation with the organisation.
  - test all changes (test tool!)
    - configuration (adding/modifying storage, changing application.)
    - software
- Use a test cluster, use a staging cluster
- Have a well documented test plan, that includes
  - Orderly startup and shutdown (without takeover) of each node in all possible order
  - Network adapter failures
  - Network failure
  - Test disk heartbeat communications
  - Test heartbeat disk replacement
  - Each node shutdown with takeover and then re-integration
  - Application failure

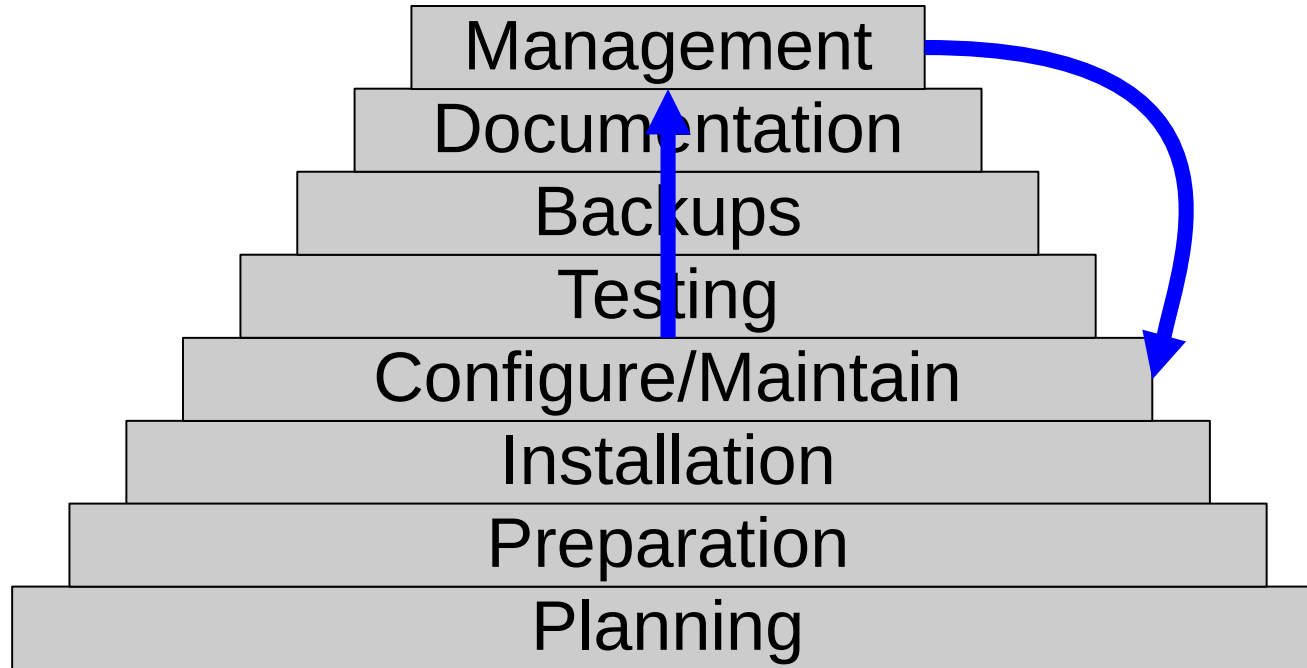
Send me an email if you would like an example test plan

## What to include in a test plan

- Have a well documented test plan (cont)
  - Failure of 1 disk subsystem component
  - Disk adapter failure
  - Site failure
  - Test PowerHA operation (optional)
    - Hung node
    - Hung application
    - Kill cluster manager
    - rootvg failure
    - partition cluster


# Monitoring

- Monitor the AIX error logs
- Use some monitoring product (nmon, njmon, ganglia, nagios...) to monitor
  - the environment
  - users
  - Performance
  - Don't wait for a problem, build a baseline
    - snap, perfpmr...
- Use the PowerHA notification methods for events that are associated with hardware or software failures
- Use AIX error notification
- PowerHA Commands
  - lssrc -ls topsvcs (heartbeat data)
  - lssrc -ls clstrmgrES (cluster manager and DNP info)
  - Cluster wide:  
`/usr/es/sbin/cluster/cspoc/cli_on_cluster -S "lssrc -ls clstrmgrES"`



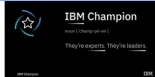
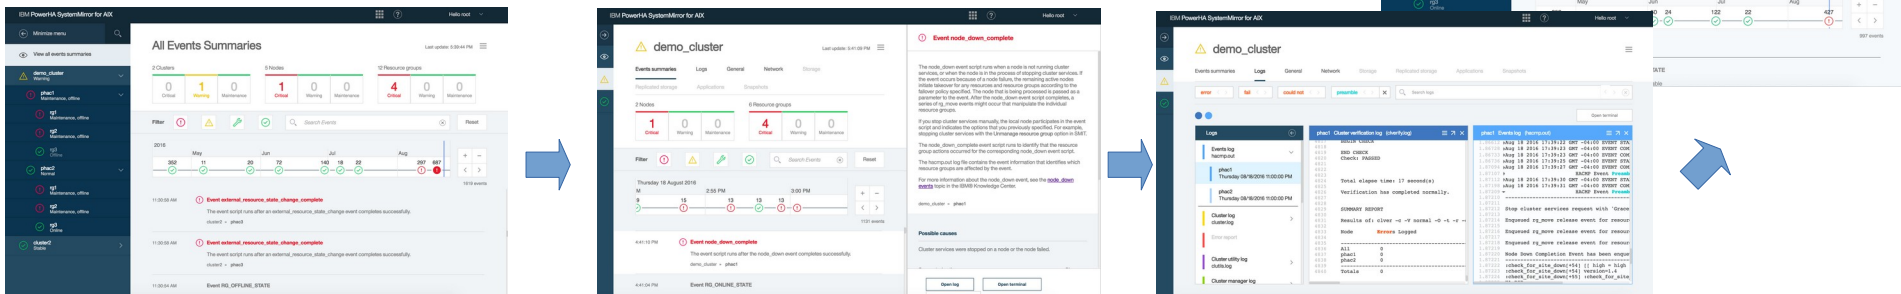
- GUI
- Cluster command line
- Cluster wide command
- Cluster test tool
- Application availability analysis tool
- AIX live kernel update compatibility
- Resource optimised high availability
- Other tools
  - File collections
  - Automatic cluster verification
  - Automatic Error Notification (can also be customised)
  - Auto-corrective/Self healing clusters
  - Custom Pager notification methods (including SMS)
  - Non-disruptive startup (create cluster around existing environment)
  - Cluster snapshots to save/restore clusters (XML format allows easy editing)

- Manage one or more clusters
  - creating, deleting, and modifying existing clusters and cluster configuration objects:
    - Resource groups (RGs)
    - Volume groups (VGs); File systems
    - Service IP addresses; Persistent IP addresses
  - Manage snapshots and clone clusters from snapshots
  - Application availability and monitoring metrics
  - Cross cluster verification
    - compare 2 cluster configurations
      - Create test cluster to test upgrade, ensure consistent with production
      - Manage an environment with many similar clusters (across regions)





- Manage one or more clusters
  - Health checking – cluster health at a glance
    - I 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?
  - Health summaries
    - determine which resource(s) had a problem
    - identify the problem
    - solve the problem
    - and watch as it returns to health state



## clmgr – cluster command line

- Replacement to Cluster Virtualisation Tool (CLVT)
  - Current Smart Assistants utilise CLVT
- Overcomes previous CLVT limitations
  - Limited trace output and logging; Difficult to use
- clmgr is a hard link to clvt, clvt is a binary
  - the *\*only\** binary in the clmgr code base
  - all other code is `ksh93 /usr/es/sbin/cluster/utilities/clmgr`
- Added 100% tracing coverage, with multiple levels
  - all STDERR output is written to `/var/hacmp/log/clutils.log`
- Fully globalised; uses the “command.cat” message catalog. Added dozens of consistent error messages, and large amounts of automatic help
  - Consolidated the set of supported actions and attributes

## clmgr – cluster command line (cont)

- Supported actions
  - add
  - delete
  - manage
  - modify
  - move
  - offline
  - online
  - query
  - recover
  - sync
  - view
- Supported object classes
  - cluster
  - site
  - node
  - interface
  - network
  - resource\_group
  - service\_ip
  - persistent\_ip
  - application\_controller
  - application\_monitor
  - dependency
  - file\_collection
  - fallback\_timer
  - volume\_group \*
  - logical\_volume \*
  - file\_system \*
  - physical\_volume \*
  - method\*
  - report
  - snapshot
  - tape

\* incomplete coverage of features

```
# clmgr online cluster WHEN=now MANAGE=auto BROADCAST=false CLINFO=true FORCE=false FIX=interactively
```

```
<snip>
```

```
s7801p22: Aug 12 2021 21:04:35 Checking for srcmstr active...
```

```
s7801p22: Aug 12 2021 21:04:35 complete.
```

```
s7801p22: Aug 12 2021 21:04:35
```

```
s7801p22: /usr/es/sbin/cluster/utilities/clstart: called with flags -m -G -i -B -A
```

```
s7801p22:
```

```
s7801p22:      Aug 12 2021 21:05:10
```

```
s7801p22: Completed execution of /usr/es/sbin/cluster/etc/rc.cluster
```

```
s7801p22: with parameters: -boot -N -A -i interactively -P cl_rc_cluster.
```

```
s7801p22: Exit status = 0
```

# Cluster wide execution

- Command is
  - `/usr/sbin/clcmd`
- Provided by CAA, distributes command to all nodes (or a subset of the nodes) in cluster (or clusters)
  - `clcmd [-n clustername] [-m node,node] command/script`
  - Similar to dsh

```
clcmd lssrc -g caa
-----
NODE s7801p22
-----
Subsystem      Group      PID      Status
cld             caa       6750432  active
clcmd          caa       7012576  active
clconfd        caa       7798794  active
solidhac       caa       6815926  active
solid          caa       8847410  active
-----
NODE s7801p20
-----
Subsystem      Group      PID      Status
cld             caa       5832952  active
clcmd          caa       6553816  active
solid          caa       7929910  active
solidhac       caa       8454150  active
clconfd        caa       8388622  active
```

# Cluster test tool

- Automated test plan
  - Important part of install process
  - Still important as regular procedure once in production
  - Many cluster administrators believe testing too time consuming and costly
  - Lack of testing leads to failures
  - Conducts a series of tests and then analyzes them
  - Will start all nodes, then preform node down with and without takeover on random nodes; network and application down. There are some limitations.
- Custom test procedure - user defined plan
- Designed to test the configuration, not the operation of the cluster manager
- Test available
  - Node tests
    - NODE\_UP: start one or more nodes
    - NODE\_DOWN\_FORCED: stop a node forced
    - NODE\_DOWN\_GRACEFUL: stop one or more nodes
    - NODE\_DOWN\_TAKEOVER: stop a node with takeover
  - Volume group test
    - VG\_DOWN: loss of VG

- Test available (cont)
  - Network tests
    - NETWORK\_DOWN\_LOCAL: stop a network on a node
    - NETWORK\_UP\_LOCAL: restart a network on a node
    - NETWORK\_UP\_GLOBAL: start a network on all nodes with interface on that network
    - NETWORK\_DOWN\_GLOBAL: stop a network on all nodes with interface on that network
  - Interface tests
    - JOIN\_LABEL, FAIL\_LABEL: Interface fail and join
  - Resource Group tests
    - RG\_ONLINE: bring a Resource Group online
    - RG\_OFFLINE: take a Resource Group offline
    - RG\_MOVE: move a Resource Group to next node in list, any node or back home
  - General tests
    - CLSTRMGR\_KILL: catastrophic software failure
    - SERVER\_DOWN: stop an application server
    - WAIT: pause testing

# Tests (cont)

- Test available (cont)
  - Site tests
    - RG\_MOVE\_SITE: Resource Group site move
    - SITE\_UP: site up
    - SITE\_DOWN\_GRACEFUL: site down graceful
    - SITE\_DOWN\_TAKEOVER: site takeover
    - SITE\_ISOLATION: site isolation
    - SITE\_MERGE: site re-integration

```

11/10/2019_07:20:24: -----
11/10/2019_07:20:24: | Validate NODE_UP
11/10/2019_07:20:24: -----
11/10/2019_07:20:24:     Event node: ALL
11/10/2019_07:20:24:     Configured nodes: ha1 ha2
11/10/2019_07:20:24:     Event 2: NODE_DOWN_GRACEFUL:
      NODE_DOWN_GRACEFUL,node1,Stop cluster services gracefully on a node
11/10/2019_07:20:24: -----
11/10/2019_07:20:24: | Validate NODE_DOWN_GRACEFUL
11/10/2019_07:20:24: -----
11/10/2019_07:20:24:     Event node: ha1
11/10/2019_07:20:24:     Configured nodes: ha1 ha2
11/10/2019_07:20:24:     Event 3: NODE_UP: NODE_UP,node1,
      Restart cluster services on the node that was stopped
11/10/2019_07:20:24: -----
11/10/2019_07:20:24: | Validate NODE_UP
11/10/2019_07:20:24: -----
11/10/2019_07:20:24:     Event node: ha1
11/10/2019_07:20:24:     Configured nodes: ha1 ha2
11/10/2019_07:20:24:     Event 4: NODE_DOWN_TAKEOVER:
      NODE_DOWN_TAKEOVER, node2,Stop cluster services with takeover on a node
  
```



# Application availability analysis tool

- SMIT / CLI
- GUI
  - Availability metrics
  - The availability metrics feature is enhanced to provide graphical information about the node, resource group, resource, and application start time and stop time. You can also view the amount of CPU and memory that is used by each application. This information can help you estimate the downtime during a planned outage, or the recovery time after a failure.

## Application Availability Analysis

Type or select values in entry fields.

```

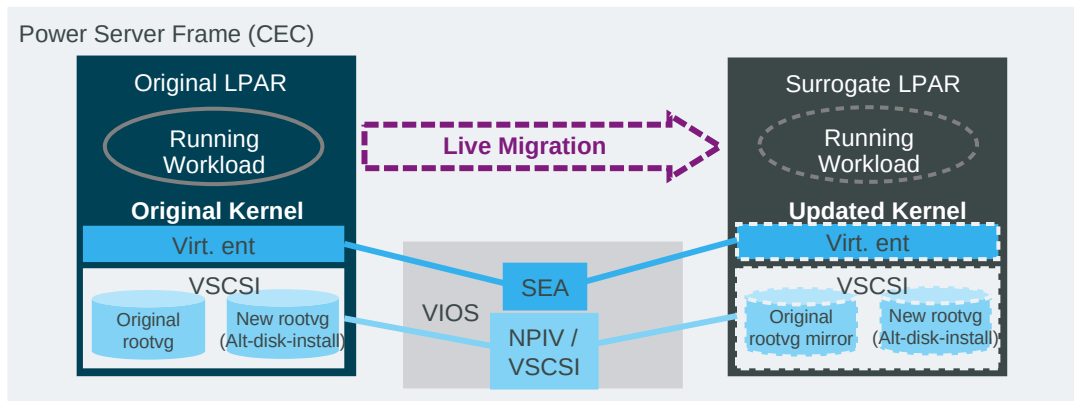
Press Enter AFTER making all desired changes.      [Entry Fields]
* Select an Application                             [test_appl01] +
* Begin analysis on YEAR (1970-2038)                [2019] #
* MONTH (01-12)                                     [01] #
* DAY (1-31)                                         [01] #
* Begin analysis at HOUR (00-23)                    [00] #
* MINUTES (00-59)                                   [00] #
* SECONDS (00-59)                                   [00] #
* End analysis on YEAR (1970-2038)                  [2019] #
* MONTH (01-12)                                     [06] #
* DAY (1-31)                                         [30] #
* End analysis at HOUR (00-23)                       [23] #
* MINUTES (00-59)                                   [59] #
* SECONDS (00-59)                                   [59] #
  
```

```

Analysis begins:           Saturday, 01-Jan-2019, 00:00
Analysis ends:             Thursday, 30-June-2019, 23:59
Application analyzed:      test_appl01
Total time:                 180 days, 23 hours, 59 minutes, 59 seconds
Uptime:
  Amount:                   180 days, 22 hours, 58 minutes, 29 seconds
  Percentage:                99.97%
  Longest period:            98 days, 16 hours, 48 minutes, 3 seconds
Downtime:
  Amount: 0 days, 1 hours, 1 minutes, 30 seconds
  
```

## AIX 7.2 - AIX Live Update (PowerHA compatible)

- Allows live update of the AIX Kernel and device drivers via i-fix without a reboot
  - No restrictions on the types or number of interim fixes (i-fixes)
  - Uses existing maintenance model and administrative tools (NIM, emgr, geninstall)
  - Introduced in: AIX 7.2 TL0 SP0



### Additional Software Requirements

- Requires Power firmware 730\_66 or later, VIOS 2.2.3.50 or later
- PowerHA 7.2 & RSCT 3.2.1.0 are required for PowerHA support
- PowerSC 1.1.4.0 for environments using PowerSC

### Additional Information

- Deeper dive into AIX Live Update  
<https://ibm.biz/BdHFKw>
- Hands on AIX Live Update article  
<https://ibm.biz/BdHFTd>

- File collections
  - use to keep important files consistent across cluster (if host dependant, test hostname, separate code path)
  - Default sets contain important AIX and PowerHA configuration files
- Automatic cluster verification and monitoring of configuration
- Automatic Error Notification (can also be customised)
  - Has two methods
    - cl\_failover is assigned if a disk or a network interface card is determined to be a single point of failure and its failure causes the cluster resources to fall over.
    - cl\_logerror is assigned for all other error types (to hacmp.out)
  - In conjunction with application monitoring => self healing clusters
- Custom Pager notification methods (including SMS)
- Non-disruptive startup (create cluster around existing environment)
- Cluster snapshots to save/restore clusters (XML format allows easy editing)

# Problem determination




## Problem determination

- PowerHA problems can be generally divided into two main areas
  - Setup / configuration
  - Operations / ongoing
- Setup problems
  - Keep up with changes in PowerHA operation (Application start scripts return codes checked, use enhanced concurrent volume groups...)
  - Do your application scripts keep logs, support debug?
  - Not testing the application starting on all nodes
  - LVM Configuration
    - Incorrect mirroring
    - Quorum left on
    - Volume group auto activated
    - VGDA and/or /etc/filesystems info different across nodes in the cluster
    - Different major numbers for NFS cross-mount
    - Mounting defined through AIX, not just PowerHA
    - Not all nodes being able to access all storage devices

## Problem determination (cont)

- Setup problems (cont)
  - PowerHA not controlling NFS (if used)
  - Resource group missing required resources
  - Failure detection rate (sensitivity) changed for 1 network only
  - User's environment not set up properly / file ownership permissions differ across nodes
  - Changes not synchronized successfully
  - Cluster ID not unique (cloned clusters)
  - Start PowerHA automatically
  - If using dual adapter and dual switch combination, make sure the PowerHA logical networks match
  - Application owners need to understand clustered environment
  - Using DNS for cluster addresses (/etc/hosts; local=hosts,...)
  - Not using slew option with time daemon (ie time going backwards)
  - No Application monitoring (both startup and long running)
  - Inconsistent scripts across the cluster.




## Problem determination (cont)

- Ongoing problems
  - Hardware not supported
    - Hardware replacement part may be different / incompatible firmware level
  - Not responding to a config\_too\_long and script failures
  - Availability outside the cluster
    - Power, network components, staff
  - Application cannot survive a crash
  - DGSP, partitioned cluster and rejoin
  - Storage considerations:
    - Clean-up after failed disk
    - Plan zoning and fabric redundancy, consider throughput requirements if there is a failure
    - Loss of rootvg used to be an issue, with PowerHA 7.1 the rootvg is monitored
    - Booting clustered node from diag (locks as devices tree walked)

## Some network issues

- All interfaces on a given HA network must be on the same physical network and can route packets between them and receive responses without any interference
- Make sure that none of the network devices between the nodes do anything but pass packets through transparently
  - Proxy ARP and MAC caching or ICMP optimisation
- Some switches interfere with the ARP response, in these cases:
  - Turn off spanning tree, portfast, uplinkfast and backbonefast.
  - If necessary to have spanning tree on, portfast should also be on
- If you are using ether channel (bonding or nib) – add entries into netmon.cf
- Set adapter speed where possible, don't use auto negotiate (less of a problem now)
- If using VIO – check all adapters on VIO speed set correctly







## What to do in the (unlikely!) event of a problem

- Check the logs (GUI makes this task simpler)
- Collect a detailed description of events and times. Any information from the users (perceptions may be useful)
- Application logs
- What has changed recently (Application, AIX, Config, LVM, networks, hardware..?)
  - Change Control!
- Test ability of staff to open support calls and upload log files (collect base snap)
- Post event review and modify scripts / procedures if required
- Data to collect (collect early as logs rotate)
  - HA log files and snapshot (snap -e)
  - System fileset and fix/efix levels (snap -g)
  - System error log and diag output (snap -g)
  - Hardware configuration and m/c levels (snap -g)
  - Cluster aware AIX (snap caa)
  - System hung or crashed? (snap -d)
  - Network (snap -t)
  - Disk/VG/LV/filesystem (snap -Lf)
  - Storage unit (configuration)

Log name	Location	size	Description
cluster.log	/var/hacmp/adm		Good high level view, output from PowerHA scripts and daemons
cluster.mmddyyyy	/var/hacmp/adm/history		daily history file
hacmp.out	/var/hacmp/log	rolled .1-.7	Output of all resource group activity with summaries
grpqlsm	/var/ha/log	trimmed	RSCT group services messages
clverify.log, working files	/var/hacmp/clverify	Previous successful and unsuccessful verification attempts.	Previous successful and unsuccessful verification attempts.
clavan.log	/var/hacmp/log	Application availability logs	Details of when each application is started or stopped
clpol	/usr/es/sbing/cluster/etc	Priority override data	Priority override data
clutils.log	/var/hacmp/log	cluster utilities and file propagation	cluster utilities and file propagation
cl_testtool.log	/var/hacmp/log	rolled	Cluster test tool logs and hacmp.out summaries for each step
autoverify.log	/var/hacmp/log	autoverification logs	autoverification logs
discovery data	/usr/es/sbin/cluster/etc	cllip_config	cllip_config
sa.log	/var/hacmp/log	Cluster discovery logs	Smart assist logs
clcomd.log / clcomddiag.log	/var/log/clcomd	Logging for the clcomd daemon	Logging for the clcomd daemon
clinfo.log	/var/hacmp/log	rolled .1-.7	output from event scripts
cspoc.log   .long   .remote	/var/hacmp/log	Summary of all cluster activity	messages and output from all CSPOC commands
clstrmgr.debug   .long	/var/hacmp/log	rolled .1-.7	



Log name	Location	Description
smui-server.log	/usr/es/sbin/cluster/ui/server/logs/	Information about the PowerHA GUI server
smui-agent.log	/usr/es/sbin/cluster/ui/agent/logs	Information about the PowerHA GUI Agent
notify-event.log	/usr/es/sbin/cluster/ui/agent/logs/	Information about events sent from the Agent to the GUI Server



# Get PowerHA certified



- A candidate pursuing the PowerHA SystemMirror Administrator Specialty has knowledge and experience deploying PowerHA clusters, both within and across data centres, and a good understanding of various deployment methods. They should have hands-on experience administering the AIX operating system, have a good understanding of SAN and network infrastructures, and have some practical experience in deploying applications that are controlled by the cluster software.
- The candidate is a technical professional who can:
  - Install, deploy, and configure PowerHA clusters on AIX.
  - Identify and troubleshoot issues in a deployment at the various levels of the solution: OS, RSCT, CAA or PowerHA cluster level.
  - Tune or customise the solution to achieve desired behaviour, including operating system, CAA, or PowerHA settings.
- Details of exam
  - Number of questions: 40
  - Number of questions to pass: 26
  - Time allowed: 75 mins.

<https://www.ibm.com/certify/specialty?id=S1000-008>

# Review of updates



- Dec 2018 PowerHA SystemMirror for AIX 7.2.3
- Dec 2019 PowerHA SystemMirror for AIX 7.2.4
- PowerHA SystemMirror for AIX 7.2.5, Announced Oct 6 2020 and GA Dec 10
  - Geographic Logical Volume Manager (GLVM) configuration wizard
  - Configure and orchestrate multiple parallel GLVM instances, particularly for cloud deployments
  - Graphical User Interface (GUI) improvements
    - Support for cross-cluster verification
    - Manage administrative authorities' users through the GUI
    - Cluster deployment and management:
      - Removal of hostname dependency for cluster management
      - Bulk cluster notification enhancement
      - Activity log server hostname support
      - SSH server key implementation
  - Increased application resiliency during manage and unmanage activity
  - Cloud tie-breaker via lock file for on premise or cloud deployments
  - Smart Assist for PowerHA currency support for IBM Tivoli Directory Server, IBM Tivoli Storage Manager, Oracle, IBM Db2, and SAP NetWeaver

## Highlights from previous releases

- Resource optimised High Availability (ROHA)
  - Support for Enterprise Pool, CoD
  - flexible Capacity management across various Power pools
  - Enterprise Pool Integration
  - Manipulate Shared Processor Pool Sizes
  - New HMC Integration Tunables
- Advanced Split handling policies
  - Node halt & disk fencing quarantine policies
  - Disk & NFS tie breaker support
    - Flexibility to avoid the need of a NAS backed device as tie breaker disk
  - Manual Fallover
- Automated repository replacement
  - Define multiple repository disks & auto replacement behaviour with AIX 7.2
- Network & Node Failure detection improvements
  - Fine granular controls to change failure detection times
  - Improved network interface failure detection methods
- Non-Disruptive Upgrade Support for HA
  - Cluster software upgrade without disruption to workload
  - Easy PowerHA Migration without downtime
- rootvg disk failure detection improvements
- AIX Live Update Support & LPM Support Enhancements
  - Adapts to LPM operation
  - Live update for AIX supported, but only one node at a time.



## Highlights from previous releases (cont)

- Enhanced AIX and detailed cluster verification checks
  - Validation of a number of new checks including AIX Runtime Expert Settings (optional)
- Graphical Management
  - Cluster administration capabilities
  - Role based UI access control management
  - Multi tenancy support through HA zones
  - Consolidated management across AIX & Linux clusters
- Log Analyser tool
  - Analyse live or offline cluster FFDC information
  - Easy problem determination
  - Pattern based analysis support
- REST based HMC communication support
  - Setup for active node halt or ROHA policies using REST
- Novalink support (P8 and above)
- PowerHA (CAA) auto restart for VM restart DR
  - PVID based repository disk management
- Easy Update management
  - Local file system or NIM based update
- Quarantine Policies
  - Define Critical resource group
  - Split handling protections
  - HMC Node Halt Policy
  - SCSI-3 Node Fence Policy

## Highlights from previous releases (cont)

- **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.
- **Enhanced Split/Merge Policies**
  - Expanded choices of Split/Merge handling
  - Consistent policies for all types of clusters
  - Standard, Stretched, & Linked Clusters




	Pre-7.2.1 Split-Merge	7.2.1 Split-Merge
Std Cluster	None - Majority	None – majority TB(disk) – TB(disk) TB(NFS) – TB(NFS) Manual - Manual
Stretched cluster	None – Majority TB(disk) – TB(disk) TB(NFS) – TB(NFS)	
Linked cluster	None – Majority TB(disk) – TB(disk) TB(NFS) – TB(NFS) Manual - Manual	

- Miscellaneous tidy
  - Some customers requested use of rsyslog (AIX uses syslog by default)
  - C-SPOC 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 as PowerHA events.
  - 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 to be reached before responding

- 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.”
    - Health dashboard with solutions
  - The ability to take actions on the managed clusters was asked for, which implied a higher level of security and organisation
    - User management
    - Role-based access control
    - Cluster zones
- Changes introduced
  - 5 customers were involved in the research
    - Design thinking principles incorporated
  - Adds more administrative capabilities
    - Modify existing cluster configuration objects
    - Delete existing cluster configuration objects
  - Installation improved
  - Highly available

# PowerHA lifecycle

- PowerHA release Life cycle strategy
  - Implement technology level release strategy
    - Major releases as necessary
    - Minor release updates (Technology Level 0 to Major release)
    - At least two technology levels per major release
  - Considering
    - Additional 2 year service offering for last TL
- New command
  - `halevel -s`



## PowerHA plans

- Version 7.2.6 (aimed for Dec)
- Hope we will see
  - More GLVM features related to replication to the cloud managed by GUI
  - Security improvements in GUI
  - Increased scaling of cross cluster verification in GUI
  - Support for PowerVS
  - Support for LVM Encryption
  - Power10 qualification
  - Updates of smart assists and Ansible enhancements

# References



- IBM PowerHA SystemMirror for AIX 7.1.3 Best Practices and Migration Guide
  - <http://www.redbooks.ibm.com/abstracts/sg248234.html?Open>
- IBM PowerHA for AIX community
  - <https://community.ibm.com/community/user/power/communities/community-home?CommunityKey=32520692-50a7-4829-91c5-038ab4df00b1>
- PowerHA SystemMirror for AIX Reference Information
  - <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=reference-information>
- IBM Support – subscribe to notifications
  - <https://www.ibm.com/support/pages/node/718119>
- IBM PowerHA SystemMirror V7.2 for IBM AIX Updates
  - <http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248278.html?Open>
- IBM PowerHA SystemMirror V7.2.3 for IBM AIX SG24-8434-00
  - <http://www.redbooks.ibm.com/abstracts/sg248434.html?Open>
- Implementing IBM VM Recovery Manager for IBM Power Systems
  - <http://www.redbooks.ibm.com/abstracts/sg248426.html?Open>
- Basic Requirements for HA and DR (on-prem and Cloud) (published soon)



## Useful links (cont)

- PowerHA Release Notes
  - <https://www.ibm.com/docs/en/powerha-aix/7.2?topic=release-notes>
- PowerHA Comments & Questions:
  - [hafeedbk@us.ibm.com](mailto:hafeedbk@us.ibm.com)
- IBM FLRT
  - <https://www14.software.ibm.com/webapp/set2/flrt/power>

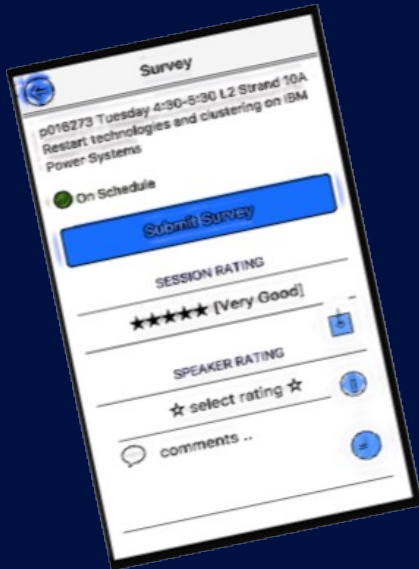


## Useful links (cont)

- Jaqui Lynch's good AIX practices and VIO maintenance
  - <https://www.youtube.com/channel/UCYH6OdgB6rV1rPxYt6FWHpw>
- Shawn Bodily's PowerHA implementation, management and upgrade
  - <https://www.youtube.com/channel/UCK-QtfzCBTQRihxQLjfDVHg>
  - Non-disruptive upgrade from 713 to 72
    - <https://www.youtube.com/watch?v=1Kzm7I2mRyE>
  - AIX Live update (PowerHA node)
    - <https://www.youtube.com/watch?v=BJAnpN-6Sno>
  - Automated Repository Replacement
    - <https://www.youtube.com/watch?v=HJZZDCXLwTk>
  - rootvg Failure monitoring
    - <https://www.youtube.com/watch?v=OZcrhVGvkBg>
- Michael Herrera's video on ROHA: Flexible Capacity Management
  - <https://www.youtube.com/watch?v=G-zefev-XYU>

# Thank you a203866 – PowerHA SystemMirror good practice

Antony (Red) Steel  
antony.steel@belisama.com.sg



**Please don't forget to complete the  
session evaluation!**

# Notices and disclaimers

© 2021 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

## **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

**Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**


Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those

customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.



- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)



# Backup slide

