Session: a111034 PowerHA 7.2.3 for AIX and Linux

Antony Steel

Thanks to:

Steven Finnes
Ravi Shankar
Paul Moyer
Dino and the Redbook team

ALTRAN

Power Systems

PowerHA™

#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.

Abstract

Belisama

- With the recent release of PoweHA 7.2.3 and the publication of the redbook, this is a good time to give an update on PowerHA new features. The session will also introduce AIX and LInux Admins to the operation of PowerHA and examples of how PowerHA assists with the management of your virtual environments - beyond maintaining their availability.

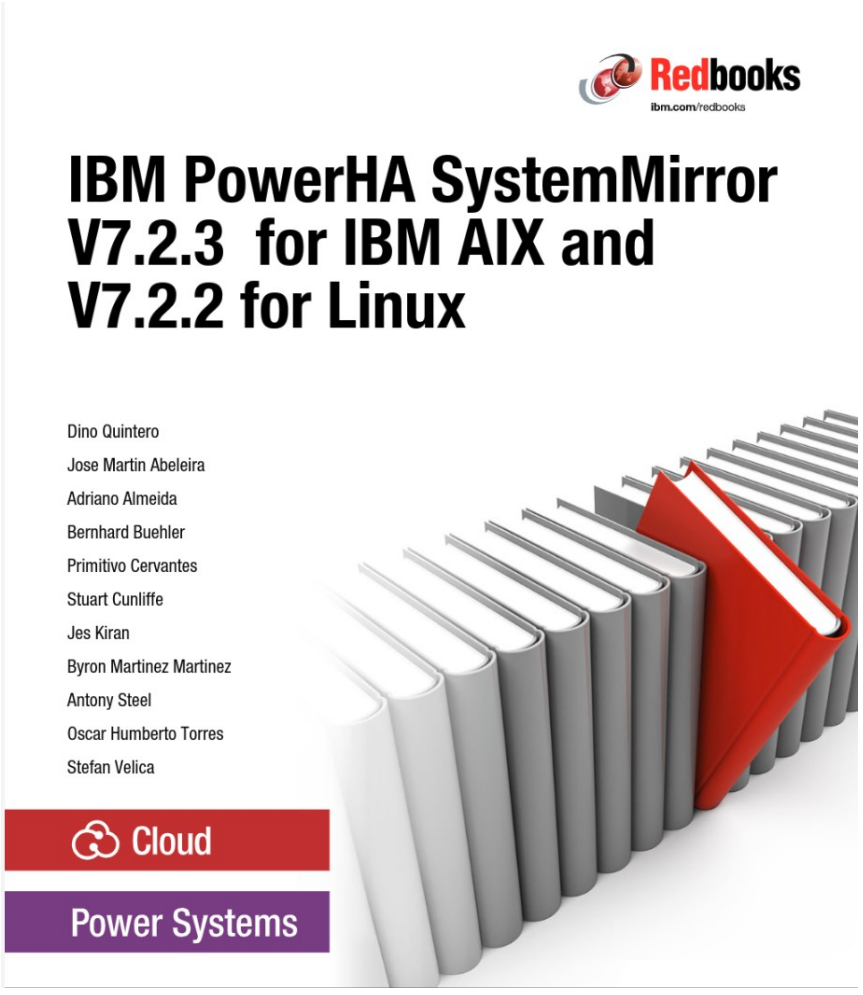

Agenda

Belisama

- Review/remind previous PowerHA release content
  - Spit/merge policy standardisation
  - Huge savings with Enterprise Pools and shared processor pool resizing
  - The CBU for Enterprise Systems
- The PowerHA what is new
- Introducing VM RM and other HA/DR options
- Resource materials

Introduction to High Availability

- PowerHA SystemMirror for AIX Standard Edition
- Integrated 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 (and Linux) – LVM mirroring; Capacity upgrade
  - Multi-cluster federated security, user management, encryption
- IBM shared storage clustering
  - Can enable near-continuous application service
  - Minimize impact of planned & unplanned outages
  - Ease of use for HA operations
- Smart Assists – application agents
  - Out of the box deployment for DB2, SAP, Oracle, Websphere and other popular applications
- Mature Product
  - 25 Major releases (averaging one a year)
  - Over 15,000 customers worldwide

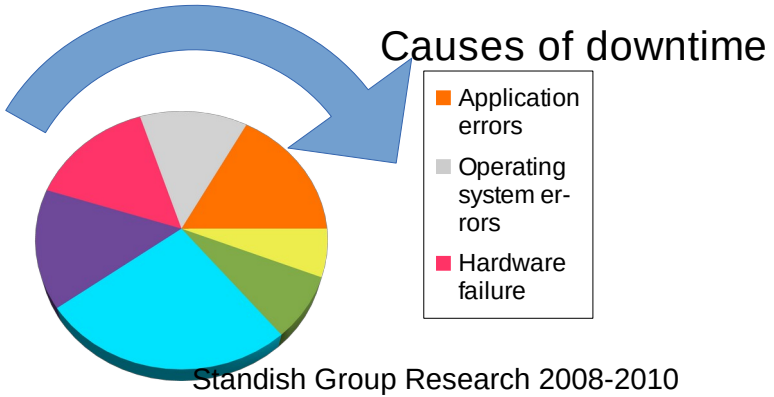

Standish Group Research 2008-2010

- PowerHA SystemMirror for AIX Enterprise Edition
  - Cluster management for the Enterprise
  - Multi-site cluster management
  - Includes the Standard Edition function
  - Low cost mirroring option supported (part of AIX)

PowerHA SystemMirror

- Standard Edition
  - High Availability within a Datacentre
    - Health management of the cluster
    - Service availability against failures in any component of the stack
- Enterprise Edition
  - 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

	Standard and Enterprise Edition
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	✓

	Enterprise Edition
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	✓

PowerHA editions

Belisama

- 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
  - 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 portion of fallover configuration
  - Provides local & extended cluster remote replication functions
- Can be configured to provide local clustering capabilities at first site and automated fallover 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

Clustered high availability solutions

- Cluster based High Availability
  - Business continuity by workload failovers
  - Provides for redundant operating environment
  - Automated workload bring ups
  - Environment specific outage actions
  - Planned and unplanned outage management
- Cluster Solutions
  - Two or more Servers in the cluster
  - Cluster management console and software
  - Communication to achieve health and synchronisation
  - Active-Passive and Active-Active models
- Disaster Recovery (DR) Management

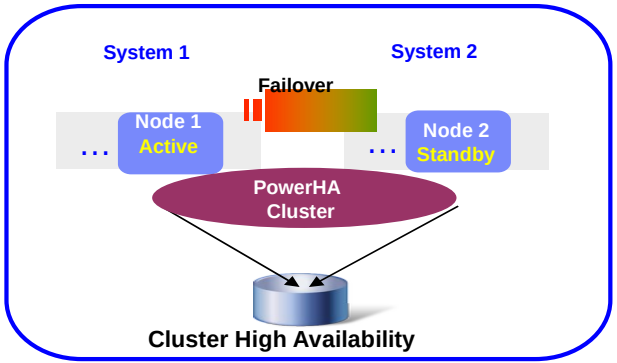

- Redundancy
- Planned activities
- Unplanned events
- S/W&H/W updates

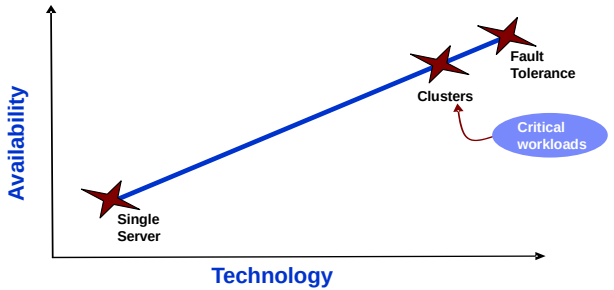

HACMP – PowerHA SystemMirror evolution

Belisama

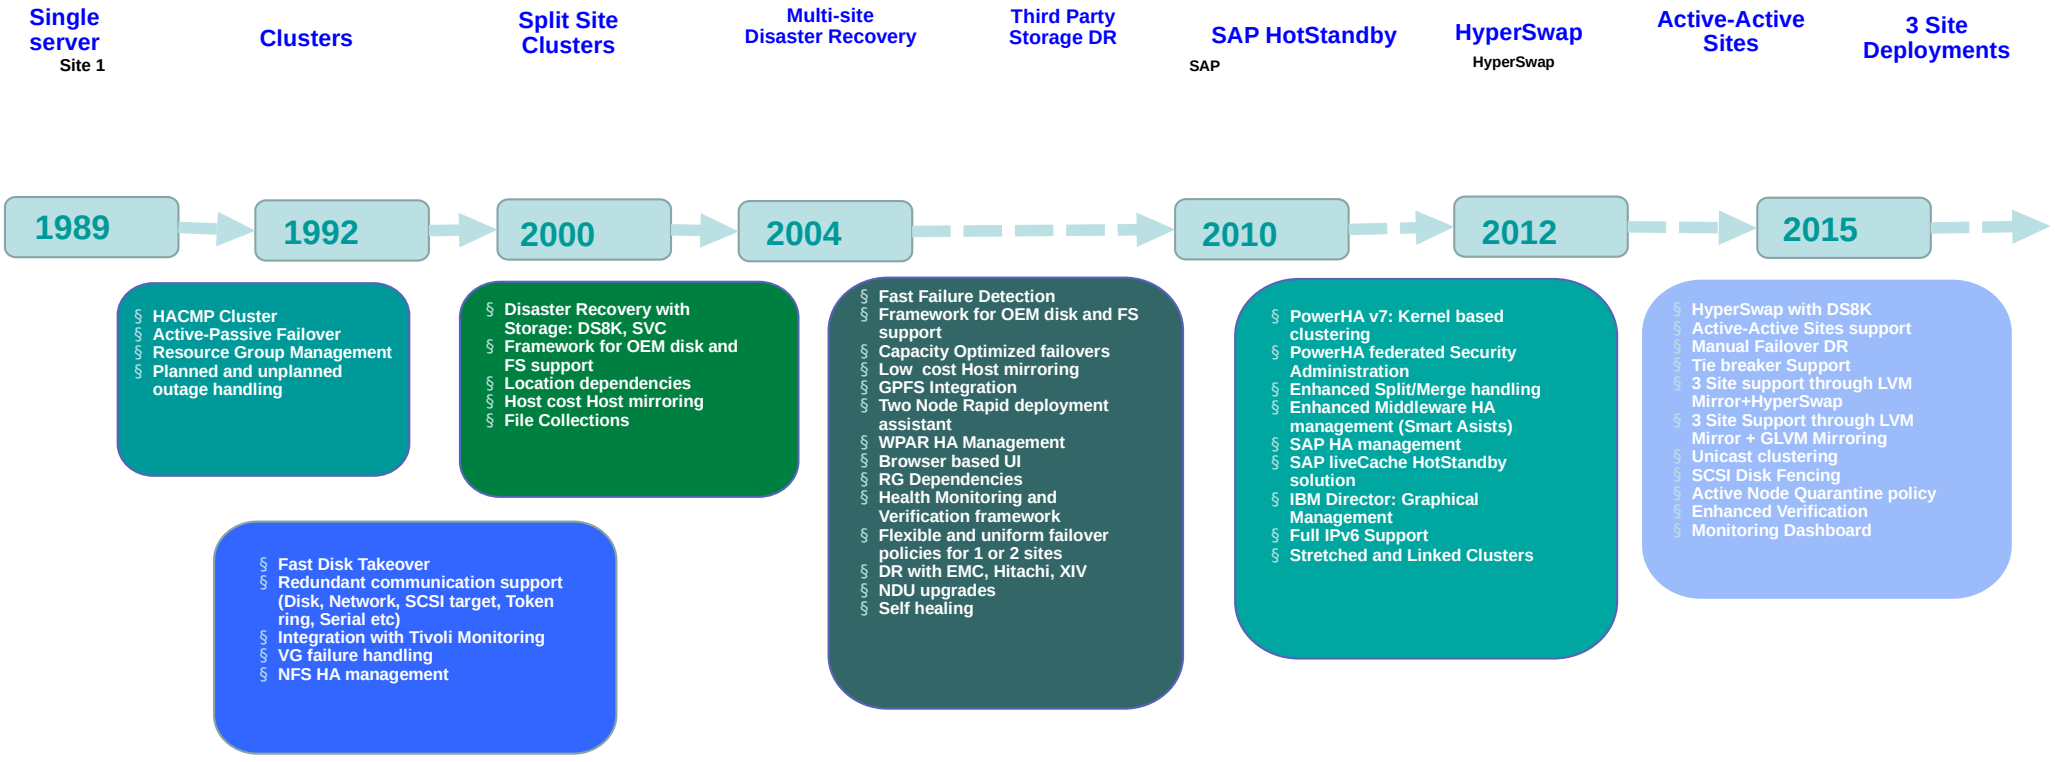

Summary of changes

- PowerHA 7.1.1
  - CAA Repository Resilience
  - JFS2 Mount Guard support
  - SAP Hot Standby Solution
  - Federated Security
  - SAP & MQ Smart Assists
  - XIV Replication Integration
- PowerHA 7.1.2
  - Cluster Aware AIX
    - IPv6, Rolling upgrade, Linked clusters
  - Enterprise Edition,
    - Linked and stretched clusters
    - Split / merge site options with tie-breaker
    - Hyperswap
      - Support for DS8k for 2 sites

PowerHA 7.1.x goes EOS 30-04-2018

- PowerHA 7.2 Updates
  - ROHA (Resource Optimised High Availability)
  - Non-Disruptive Upgrade Support (PowerHA code)
  - Automatic Repository Disk Replacement
  - AIX Live Update Support & LPM Support Enhancements
  - Cluster Detailed Verification Checks
  - (optional) Validation of a number of new checks including AIX Runtime Expert Settings
  - Quarantine Policies (Critical RG)
  - NFS Backed Tie Breaker Disk support
- Announce & GA planned for 1half 2017
  - PowerHA Health Dashboard support for earlier PowerHA releases
  - Agent support for PowerHA 713 and 720
  - Easy update management for PowerHA software stack
  - Update one node or all nodes in the cluster with ease
  - REST based HMC communication
  - Support for Enterprise pool management (ROHA) and Active Node Halt Policy in Cloud environments

When is PowerHA a good idea?

- High availability is:
  - The reduction to close to zero for downtime (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 single points of failure (SPOF)

- PowerHA is not the answer if
  - Cannot afford any downtime – life critical systems - Need a fault tolerant solution
  - Environment is not secure
    - Many users with root access
  - Then 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

Eliminate single points of failure by:

Belisama

- Node
  - Using multiple nodes
- Power source
  - Using multiple circuits or un-interruptible power supplies
- Network adapter
  - Using redundant network adapters and bonding (etherchannel etc)
  - Network Using multiple networks to connect nodes / clients
- TCP/IP subsystem
  - Using non-IP networks to connect nodes
- Disk adapter
  - Using redundant disk adapter or multipath hardware
  - Disk Using multiple disks with mirroring or raid
- Application
  - Adding node for takeover; configuring application monitor
  - VIO server Implementing dual VIO servers
- Site
  - Adding an additional site

Setting realistic expectations

- What is considered an outage in your environment?
- Unexpected downtime
- Maintenance Tasks

Availability	Downtime
90% (1-nine)	36.5 days/year
99% (2-nines)	3.65 days/year
99.9% (3-nines)	8.76 hours/year
99.99% (4-nines)	52 minutes/year
99.999% (5-nines)	5 minutes/year
99.9999% (6-nines)	31 seconds/year

- What are the desired:
  - RTO – Recovery Time Objective
  - RPO – Recovery Point Objective
- Environment
  - Well managed and change control
  - Scripts can start / stop / monitor applications

Building for availability

- Infrastructure planning
  - Power Redundancy; I/O Drawers; SCSI Backplane; SAN HBAs / Multipathing; Virtualized 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
- CHARM (CEC Hot Add & Repair Maintenance)
  - Available on high end models (>= *70)
  - Perform CHARM during low-use periods
  - LPM critical partitions to other servers if possible
  - Depending on the repair, IBM may recommend quiescing critical applications on running partitions
  - Have current backups before beginning, and make sure all configuration redundancy requirements have been met
  - Use PowerVM Suspend / Resume to reduce CPU and active memory

- Planned
  - Maintenance
  - Upgrades
  - Testing
  - Development
- Unplanned
  - User Error
  - Application Failure
  - Component Failure
  - Operating System Failure
  - Environmental Disasters

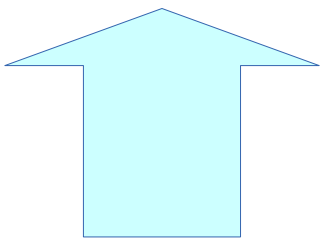

Becoming a more important area - PowerHA as an administration tool

LPM is an alternative for, … but not for (or software upgrades etc)

PowerHA will help to mask or eliminate

You cannot let sleeping clusters lie

- Why touch the system ?? – has been working now for 2 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
  - OS and/or application out of support
  - Business expands
  - PowerHA designed to manage/support upgrade process
    - Rolling upgrades
    - Snapshot conversions

High Availability options

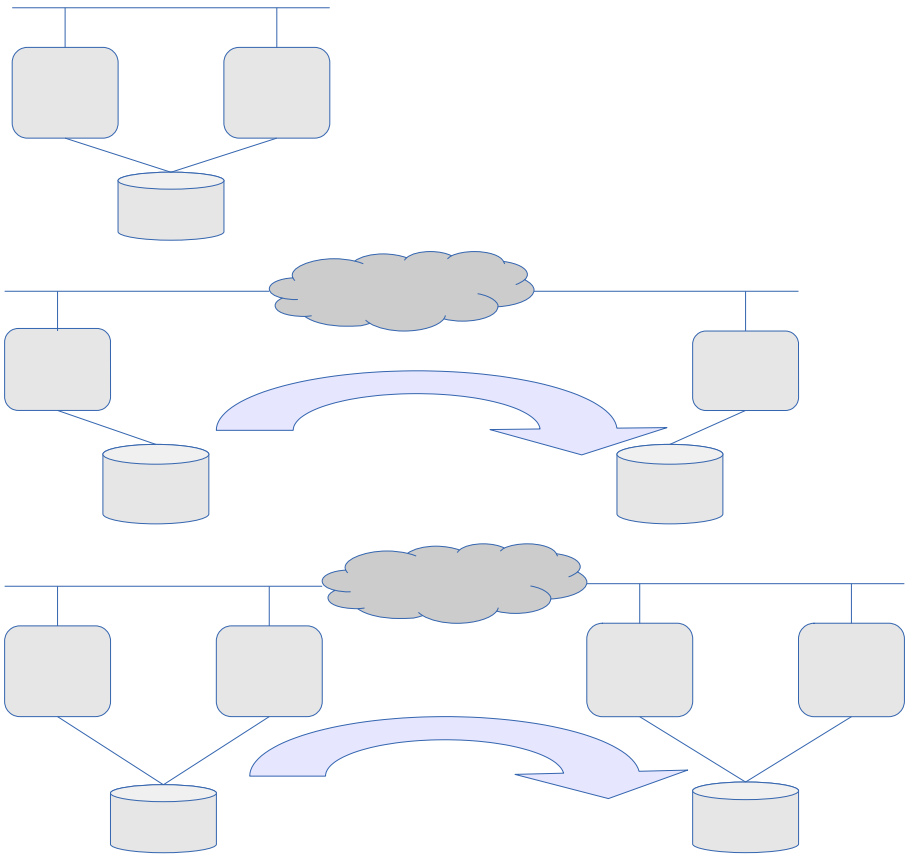

>> Planning and preparation

- One site – HA
  - PowerHA SystemMirror
    - Dual servers, shared storage
    - Site only single point of failure
- Disaster Recovery
  - Replication
    - GLVM
    - Storage / Database
    - SRR
    - VM Recovery Manager DR
  - PowerHA SystemMirror Enterprise Ed.
    - PowerHA managing application and storage replication
      - GLVM
      - SVC; Storewise; MetroMirror; GlobalMirror
      - EMC SRDF / Hitachi TrueCopy/HUR

Cluster topology

Belisama

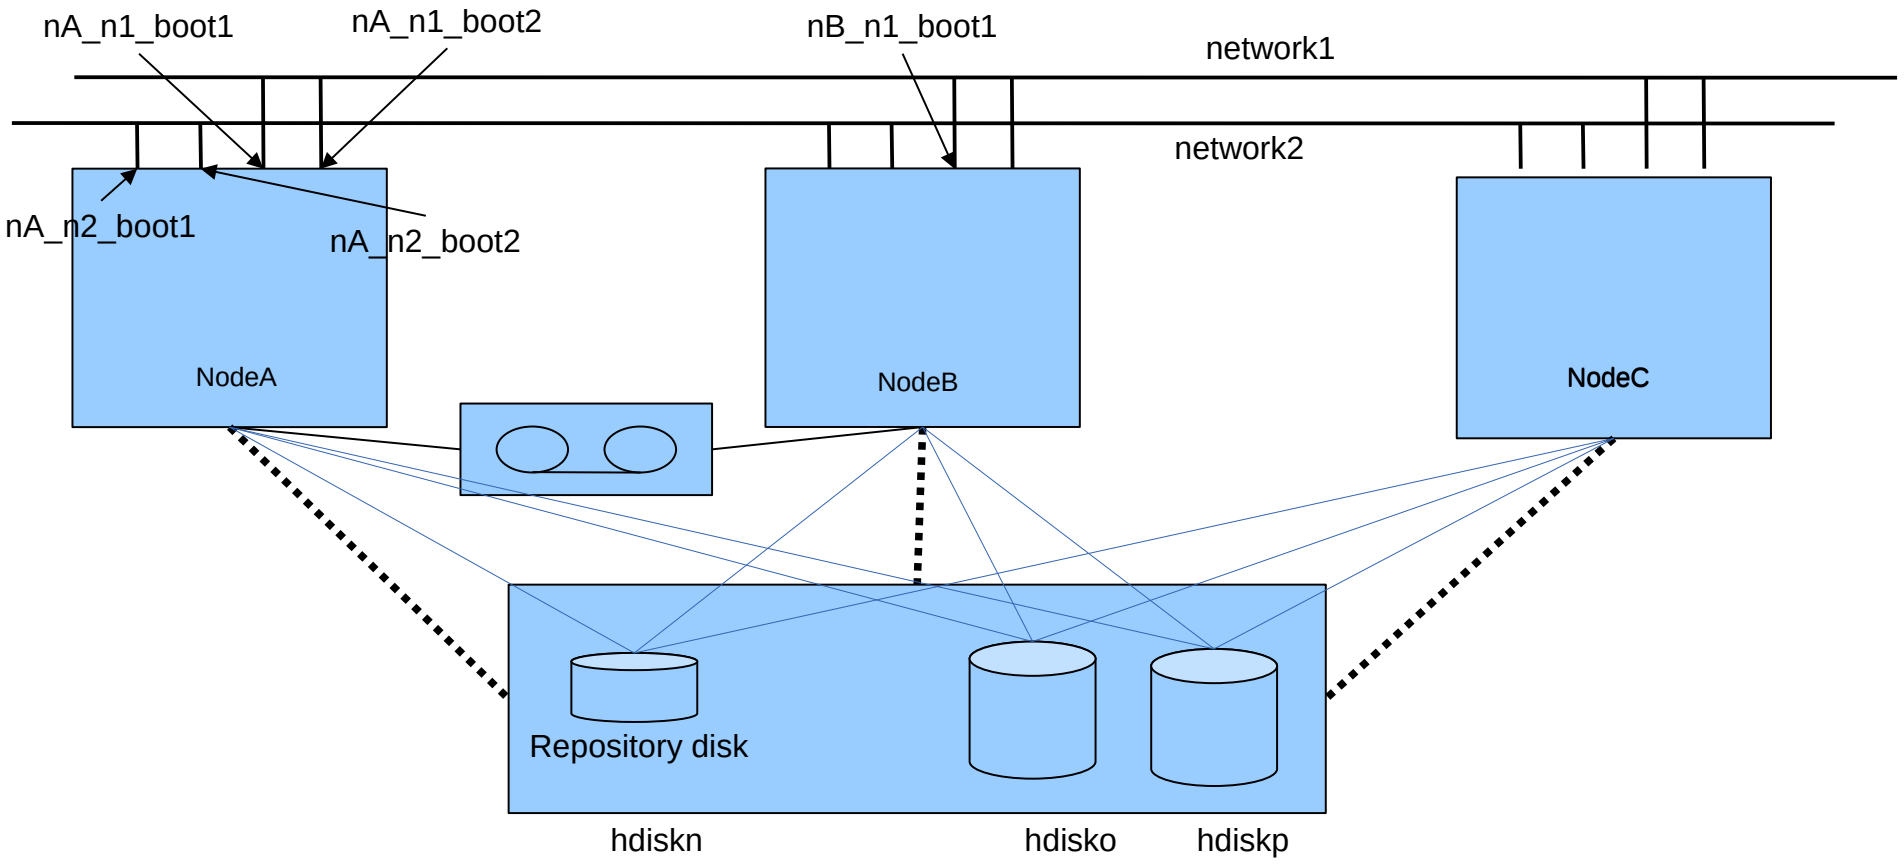

Cluster topology

Belisama

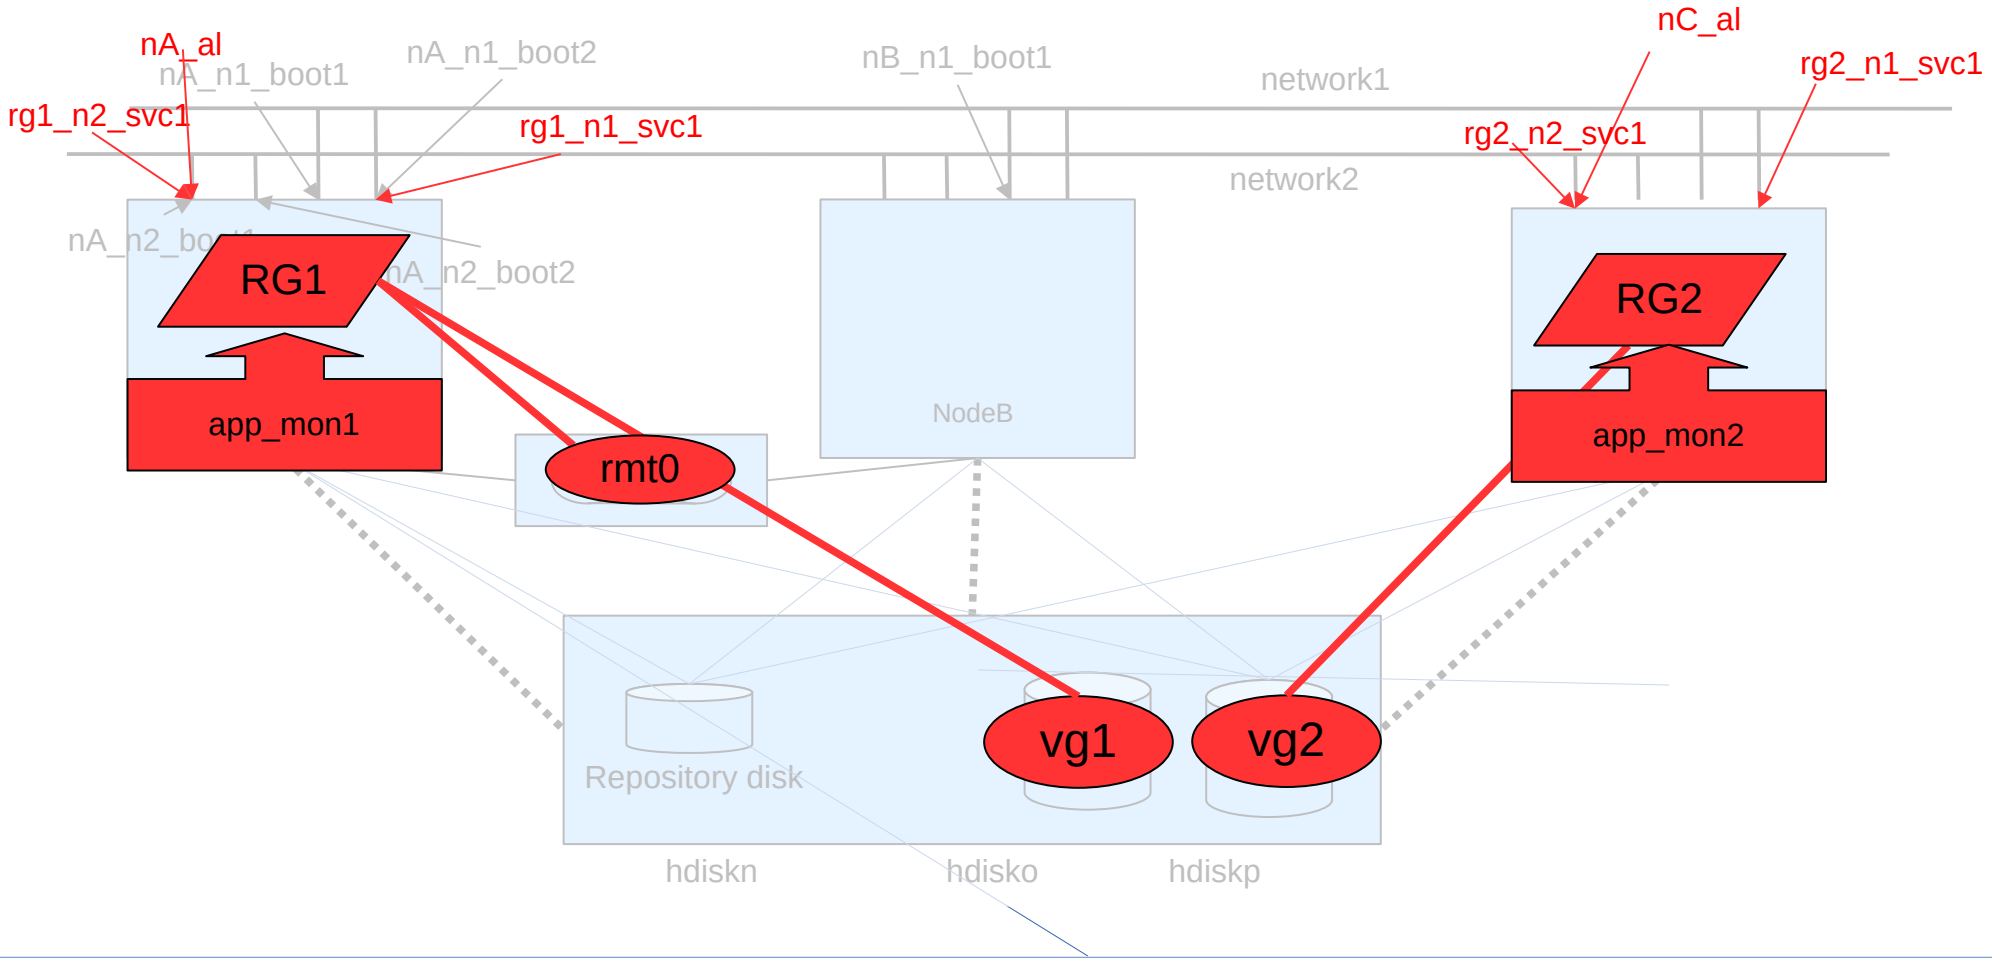

Cluster topology

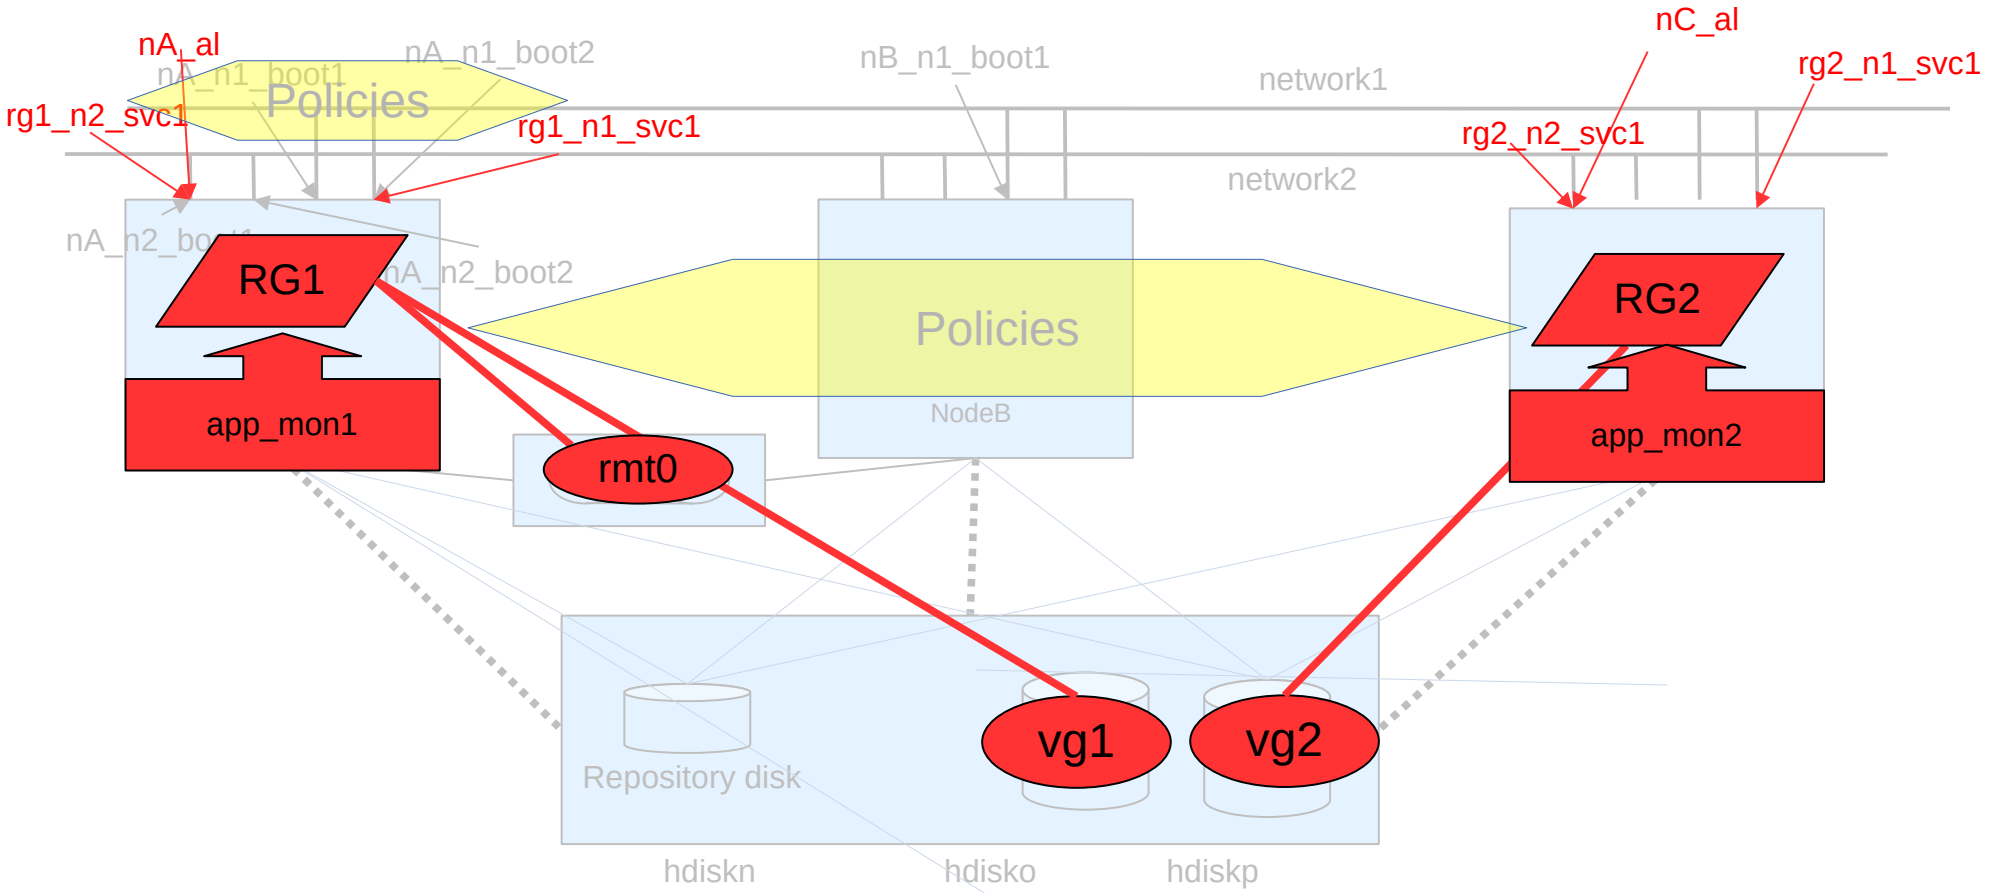

HA – local clustering

- Supported topology configurations
  - Active | Standby
  - Active | Active (Independent Workloads)
  - Active | Active (Concurrent or online on all available nodes)
- Supported Features:
  - Resource Dependencies
  - Application Monitoring
  - Custom Events
  - Integrated cpu/memory capacity management
- Supported Resource Configurations:
  - Dedicated resources
  - Virtualized (NPIV, VSCSI, SSP)
  - Live Partition Mobility awareness
  - AIX 7.2 Live Update awareness

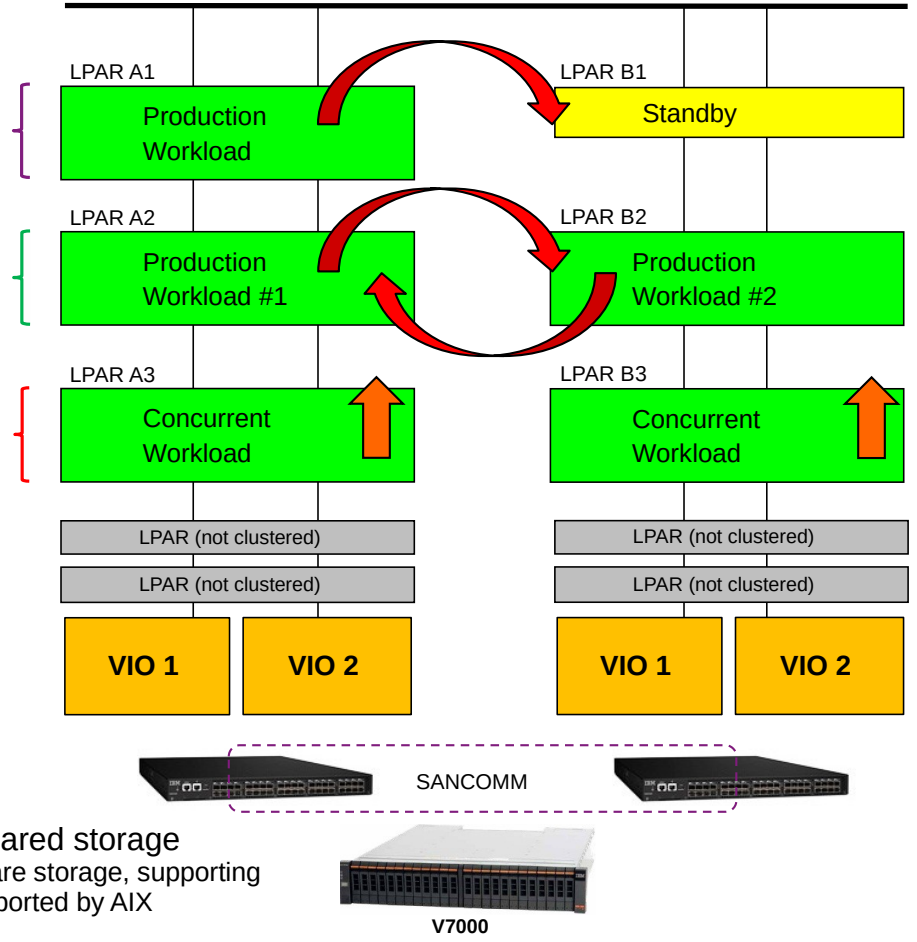

Supported shared storage
- Clusters share storage, supporting storage supported by AIX

Cluster behaviour

Belisama

- Resource Group Policies
  - Startup
    - Online on home node only
    - Online on first available
    - Online on all available
    - Start up distribution
  - Failover
    - Failover to next node in the list
    - Failover using Dynamic node priority (CPU, Paging space, Disk IO, Adaptive (user defined))
    - Bring offline
  - Fallback
    - Fallback to higher priority node
    - Never fallback
- Resource group dependencies
- IP distribution preferences
  - See next page
- Inter site management policies
  - Online on Both Sites
  - Online on Either Site
  - Prefer Primary Site
  - Ignore

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 1st 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.

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
  - Memory Utilisation
  - Disk I/O Utilisation
- Dynamic Node Adaptive Fallover
  - cl_lowest_nonzero_udscript_rc
  - cl_highest_udscript_rc

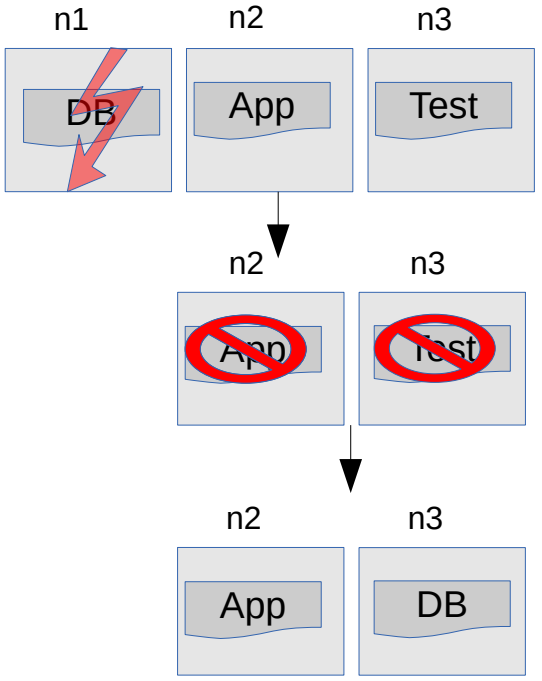

On same node dependency

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

Parent / Child

DB – parent; App - Child

Are you using PowerHA features

- Are you aware of / using
  - Fast failure detection
  - File collections
  - Application monitoring
    - Startup, long running or both
    - Process or custom
  - CSPOC
  - Cluster Test tool
- Remember that in the new versions of PowerHA, the developers used feedback from the field/PMRs to fix common problems

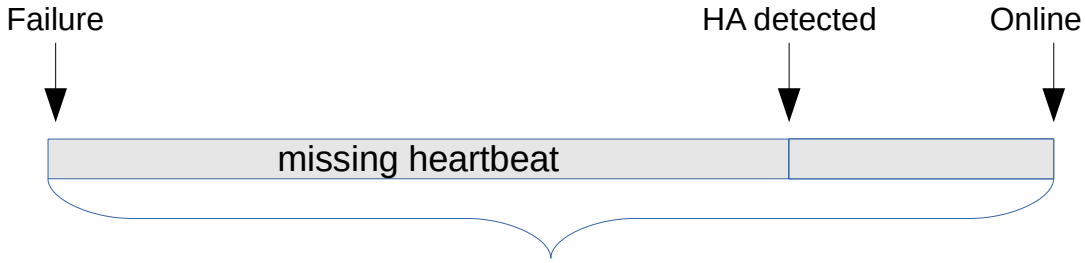

Cluster Aware AIX

- 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 20 nodes of a 21 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 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
- snap caa to collect PD data

Cluster Aware AIX (cont)

- Cluster disks.
  - CAA has information on all disks in the cluster – including their state. (3rd party disks do not participate in the monitoring).
- SolidDB and cluster disk naming dropped in 2010
- In 2011 added:
  - Deadman switch for isolated nodes – tuneable and response options.
  - 3rd party disk support added
  - Synchronous changes allowed across the cluster
  - Improved logging and RAS tools
- In 2012 added:
  - 2 sites – Linked or stretched clusters
    - Stretched Cluster (Single CAA cluster; Single Repository Disk; Require multi-cast across 2 sites; Cluster communication:- Networks, SAN, or Disk)
    - Linked Cluster (Linked CAA cluster; 2 Separate Repository Disks; One local repository on each site; Synchronised between sites; Cluster communication:- Networks)

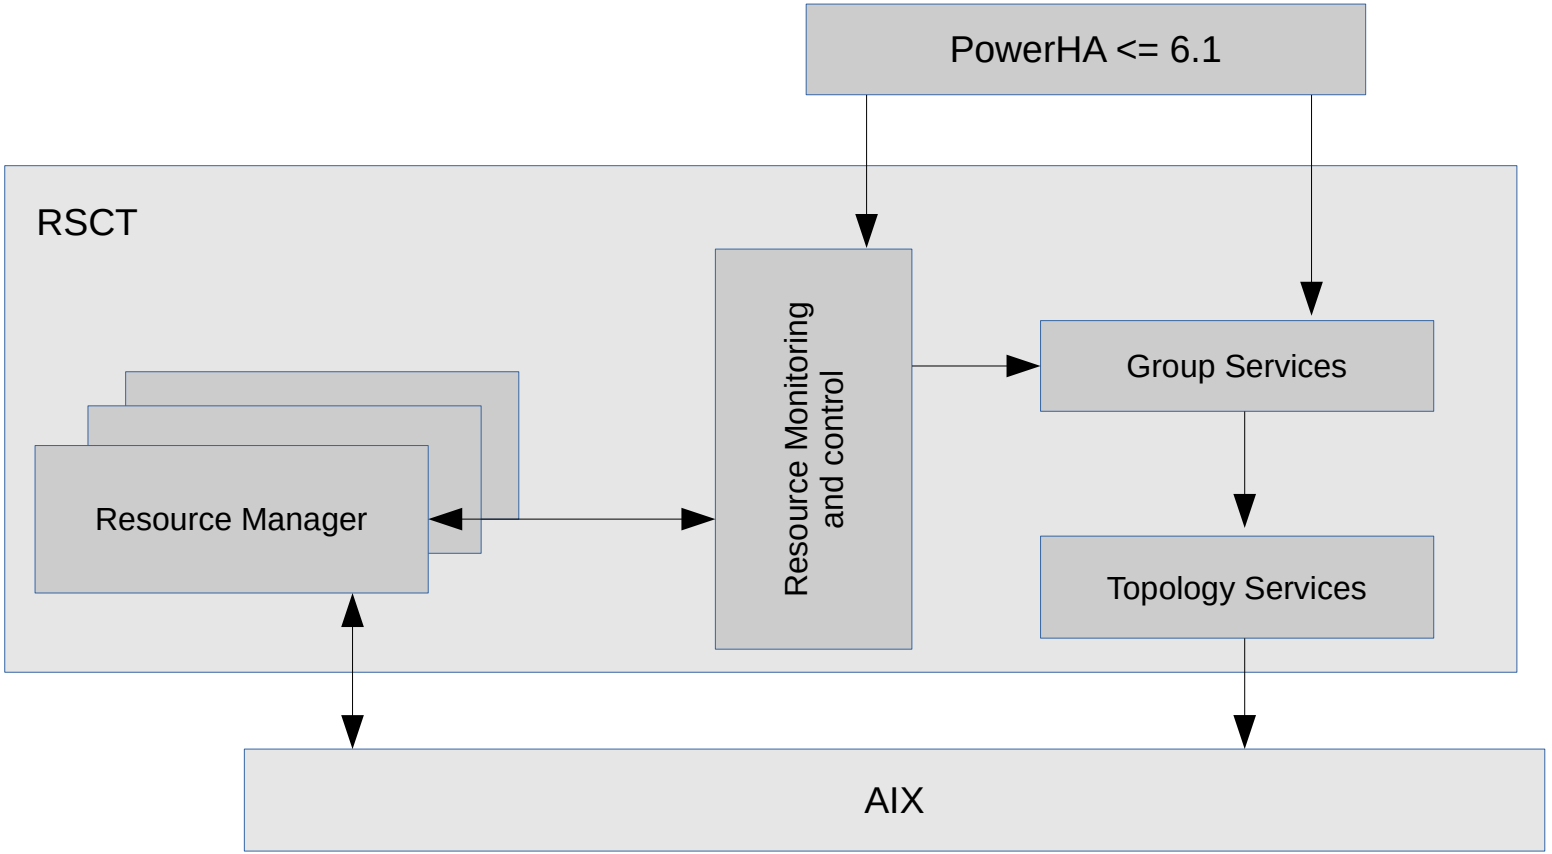
PowerHA <= 6.1
RSCT
Resource Monitoring and control
Group Services
Resource Manager
Topology Services
AIX

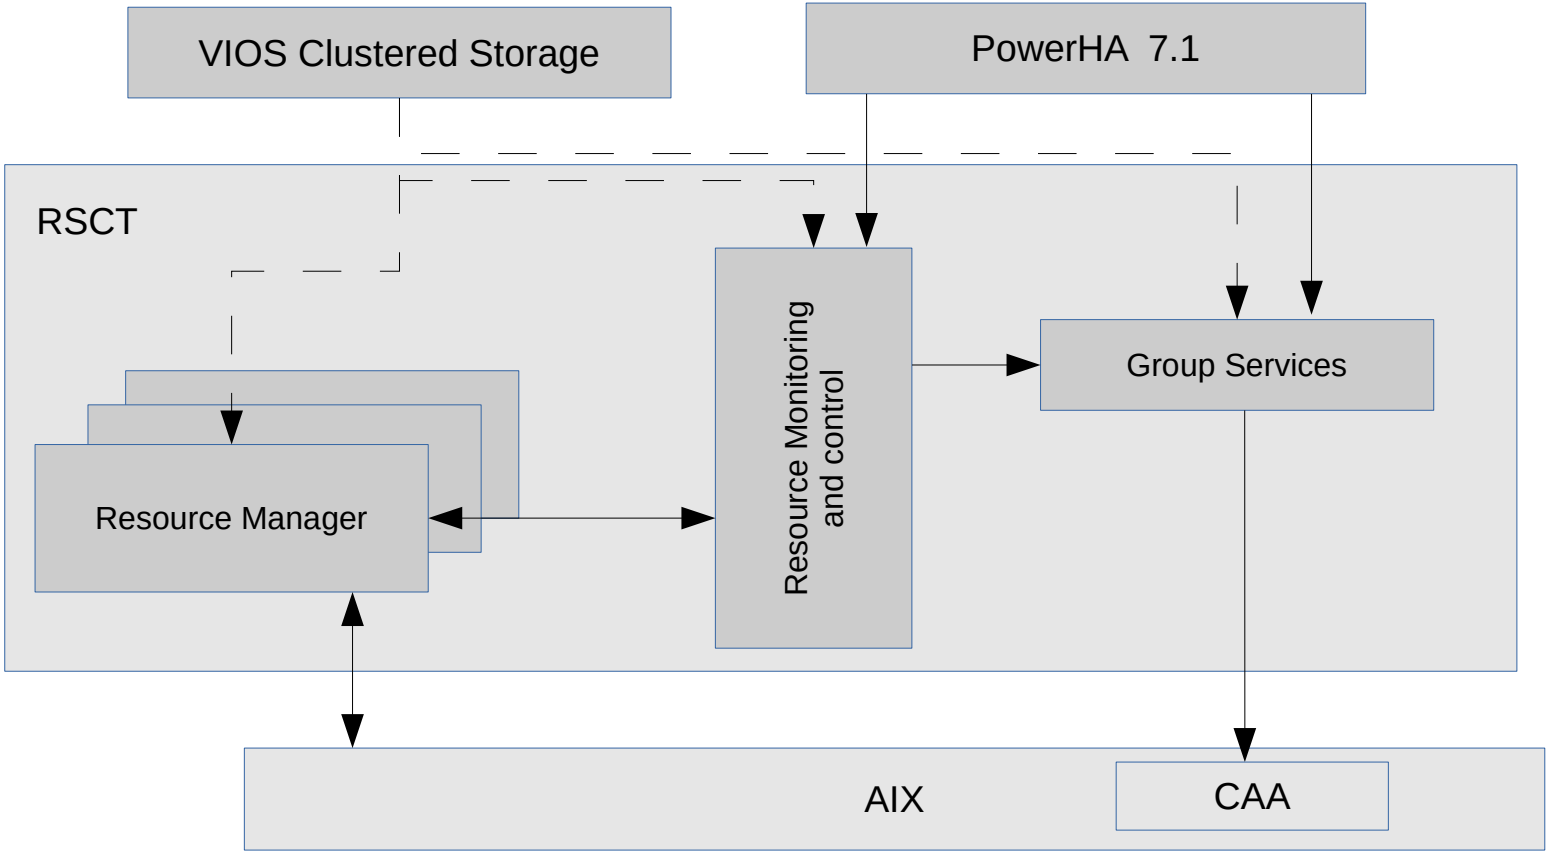
VIOS Clustered Storage
PowerHA 7.1
RSCT
Resource Monitoring and control
Group Services
Resource Manager
AIX
CAA

Cluster aware AIX – Topology management

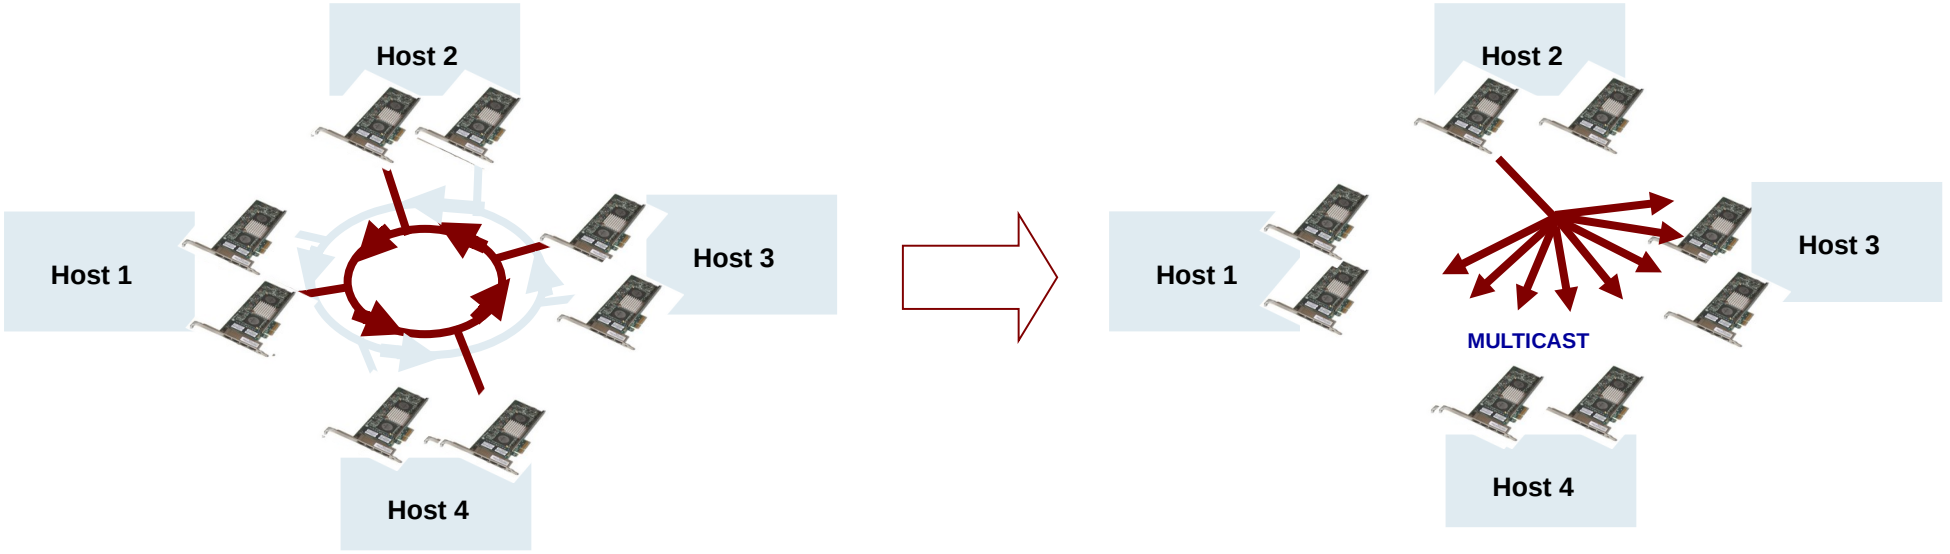

- PowerHA 6.1 & > 7.1.3
- Heartbeat Rings: detailed protocol
  - Leader, Successor, Mayor etc
  - Difficult to add/delete nodes
- Requires IP aliases management in the subnet

- PowerHA >7.1
- Multicast based protocol
  - Discover and use as many adapters as possible
  - Use network and SAN as needed
  - Adapt to the environment: delay, subnet etc
- Kernel based cluster message handling

- Minimal Setup
- Multiple channels of communication
  - **Network**
  - **SAN**
  - **Central Repository**

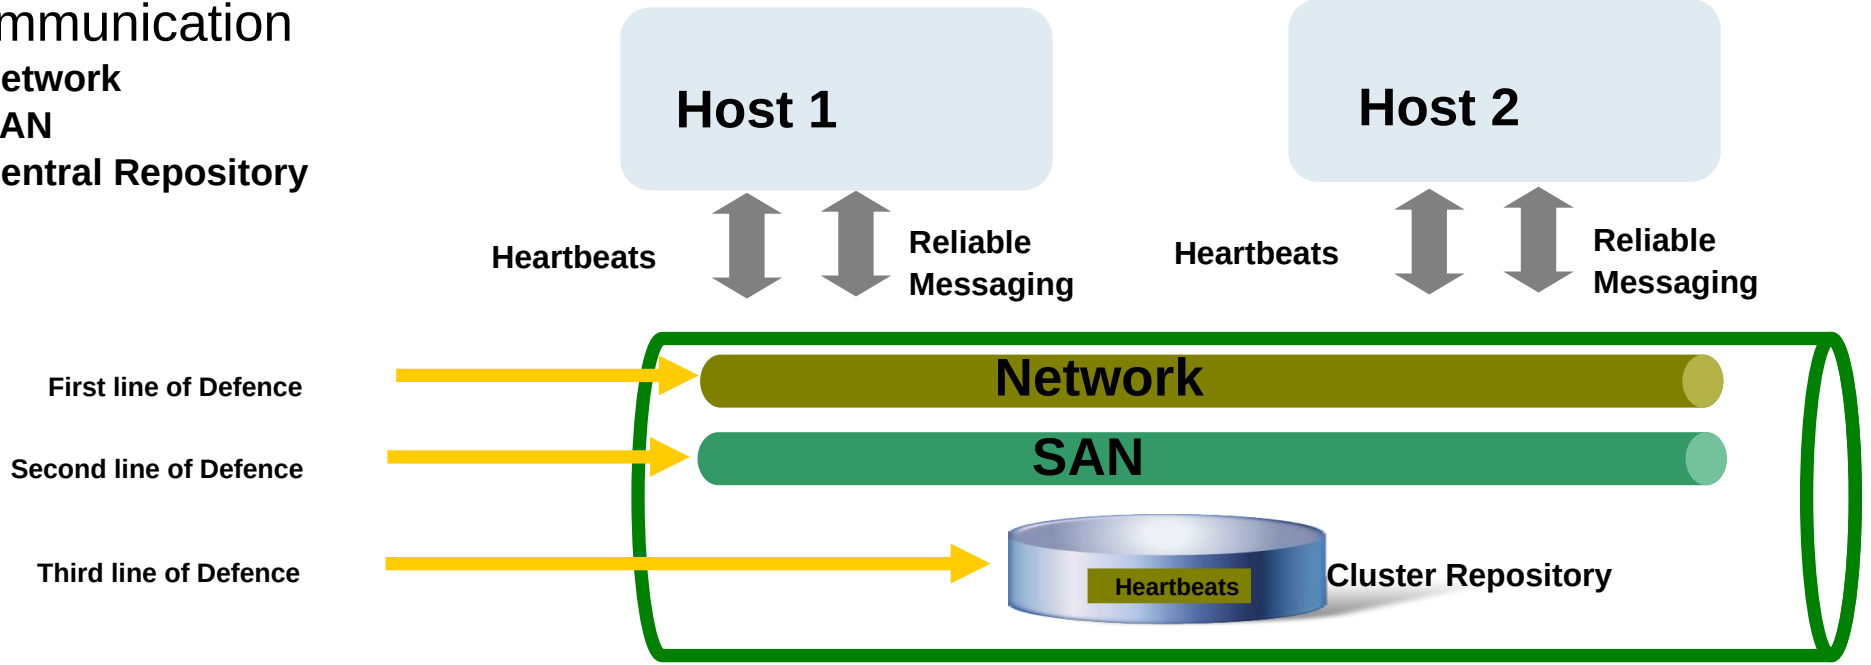

3 lines of (redundant) independent communications

SAN heart beating in virtual environment

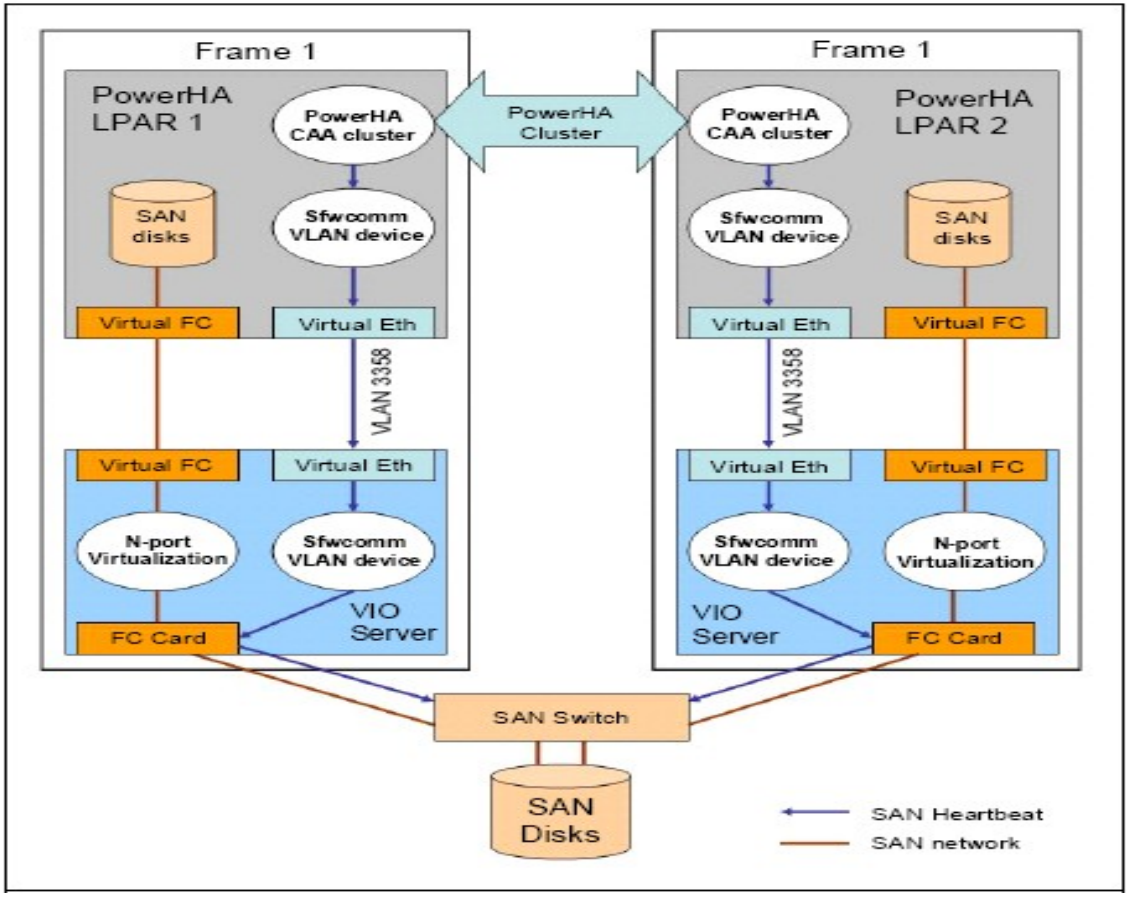

https://www.ibm.com/developerworks/aix/library/au-aix-powerha-heartbeat/index.html

HA – Application Monitoring

- Some monitors are provided in the Smart Assists
  - The application monitor is bound to the application controller
    - Startup Monitor
    - Long running Monitors:
      - Process Monitor
      - Custom Monitor
  - Startup Monitor
    - Invoked when application started and used to confirm application running for other dependencies
  - Process Monitor
    - Checks the process table
  - Application Monitor
    - Uses custom logic (read record in table)
  - Can be configured to perform
    - Application restarts / notification / Resource Group failover

HA – Application controllers

- Application controller (script) can be started:
  - 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 RC=1

Linked / stretched

Belisama

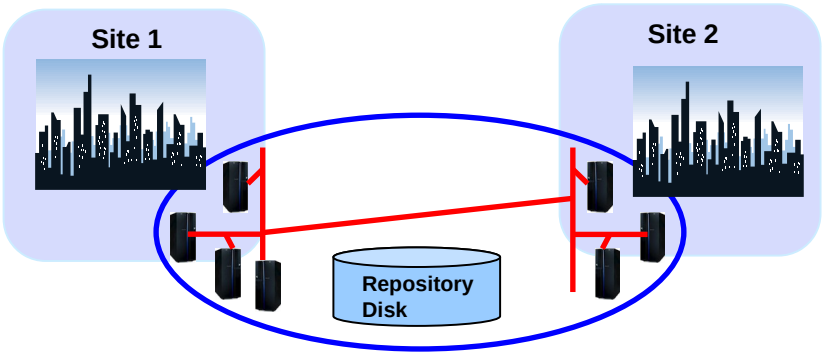

Multi Sites with Stretched Cluster

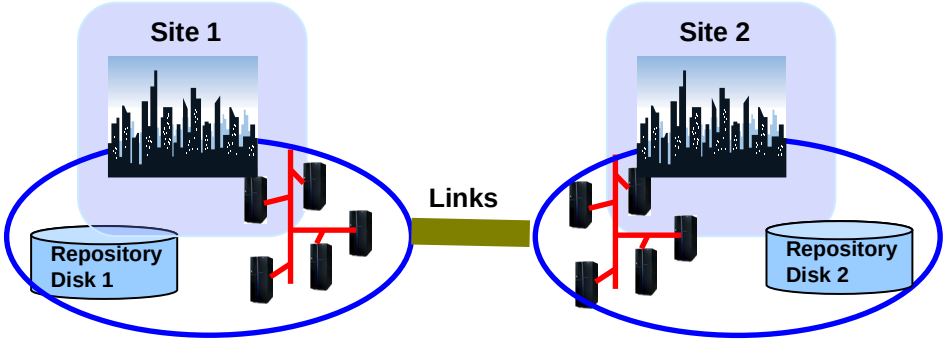

Multi Sites with Linked Clusters

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

	Standard	Enterprise
Multi Site Definition • Site Service IP • Site Policies	✓	✓
Stretched Clusters	✓	✓
Linked Clusters	✓	✓
HADR with storage replication management	✗	✓
Hyperswap	✗	✓

Tie breaker support

- PowerHA 7.1.2 Tie Breaker Support
  - Separate Site Split and Merge policies
    - Split/Merge: Tie Breaker policy
    - FC/iSCSI Tie Breaker
    - SCSI 3 reservation disk
    - Losing side is quiesced

More suited for Linked Clusters

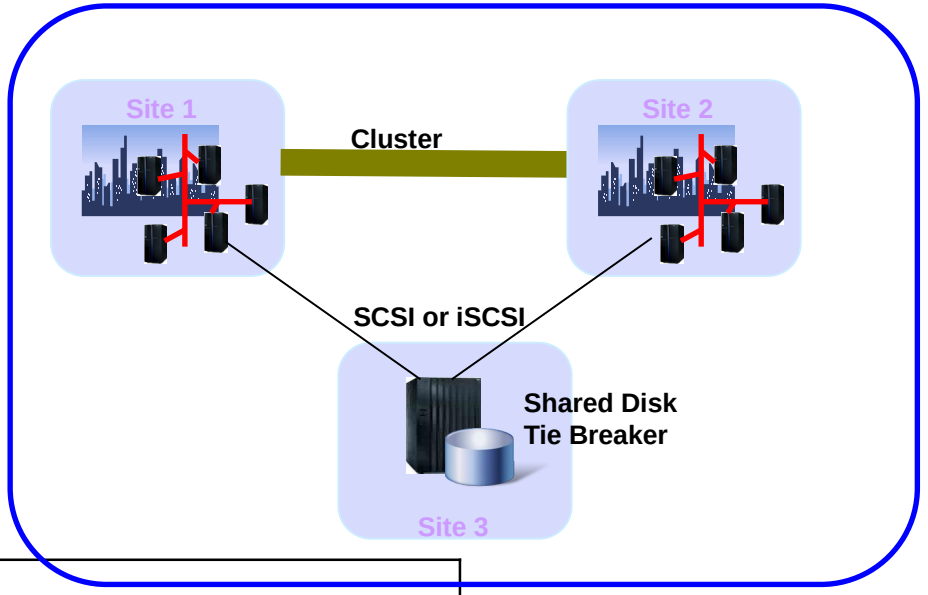

Policy Setting	Split	Merge	Comments
Tie Breaker	✔	✔	▪ **Tie break Holder side wins**
Majority Rule	✔	✔	▪ **>N/2 side wins** ▪ **In case of N/2, side that includes node with the smallest node id**
Manual	✔	✔	▪ **Manual steps needed for recovery to continue**

DS8K Series Storage: Metro Mirror Capabilities

- PowerHA HA/DR with IBM DS8K Mirroring
  - Metro Mirror uses Synchronous mirroring technology
  - No Data Loss (Recovery Point Objective = 0)
  - Distances up to 300 km supported
    - Response time/latency impacted by distance - for light to travel 100 km round trip (=200 km distance) takes 1 ms
- PowerHA
  - Automation of workload to mirror relationship management
    - PowerHA supports Metro & Global Mirror
    - Hyperswap support for Metro Mirror
  - Support for:
    - Consistency Groups provide data integrity for databases
    - Consistency Groups can span multiple disk subsystems
    - Multiple consistency groups supported between sites
    - Adapts to failures and adjusts the mirrors as needed

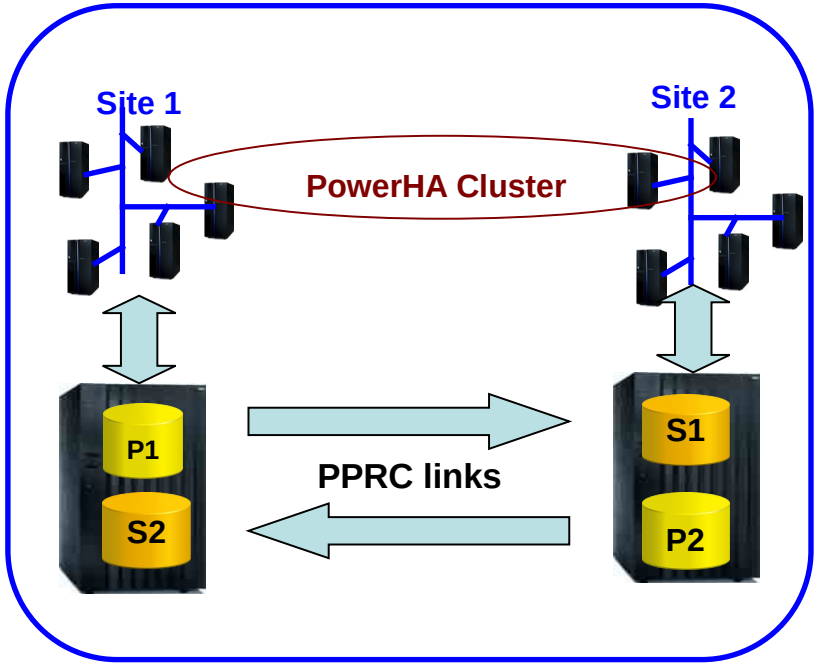

PowerHA HyperSwap Background

Belisama

- 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 host to transparently switch an applications I/O operation to the secondary volumes, provided physical connectivity exists between the host and the secondary storage subsystem.

HyperSwap Technology

- Continuous Availability against Storage failures
- Substitutes storage secondary to take the place of failed primary device
  - Non-disruptive - applications keep running
  - Key value add to HA/DR deployments

Customer Benefits

- **Unplanned HyperSwap:**
  - **Continuous Availability against storage failures**
- **Planned HyperSwap:**
  - **Storage Maintenance without downtime**
  - **Storage migration without downtime**

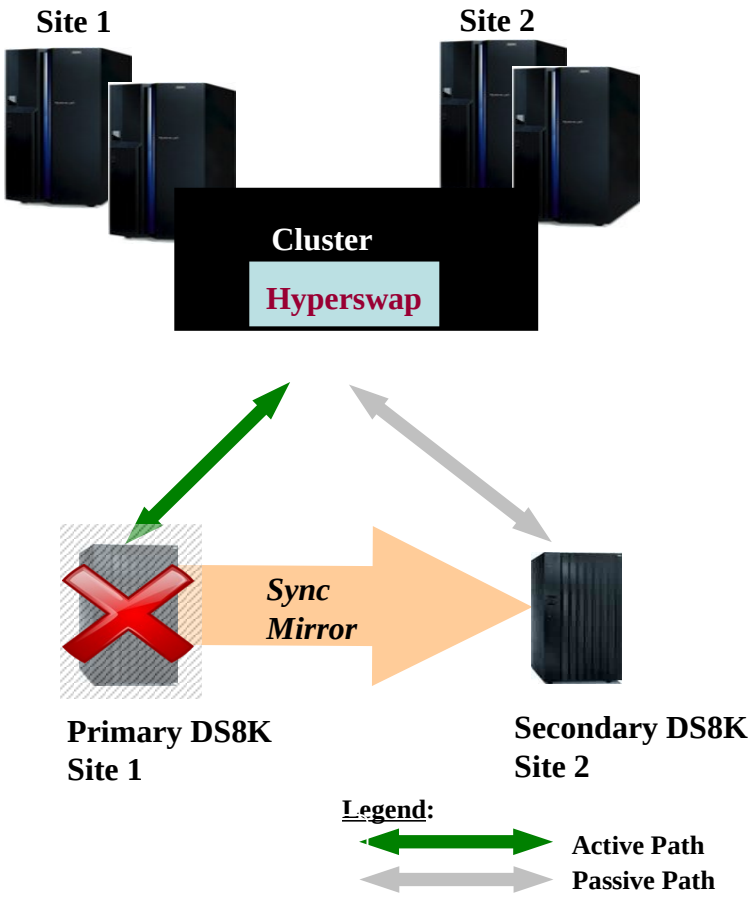

- **HyperSwap device configuration transparent to application**
  - ***Application can continue to use the device as before***

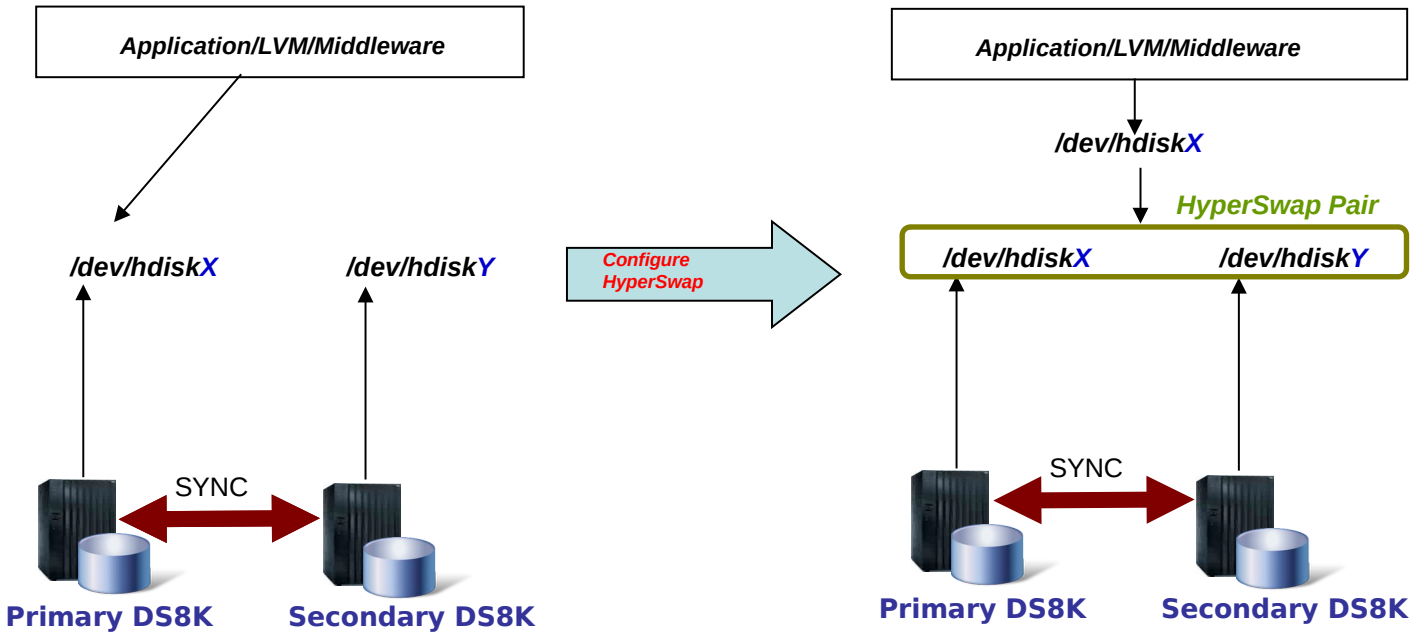

Auto-Resync Support

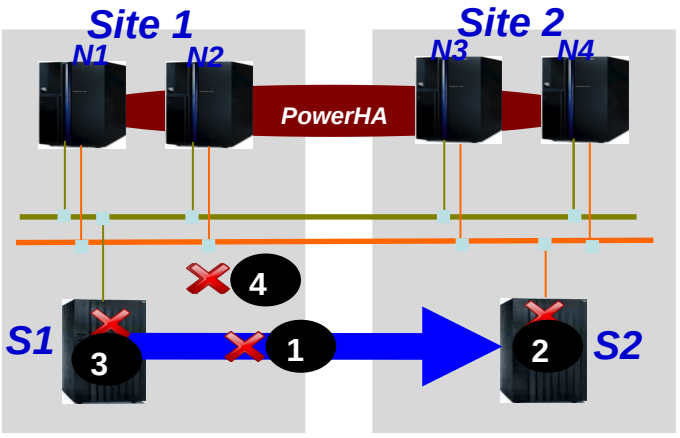

- PowerHA 2013 supports Auto-Resync Policy
  - Administrator can turn on/off “auto resync”

Events	Description	Auto-Resync
1	Replication links failed and then recovered	Primary stays active and secondary is synchronized
2	Secondary Storage failed. Later on Secondary Storage Recovered	Primary stays active and secondary is synchronized
3	Primary Storage Fails Primary Storage Recovers	HyperSwap occurs to Secondary Resync from Site 2 to Site 1 is initiated
4	Site 1 Fails Site 1 Recovers	Secondary storage is made active Resync from Site 2 to Site 1 is initiated

PowerHA installation

- Recommended levels
  - https://aix.software.ibm.com/aix/ifixes/PHA_Migration/ha_install_mig_fixes.htm
- Standard edition filesets
  - CAA packages (part of the AIX install media)
    - bos.cluster.rte
    - bos.ahafs
    - bos.clvm.enh
    - devices.commom.IBM.storfwork
  - PowerHA packages
    - cluster.license electronic license file
    - cluster.es.server base cluster filesets
      Clinfo and Clstat samples and include files and a Web Based Monitor
    - cluster.doc.en_US.es PowerHA SystemMirror PDF Documentation
    - cluster.es.client cluster client binaries and libraries, plus Web based Smit for PowerHA
    - cluster.es.cspoc CSPOC and Dsh
    - cluster.es.migcheck Migration support
    - cluster.es.nfs NFS Server support
    - cluster.msg.en_US.es U.S. English message catalogue
    - cluster.man.en_US.es man pages - U.S. English

PowerHA installation (cont)

Belisama

- PowerHA packages (optional)
  - cluster.doc.en_US.assist — Smart Assist PDF documentation
  - cluster.hativoli — PowerHA SystemMirror Tivoli Server and Client
  - cluster.es.assist — Smart Assist filesets
  - cluster.msg.en_US.assist — U.S. English Smart Assist messages
  - cluster.es.cfs — GPFS support
  - cluster.es.worksheets — Online Planning worksheets

PowerHA installation (cont)

Belisama

Replication Type	Replication Type File Sets to Install
ESS Direct Management PPRC	cluster.es.pprc.rte cluster.es.pprc.cmds cluster.msg.en_US.pprc
ESS DS6000/DS8000 Metro Mirror DSCLI PPRC	cluster.es.spprc.cmds cluster.es.spprc.rte cluster.es.cgpprc.cmds cluster.es.cgpprc.rte cluster.msg.en_US.svcpprc
San Volume Controller (SVC) & Storwize Family	cluster.es.svcpprc.cmds cluster.es.svcpprc.rte cluster.msg.en_US.svcpprc
XIV, DS8800 in-band and Hyperswap, DS8700/DS8800 Global Mirror	cluster.es.genxd.cmds cluster.es.genxd.rte cluster.msg.en_US.genxd
Geographic Logical Volume Mirroring (GLVM)	cluster.doc.en_US.glvm.pdf cluster.msg.en_US.glvm cluster.xd.glvm glvm.rpv *(file sets in base AIX)
EMC SRDF	cluster.es.sr.cmds cluster.es.sr.rte cluster.msg.en_US.sr
Hitachi TrueCopy / Universal Replicator	cluster.es.tc.cmds cluster.es.tc.rte cluster.msg.en_US.tc

Plus: cluster.xd.license

At-a-Glance Heath

Belisama

- 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?

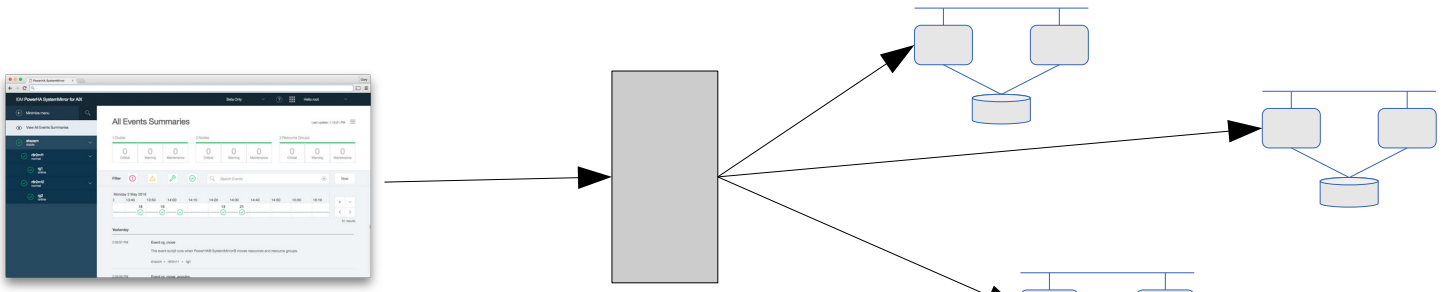

Browser (Chrome or Firefox)

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

Using the UI to resolve PowerHA problems

Belisama

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

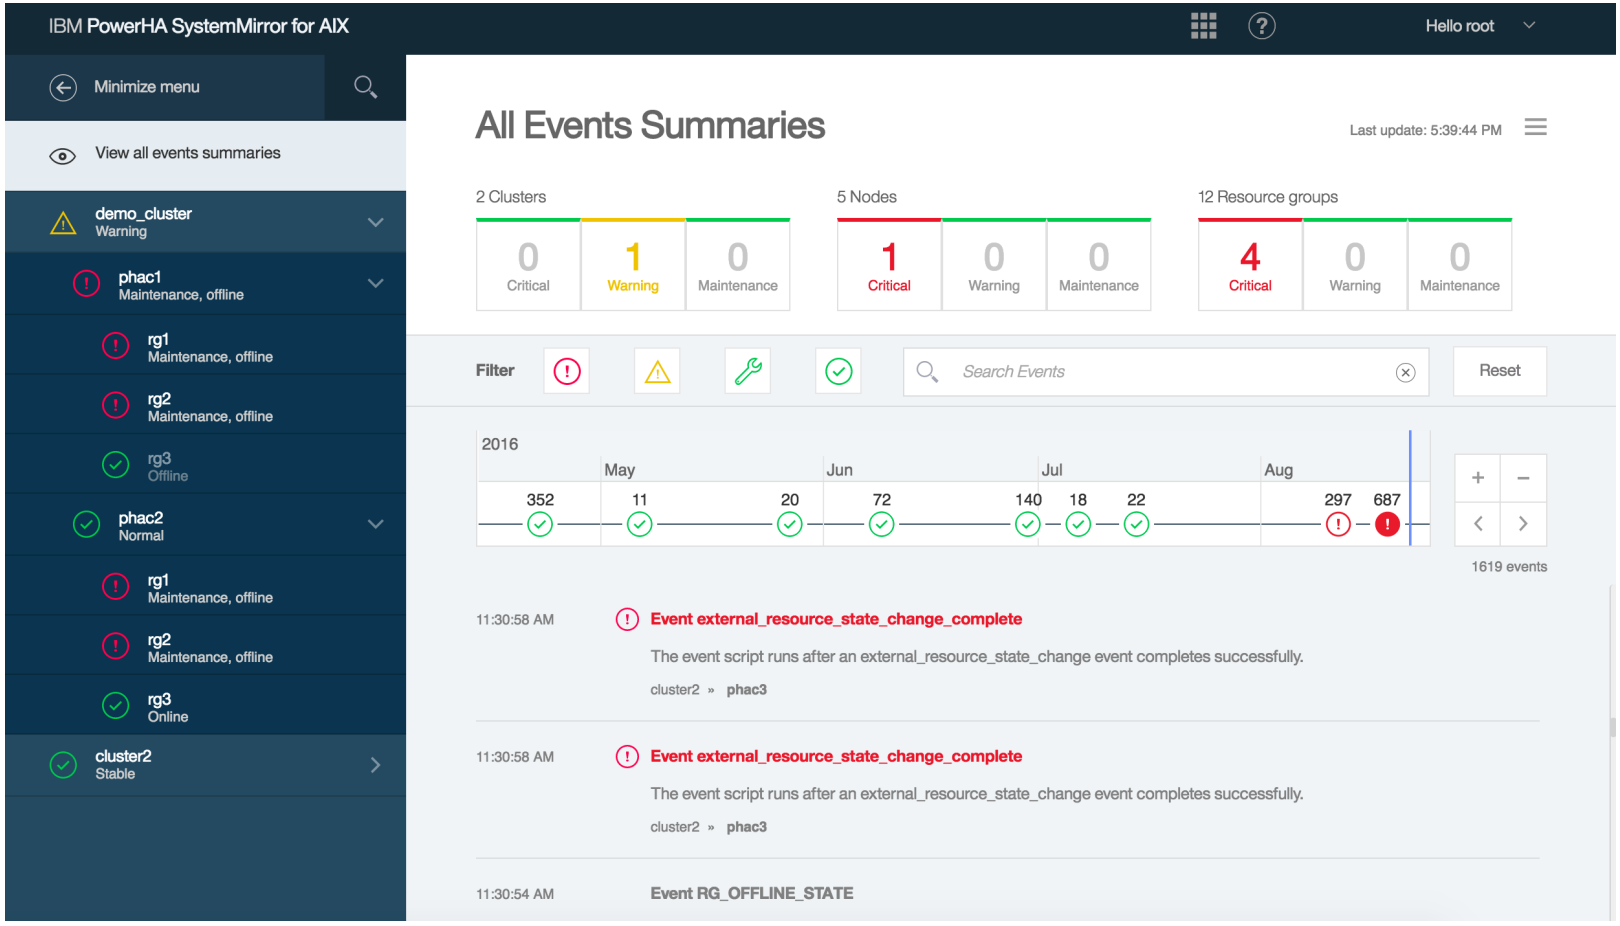

Using the UI to resolve PowerHA problems

Belisama

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

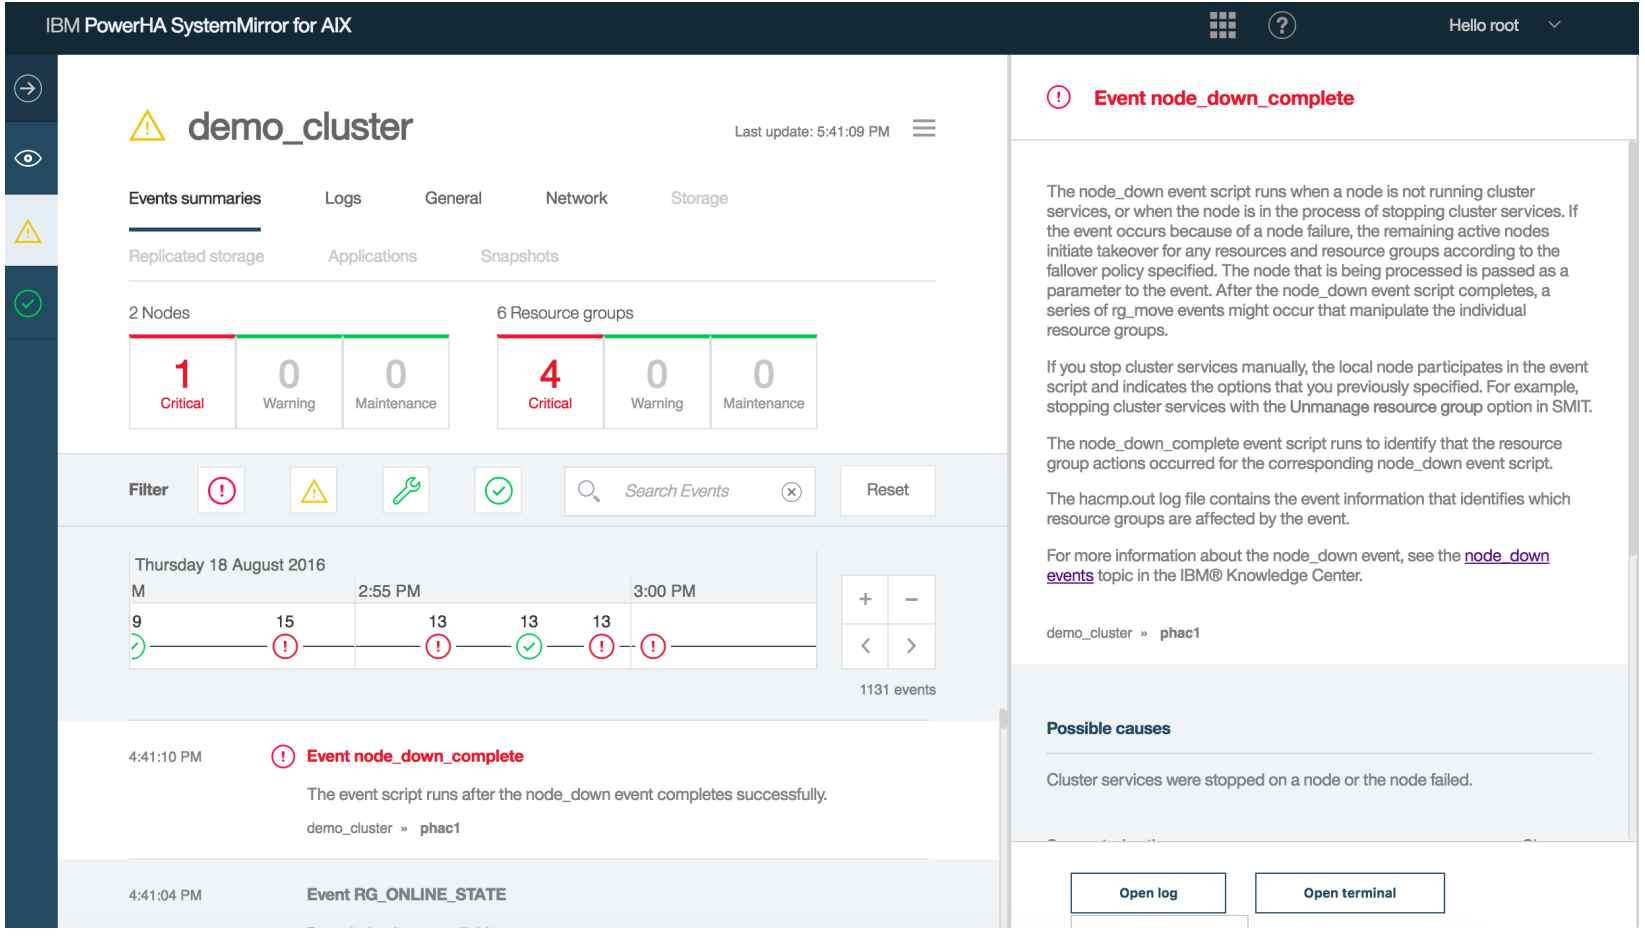

Using the UI to resolve PowerHA problems

Belisama

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

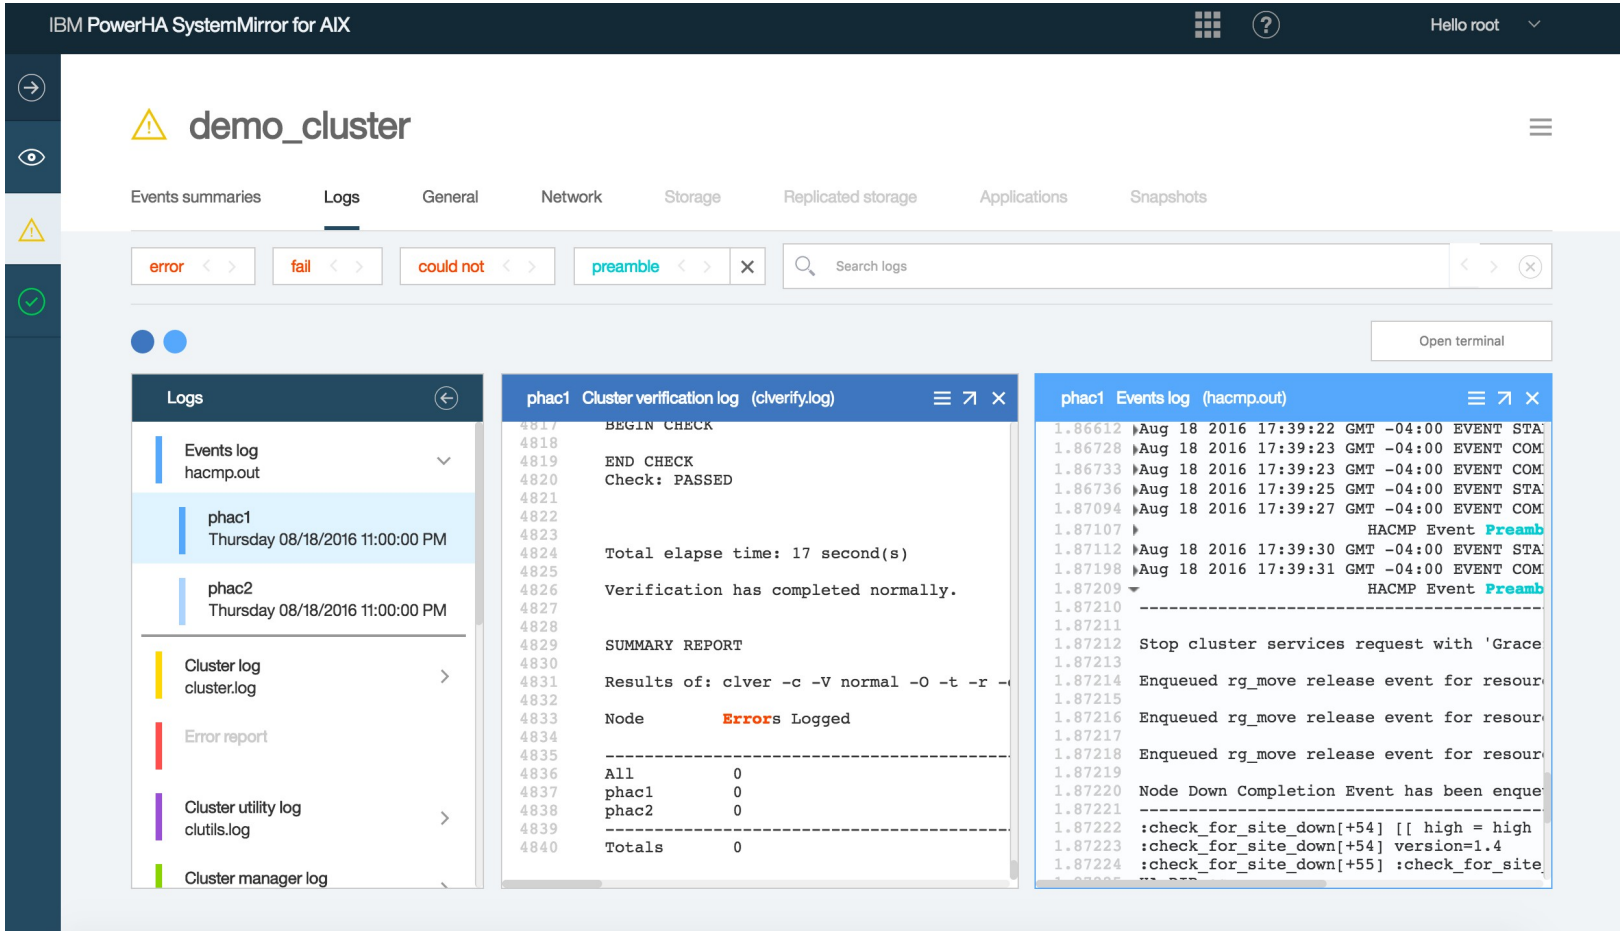

Using the UI to resolve PowerHA problems

Belisama

- 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.

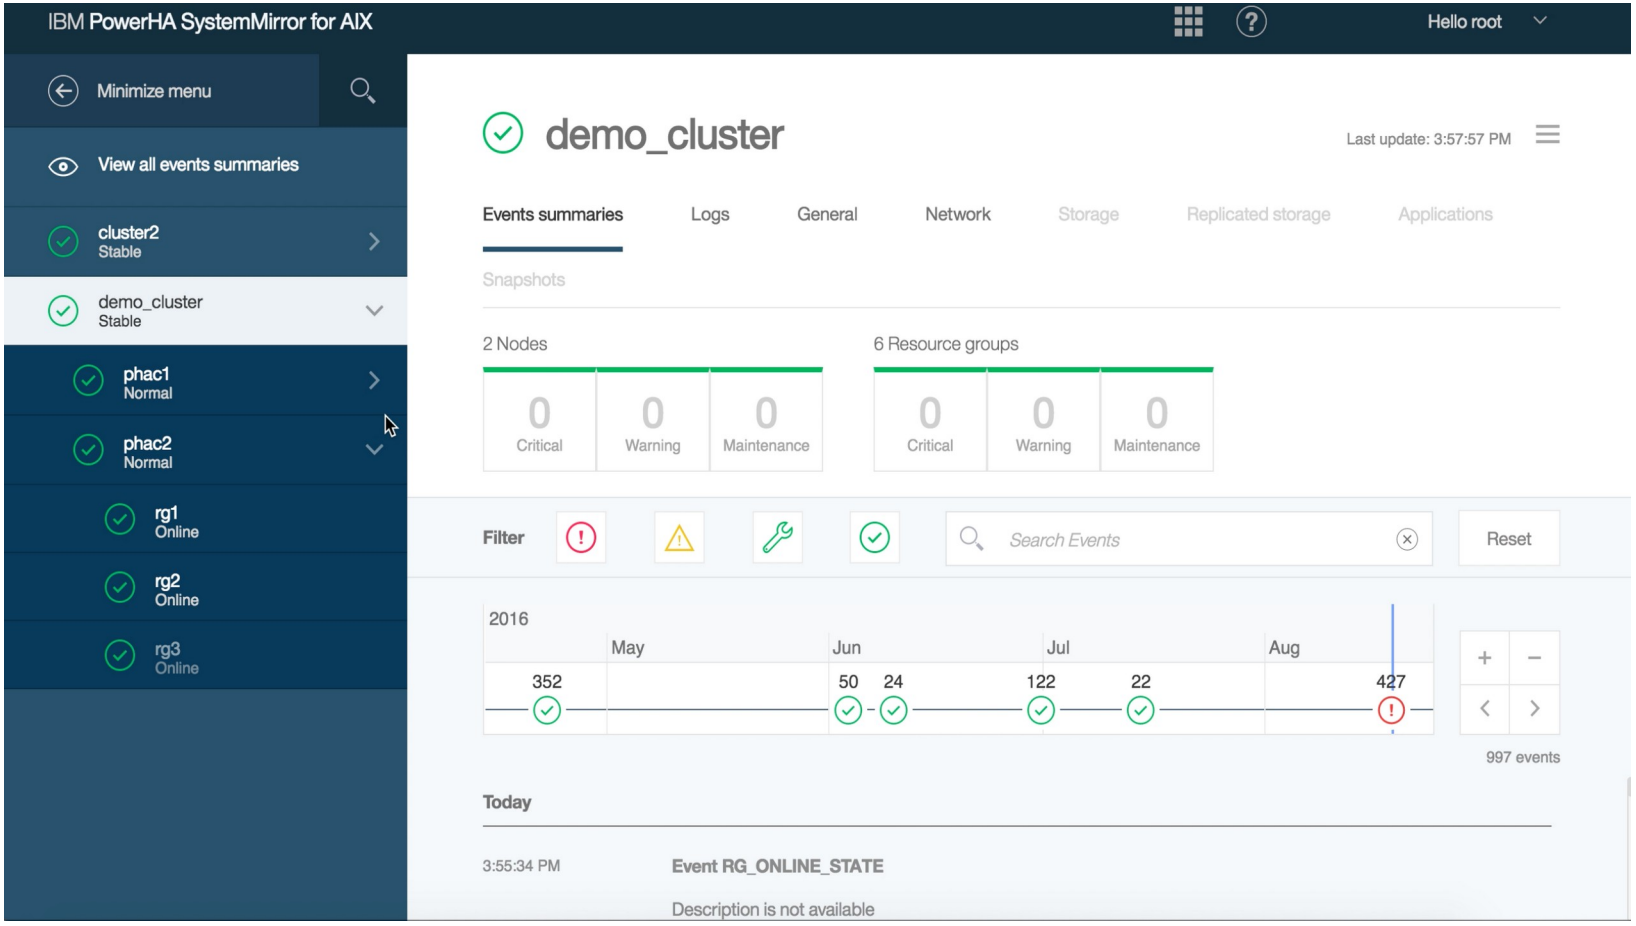

PowerHA SystemMirror Release Life Cycle

Belisama

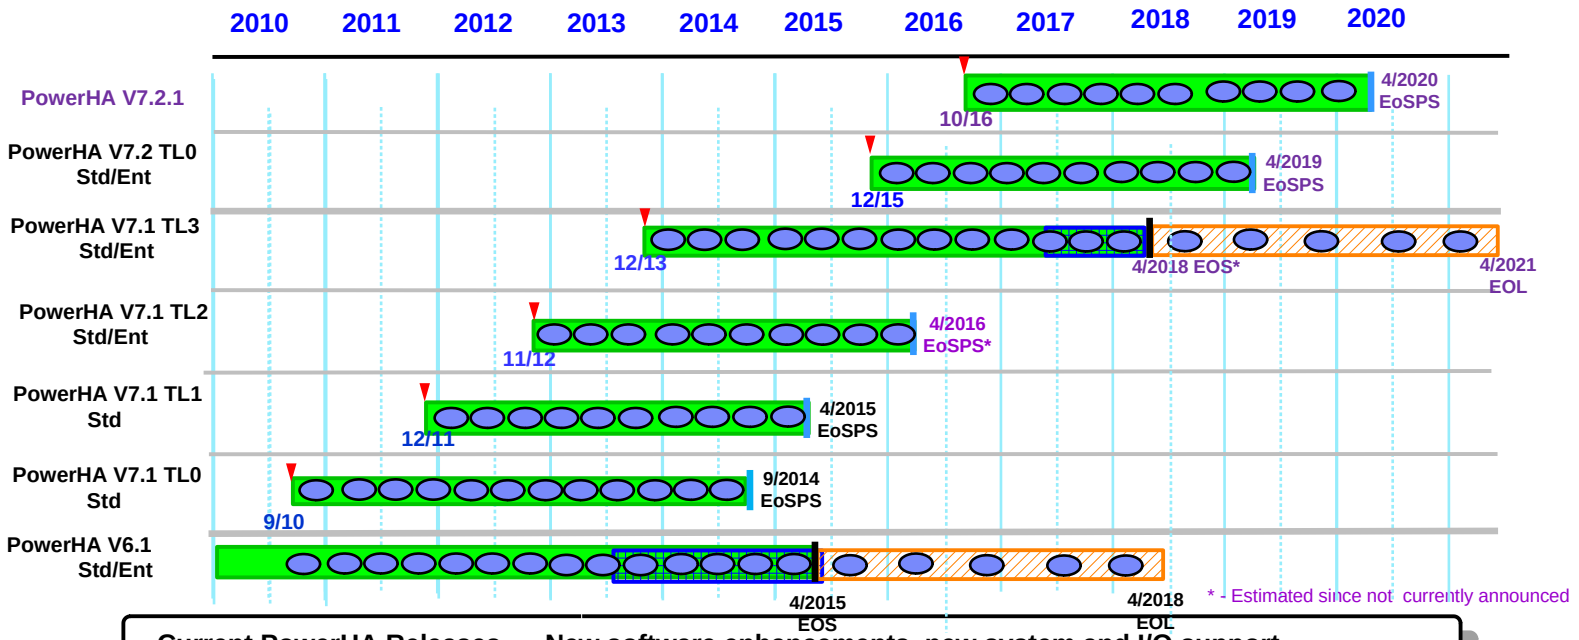

Current PowerHA Releases - New software enhancements, new system and I/O support
PowerHA Service Packs - Software updates, New hardware support

- Product Marketing and Service Available
- Withdrawn from Marketing; Standard Program Services Available
- End of Service marker (without additional fees)
- EoSPS: End of Service Pack Support marker (Debug Support only after this date)
- Fee-based support offering
- SP update for software
- Significant new functionality added

New features summary

- PowerHA 7.2 summary
  - ROHA (Resource Optimised High Availability)
    - flexible Capacity management across various Power pools
    - Enterprise Pool Integration
    - Manipulate Shared Processor Pool Sizes
    - New HMC Integration Tunables
  - Non-Disruptive Upgrade Support (PowerHA code)
    - Ability to upgrade HA to 7.2 from 7.1.3 without requiring a Rolling Upgrade or interruption of service
  - Automatic Repository Disk Replacement
    - Define multiple repository disks & auto replacement behavior with AIX 7.2
  - AIX Live Update Support & LPM Support Enhancements
    - (optional) Validation of a number of new checks including AIX Expert Settings
  - Cluster Detailed Verification Checks
  - (optional) Validation of a number of new checks including AIX Runtime Expert Settings
  - Quarantine Policies (Critical RG)
    - Split handling protections
    - HMC Node Halt Policy
    - SCSI-3 Node Fence Policy
  - NFS Backed Tie Breaker Disk support
    - New support flexibility to avoid the need of a NAS backed devicewhen using Tie Breaker Disk function

New features summary

- Non-disruptive upgrade (NDU) support
  - Cluster software upgrade without disruption to workload
  - Easy PowerHA Migration without downtime
- Enhanced Split/Merge Policies
  - Expanded choices of Split/Merge handling
  - Consistent policies for all types of clusters
    - Standard, Stretched, & Linked Clusters
- Health Dashboard
  - Customer collaboration based graphical interface
  - Design thinking principles incorporated

	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	

What's new V7.2.1 1H 2017 enhancements

- Announce & GA planned for 1half 2017
- PowerHA Health Dashboard support for earlier PowerHA releases
  - Agent support for PowerHA 713 and 720
- Easy update management for PowerHA software stack
  - Update one node or all nodes in the cluster with ease
- PowerHA compatible with AIX Live Kernel update
  - No timing issues with PowerHA during LPM and Kernel Update
- REST based HMC communication
  - Support for Enterprise pool management (ROHA) and Active Node Halt Policy in Cloud environments

- New features for PowerHA 7.2.3

GA: 14/12/18

  - 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 analyzer and miscellaneous features
  - GUI enhancements (Power UI Server can be AIX 7.1,7.2; SUSE 12 le; RHEL 7 le/be)
  - Miscellaneous tidy
    - 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

New features - Linux

Belisama

- Previous…

GA: 14/12/18

  - Support for RHEL (RHEL 7.2, 7.3, 7.4) and SLES (SLES 12 SP1, SP2, SP3) on PowerVM platform and 4 nodes with GUI administration
  - Network and Disk heartbeat management
  - Support for core functions such as: Tie breaker split handling etc
  - SAP HANA & SAP NW HA support
    - SLES based support (based on current support from SAP for Power platform)
    - CLI based smart assist wizard support for SAP HANA
- 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)

GUI update

- 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

Other availability options

- PowerHA is managed within the operating system environment
- Options for external management
  - Simplified Remote Restart (SRR)
  - PowerVC Managed SRR
  - 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
    - Using *ghostdev* and scripts for network changes at DR site

Comparison

Belisama

	Storage Requirements	Automated Failover	Source Server Status	Source VIO Server Status	VM/Application Outage
Live Partition Mobility	Shared	No	Active	Active	No (if VM is active)
Simplified Remote Restart	Shared	No	Inactive	Inactive	Yes
VM RM HA	Shared	Yes	Active or Inactive	Active or Inactive	Only if a server/VM outage occurs
PowerHA	Remote Copy	Yes	Active or Inactive	Active or Inactive	Yes
PowerHA EE	Remote Copy	Yes	Active or Inactive	Active or Inactive	Yes
VM RM DR	Remote Copy	No	Active or Inactive	Active or Inactive	Yes

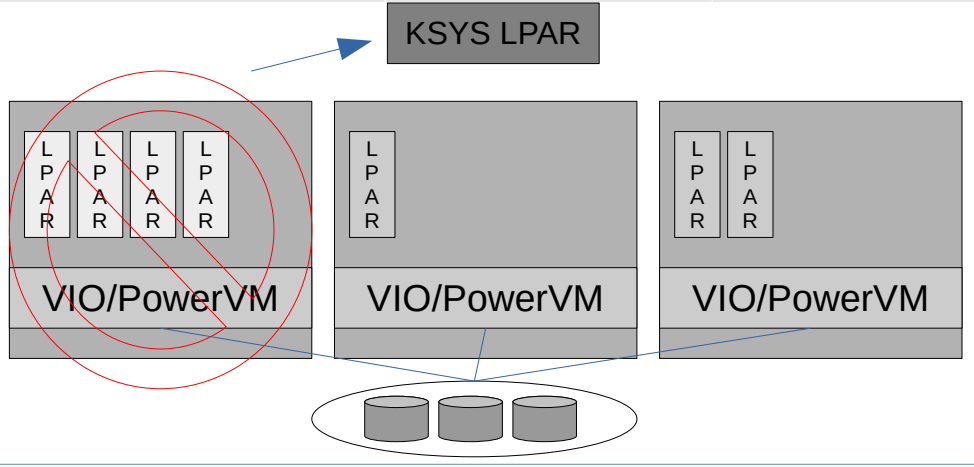

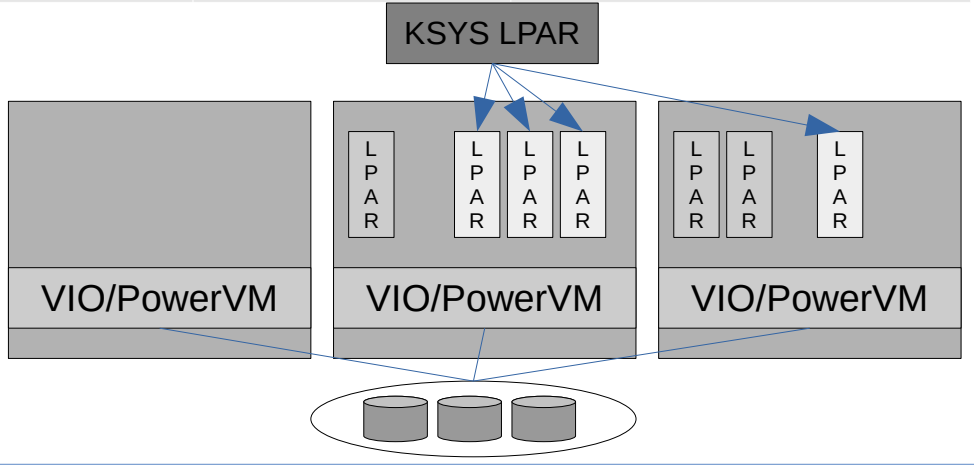

Useful links

- 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
- Developerworks Wiki: http://ibm.biz/developerworks-PowerHA-wiki
  - QA Forum: http://ibm.biz/developerworks-PowerHA-Forum
- 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

Some older links

- PowerHA 7.2 Release Notes
  - https://ibm.biz/BdHaRM
- PowerHA 7.2 Redbook
  - http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248278.html?Open
- 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: a111034

PowerHA 7.2.3 for AIX and Linux

¿ Questions ?

Thanks!

Your feedback about this session is very important to us.

Please remember to submit a survey

For further information….
Contact:

Antony (Red) Steel antony.steel@belisama.com.sg
+65 9789 6663

Power Systems

Backup Slides

Technical Details

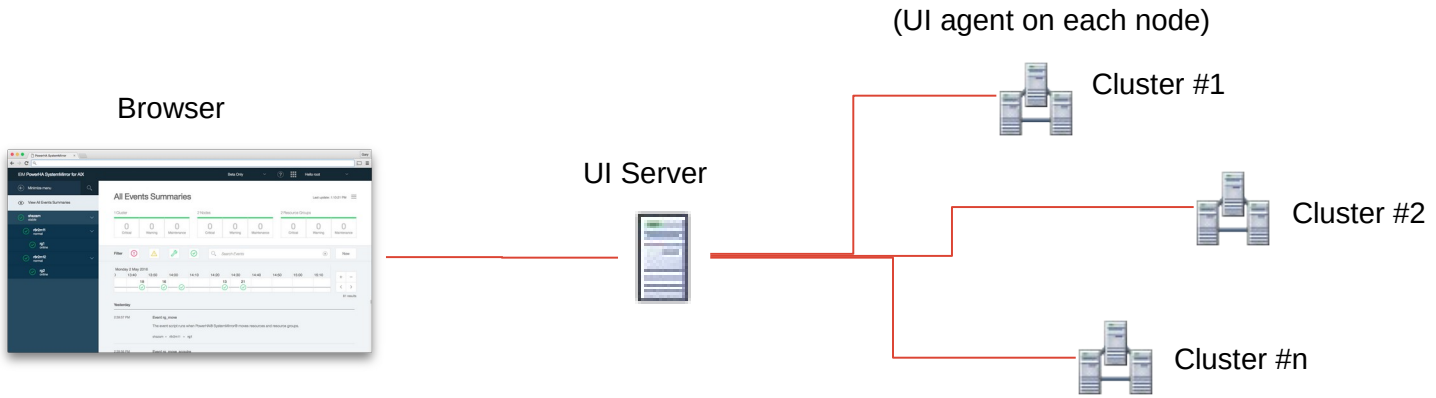

- Supported PowerHA Levels
  - 7.2.1 GA
  - 7.2, 7.1.3 via service packs
- Supported AIX Levels
  - 6.1, 7.1, 7.2
  - Minimum required by the associated PowerHA version

Supported Browsers

1. Google Chrome
2. Mozilla Firefox

PowerHA SystemMirror: Easy update tool

- Simplified IFIX, SP and TL 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

Supported Updates

Update Type	Reboot Required	Non-disruptive
PowerHA IFIX, SP, or RL	No	Yes
AIX (CAA & RSCT) IFIXes	No	Yes
AIX SP (all of AIX changes)	Yes(mostly)	No
~~AIX TL~~	~~Yes~~	~~Outside the tool~~

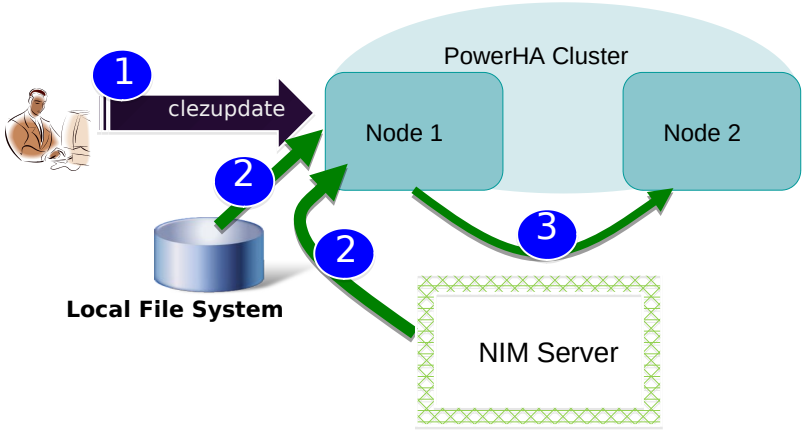

PowerHA SM: REST communication with HMC

- PowerHA to HMC via REST based management or SSH
  - PowerHA implemented key features such as ROHA (Enterprise pool management) & Active Node Halt Policy(ANHP for Split related quarantine management)
  - PowerHA supported these features by integrating with HMC
  - PowerHA worked with HMC using SSH protocols
  - HMC also supports REST API based management which is more popular in Cloud deployments
    - Admin can now deploy PowerHA clusters with HMC using SSH or REST based communication

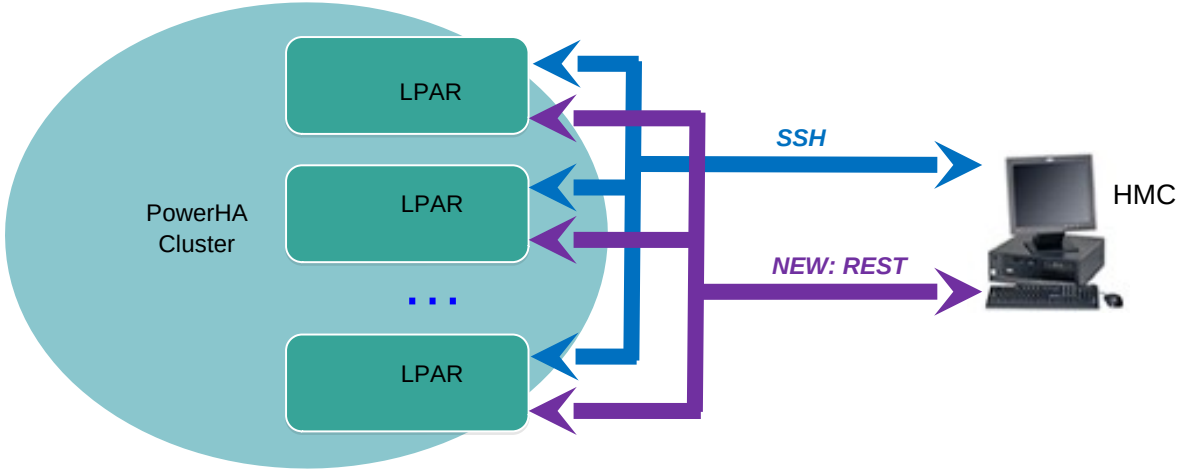

At-a-Glance Heath

- 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.
- here might be 100 lines relevant to the event how do I know which one I look at?

Using the UI to resolve PowerHA problems

Belisama

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

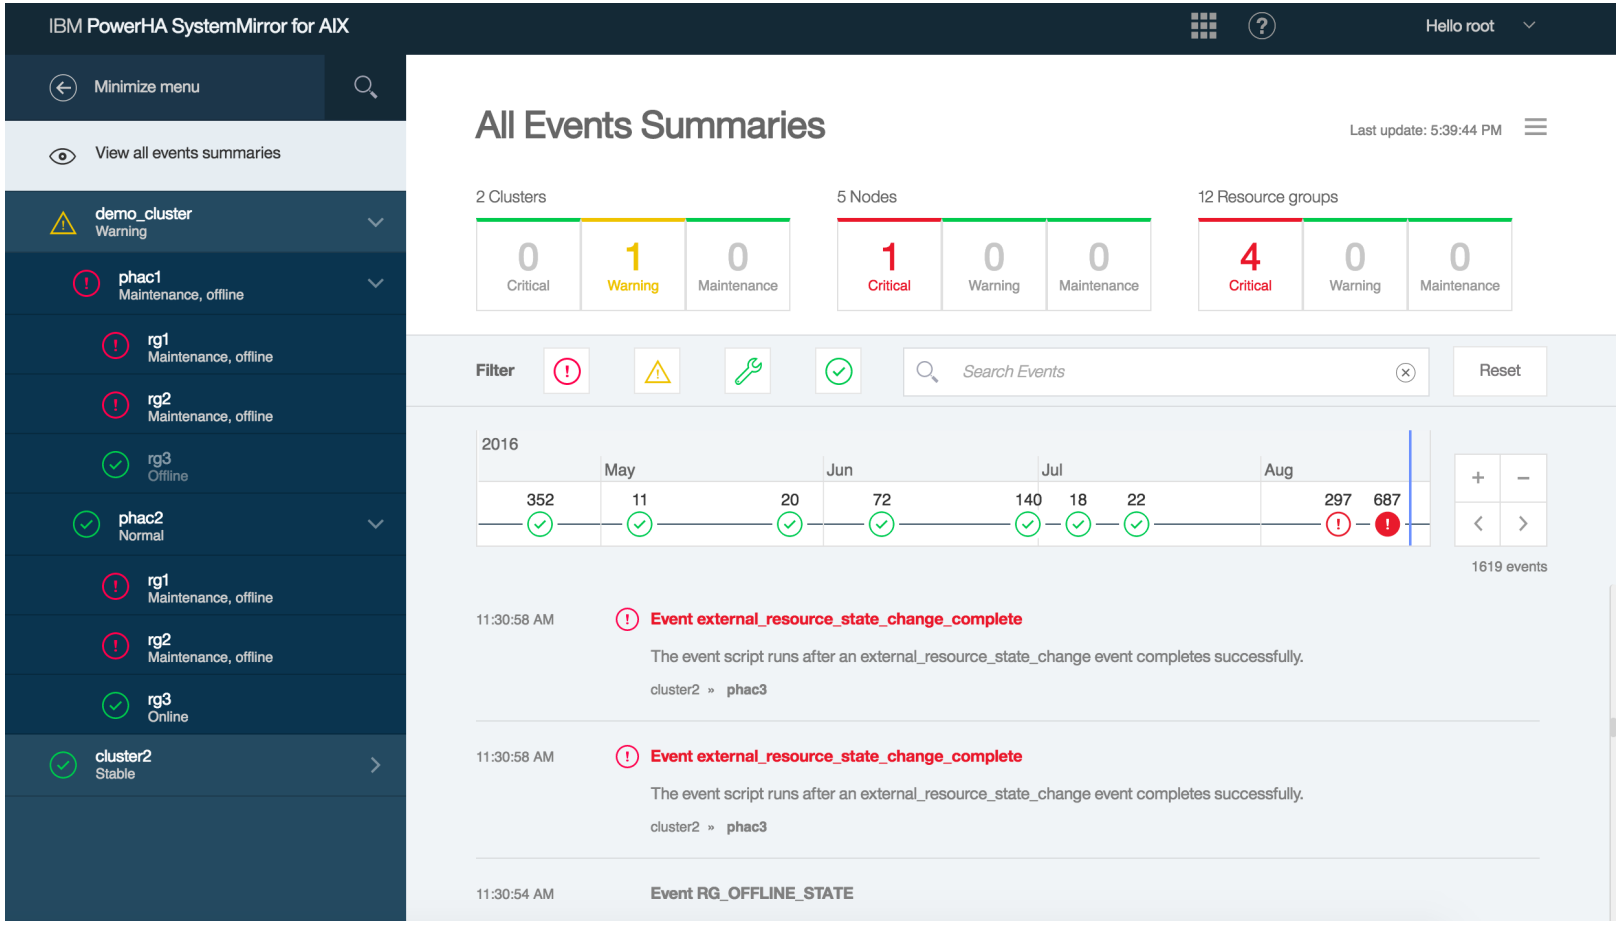

Using the UI to resolve PowerHA problems

Belisama

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

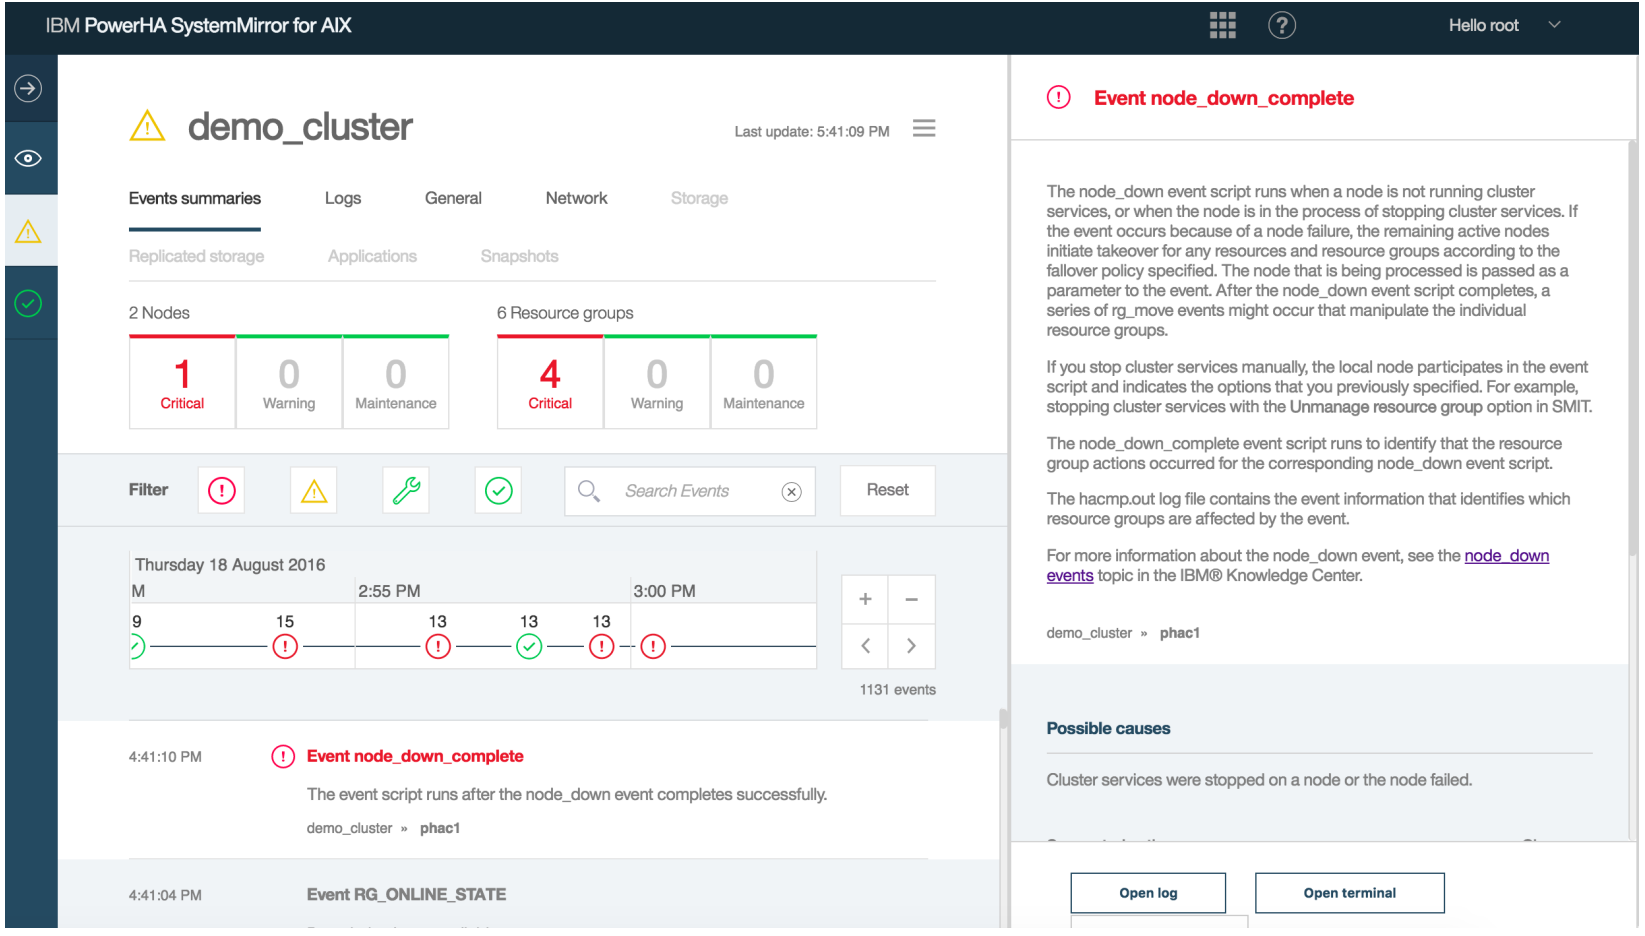

Using the UI to resolve PowerHA problems

Belisama

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

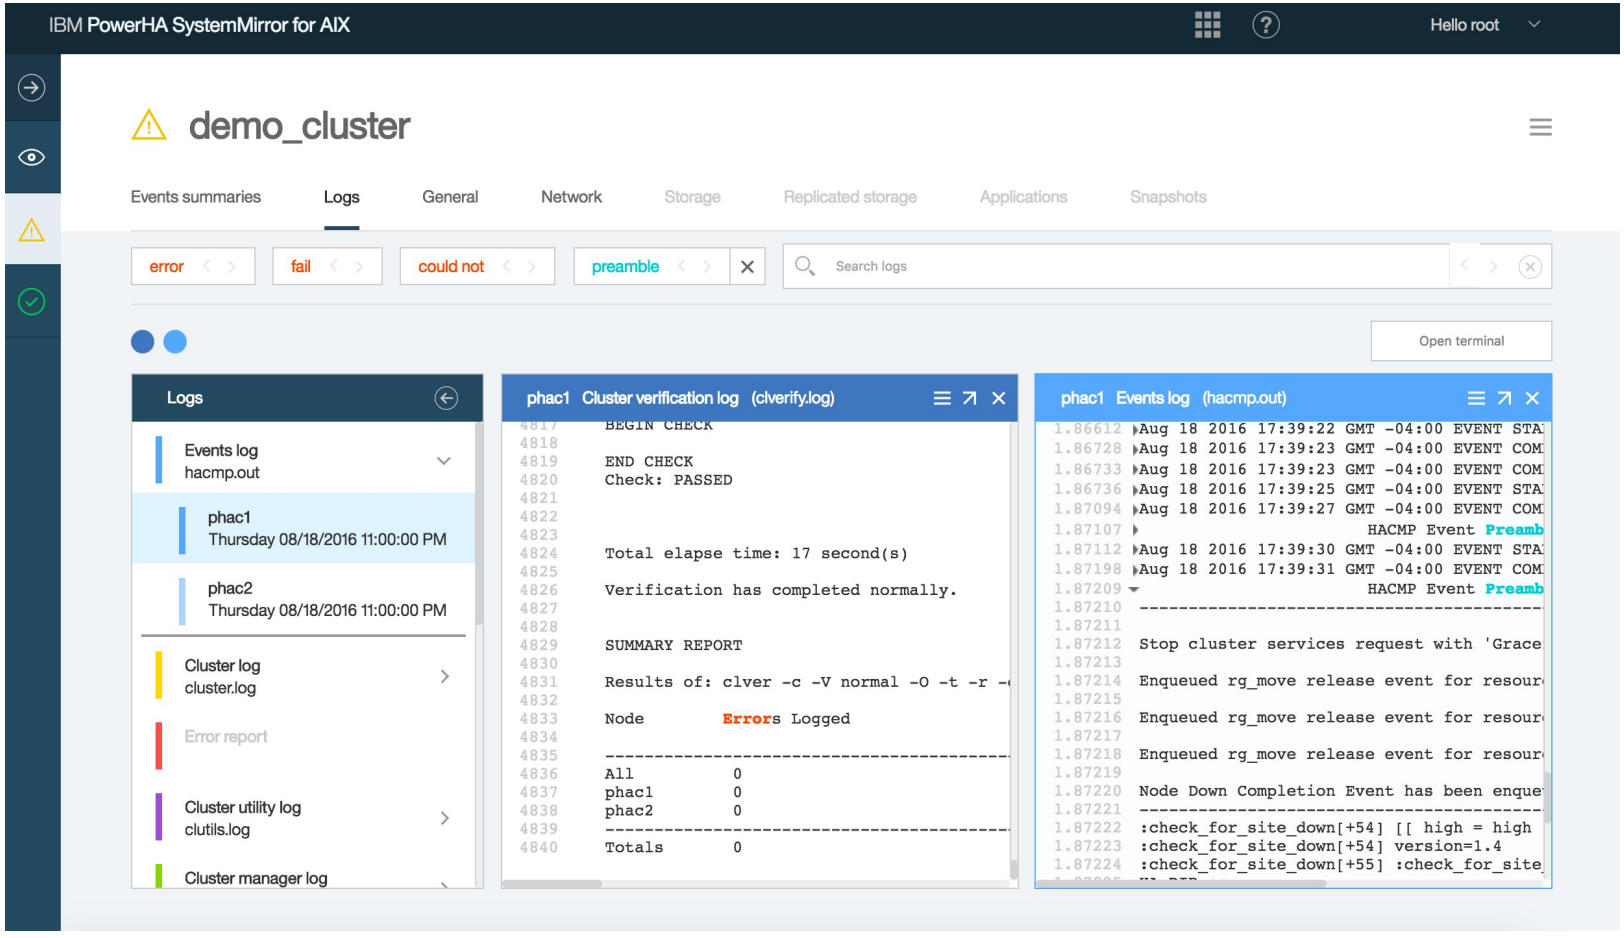

Using the UI to resolve PowerHA problems

Belisama

- 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.

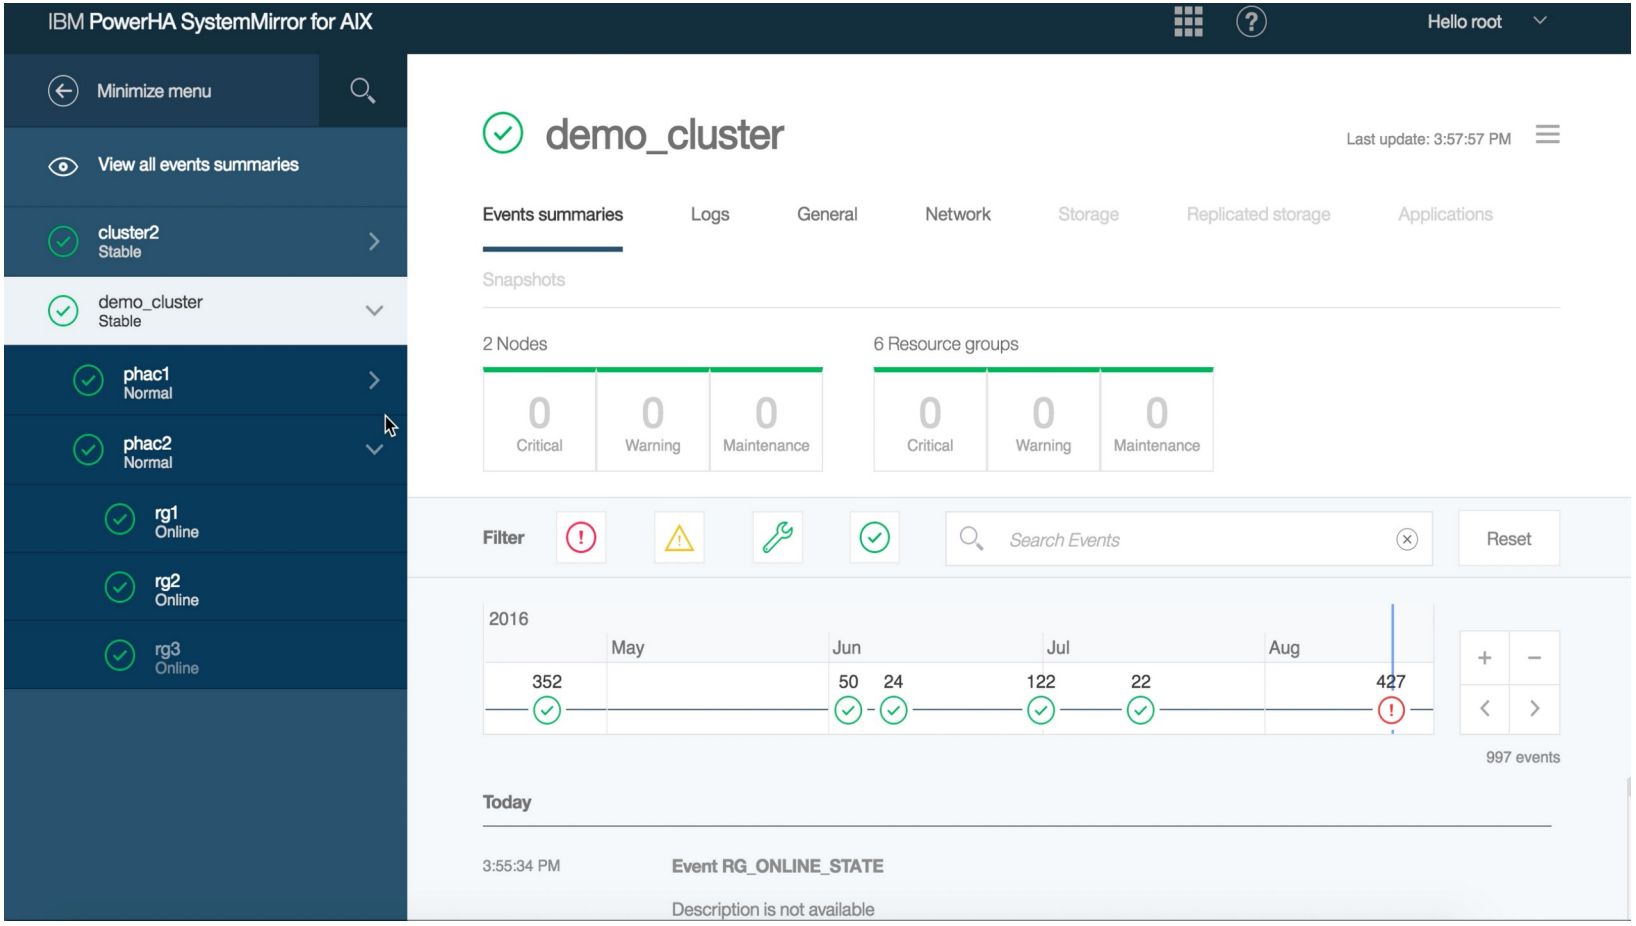

PowerHA SystemMirror split/merge policies

2015 PowerHA 7.2

	Split – Merge combination
Std Cluster	None - Majority
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

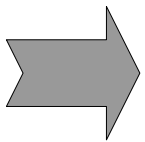

2016 PowerHA 7.2.1

	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	

- Quarantine Policies (standard/stretched/linked clusters)
  - Active Node Halt Policy (ANHP) : HMC based node shoot down
  - Disk fencing: SCSI-3 reserve based fencing

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

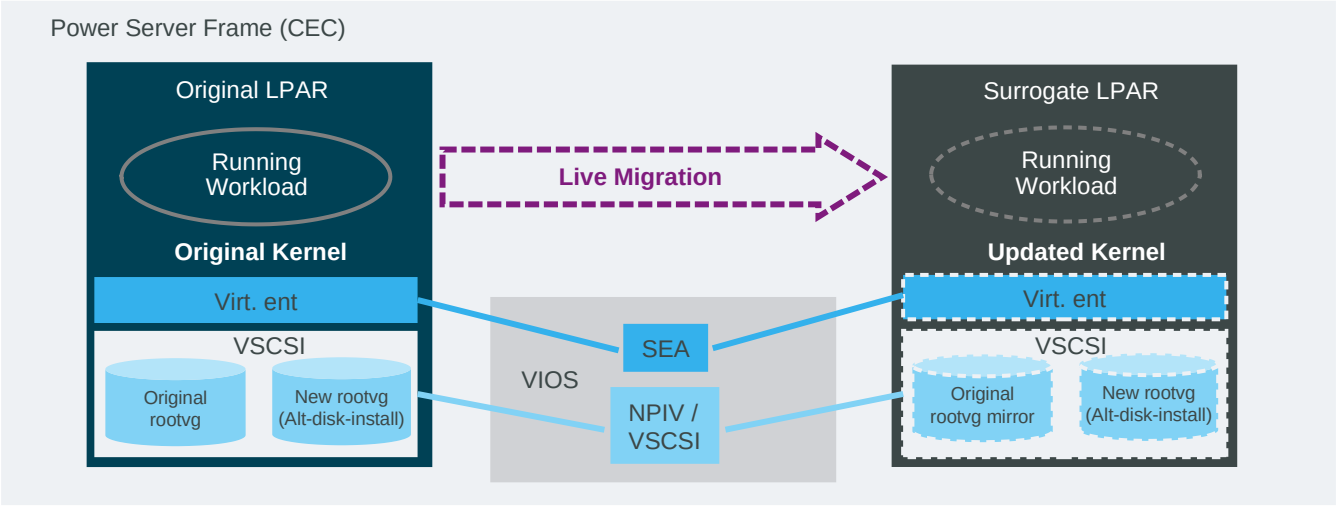

- Additional Software Requirements
  - Requires Power firmware 730_66 or later, VIOS 2.2.3.50 or later
  - PowerHA 7.2 & RSCT 3.2.1.0 are required for PowerHA support
  - PowerSC 1.1.4.0 for environments using PowerSC

- Additional Information
  - Deeper dive into AIX Live Update
    https://ibm.biz/BdHFkW
  - Hands on AIX Live Update article
    https://ibm.biz/BdHFtd

Capacity Backup For Power Enterprise Systems

New offering

- Eliminates many of the configuration requirements
- Provides support for the Power E870 and new Cloud models
- Provides lower hardware and hardware maintenance pricing for clients looking for HA/DR solution
- Fast failover to ***active*** memory on the CBU
  - To add 256GB of memory to the partition it can take 3 or 4 minutes, the more being added the longer the time

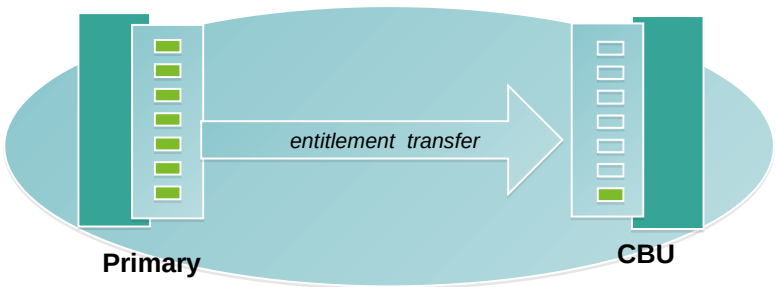

Power Systems CBU for Enterprise Servers Offering

Belisama

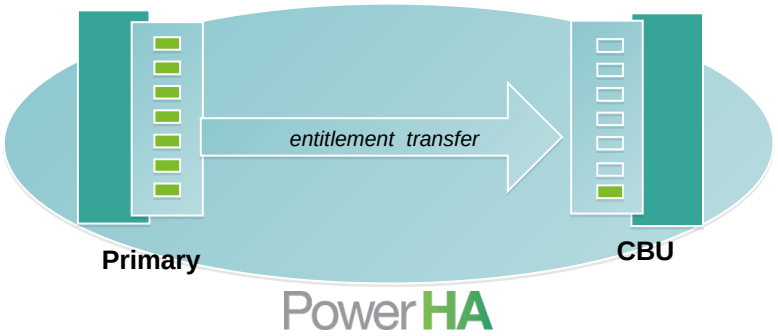

- All temporarily transferable entitlements must originate on the primary system and may not run concurrently on the primary system and the CBU system
- Subsequent to the initial workload deployment, production partitions may be moved to the CBU system for workload balancing etc
- The total number of processor entitlements running production across both servers can not exceed the original total licensed entitlements.
- Temporarily transferred entitlements return to the production server if either server leaves the registered pairing

Offering for

- Power System E880, E880C, E870 & E870C customers
- HA/DR deployments requiring fast failover via active standby memory

CBU offering Features:

- Deeply discounted processor nodes matching the installed production server processor nodes
- No charge active standby memory = 365xNx32 GB, where N is the number of active mobile cores on the production server **renewable annually**
- Mobile processor activations are transferred from production to CBU via Enterprise Pool transfers

Offering requirements overview

- Primary system and CBU system must be the same machine model (only one CBU to one production server for registration purposes but multiple production servers to one CBU is allowed.
- More than one Enterprise CBU to a primary or production server not allowed.
- Primary can be new or installed box, CBU must be a new box
- A minimum of one entitlement of AIX or IBM i & PowerHA on the CBU or if alternative HA/DR solution is used, as many IBM i or AIX entitlements as needed to support the workload (such as a middleware replication workload)
- 8 processor static activations on the CBU (no more no less)
- Minimum of 25% of DIMM memory active on the CBU
- No charge Memory ECOD days must be activated upon install of CBU system and remain active for 365 days
- Registration of primary system and CBU prior to CBU order fulfillment (registration validates configuration)Primary and CBU must be within the same enterprise

PowerVC: automated VM ‘remote’restart

Belisama

Automatically detect host failure and rebuild VMs on healthy hosts

- Shared storage data center solution
- Can be enabled/disabled at host group, host and VM level
- VMs are placed based on the host group’s placement policy
- Supported on PowerVM (NovaLink and HMC) and PowerKVM
- Works on AIX, Linux and IBM i VMs; requires Power 8