

Advanced  
Technology  
Group



# IBM Power Virtual Server &

# Resource Optimized High Availability

The image shows a perspective view of a server room with rows of server racks on both sides. A large, white, fluffy cloud is superimposed in the center of the room. Overlaid on the cloud is the IBM logo, which consists of eight horizontal blue stripes of varying lengths, forming the letters 'IBM'. The logo is centered horizontally and vertically within the cloud. The server racks are dark grey or black, and the floor is a light, reflective surface. The ceiling has recessed lighting fixtures.

IBM



---

# IBM - Advanced Technology Group

---

## *ISV on Power Team*



**Michael Eric Herrera**  
**HA & DR Solutions SME**  
mherrera@us.ibm.com



**Ralf Schmidt-Dannert**  
**Oracle on Power SME**  
dannert@us.ibm.com

# IBM Advanced Technology Group

Our mission is to provide **technical pre-sales** assistance to IBM and Business Partner hardware sellers within North America for **IBM Power and IBM Storage** opportunities

## What ATG offers:

- Competitive engagements (Server and storage)
- Proof of Concepts, Proof of Technologies, performance benchmarks
- Complex solution architecture and design
- Articulate IBM value propositions
- Establish credibility through best practices guidance
- Subject Matter Expertise to improve client experience on IBM platforms
- Customer presentations, Lunch & Learns, canned & live demonstrations
- System Health Check, best practices review
- Tools, collateral & cookbooks

## ATG Technical Leads:

### ▪ ISV on Power:

- Oracle: Ralf Schmidt-Dannert ([dannert@us.ibm.com](mailto:dannert@us.ibm.com))
- HA & DR: Michael Herrera ([mherrera@us.ibm.com](mailto:mherrera@us.ibm.com))
- SAP: Louis Lamprinakos ([lnlampri@us.ibm.com](mailto:lnlampri@us.ibm.com))
- SAS: Harry Seifert ([seifert@us.ibm.com](mailto:seifert@us.ibm.com))

### ▪ IBM Tape:

- Roy Peek ([rhpeek@us.ibm.com](mailto:rhpeek@us.ibm.com))

### ▪ IBM Storage:

- Carl Brown ([carltb@us.ibm.com](mailto:carltb@us.ibm.com))

## ATG team engaged via:

- IBM seller
- IBM Business partner

**I. Power Virtual Server – Deploying AIX Instances**

**II. ORACLE Database Installation Options**

**III. PowerHA SystemMirror Cluster Configuration**

**IV. PowerVS ROHA Demonstration**

**V. Lessons Learned & Wrap-Up**

---

# **Power Virtual Server**

---

# On-Premise Solutions

## Power Server Standalone



Traditional Infrastructure Deployment

*Consume entire Server as an application platform*

## Power Private Cloud with Dynamic Capacity



Pay only for what you use

*Cloud Capabilities & Advanced Monitoring*

- Performance and Scale
- Best in class Availability
- Flexibility & Agility
- Security
- Flexible options to lower cost
  - Optimize TCO with shared capacity
  - Low initial TCA with pay-per-use
- Simple agility to respond to business demands
- Metering by the minute
- Real time visibility and control
- Scale up or Scale out

SAN attached storage

# IBM Power Virtual Server

Infrastructure as a Service [IaaS]

## Power Virtual Server on IBM Cloud



Burst to Public Cloud

*Same mission Critical Infrastructure as on-prem, for Dev/Test, HA/DR, Modernization*

- VM as a Service
- Consistent architecture to On Prem infrastructure
- AIX, IBM i and Linux
- Access to other cloud services
- Expanding locations and capability

SAN attached storage

---

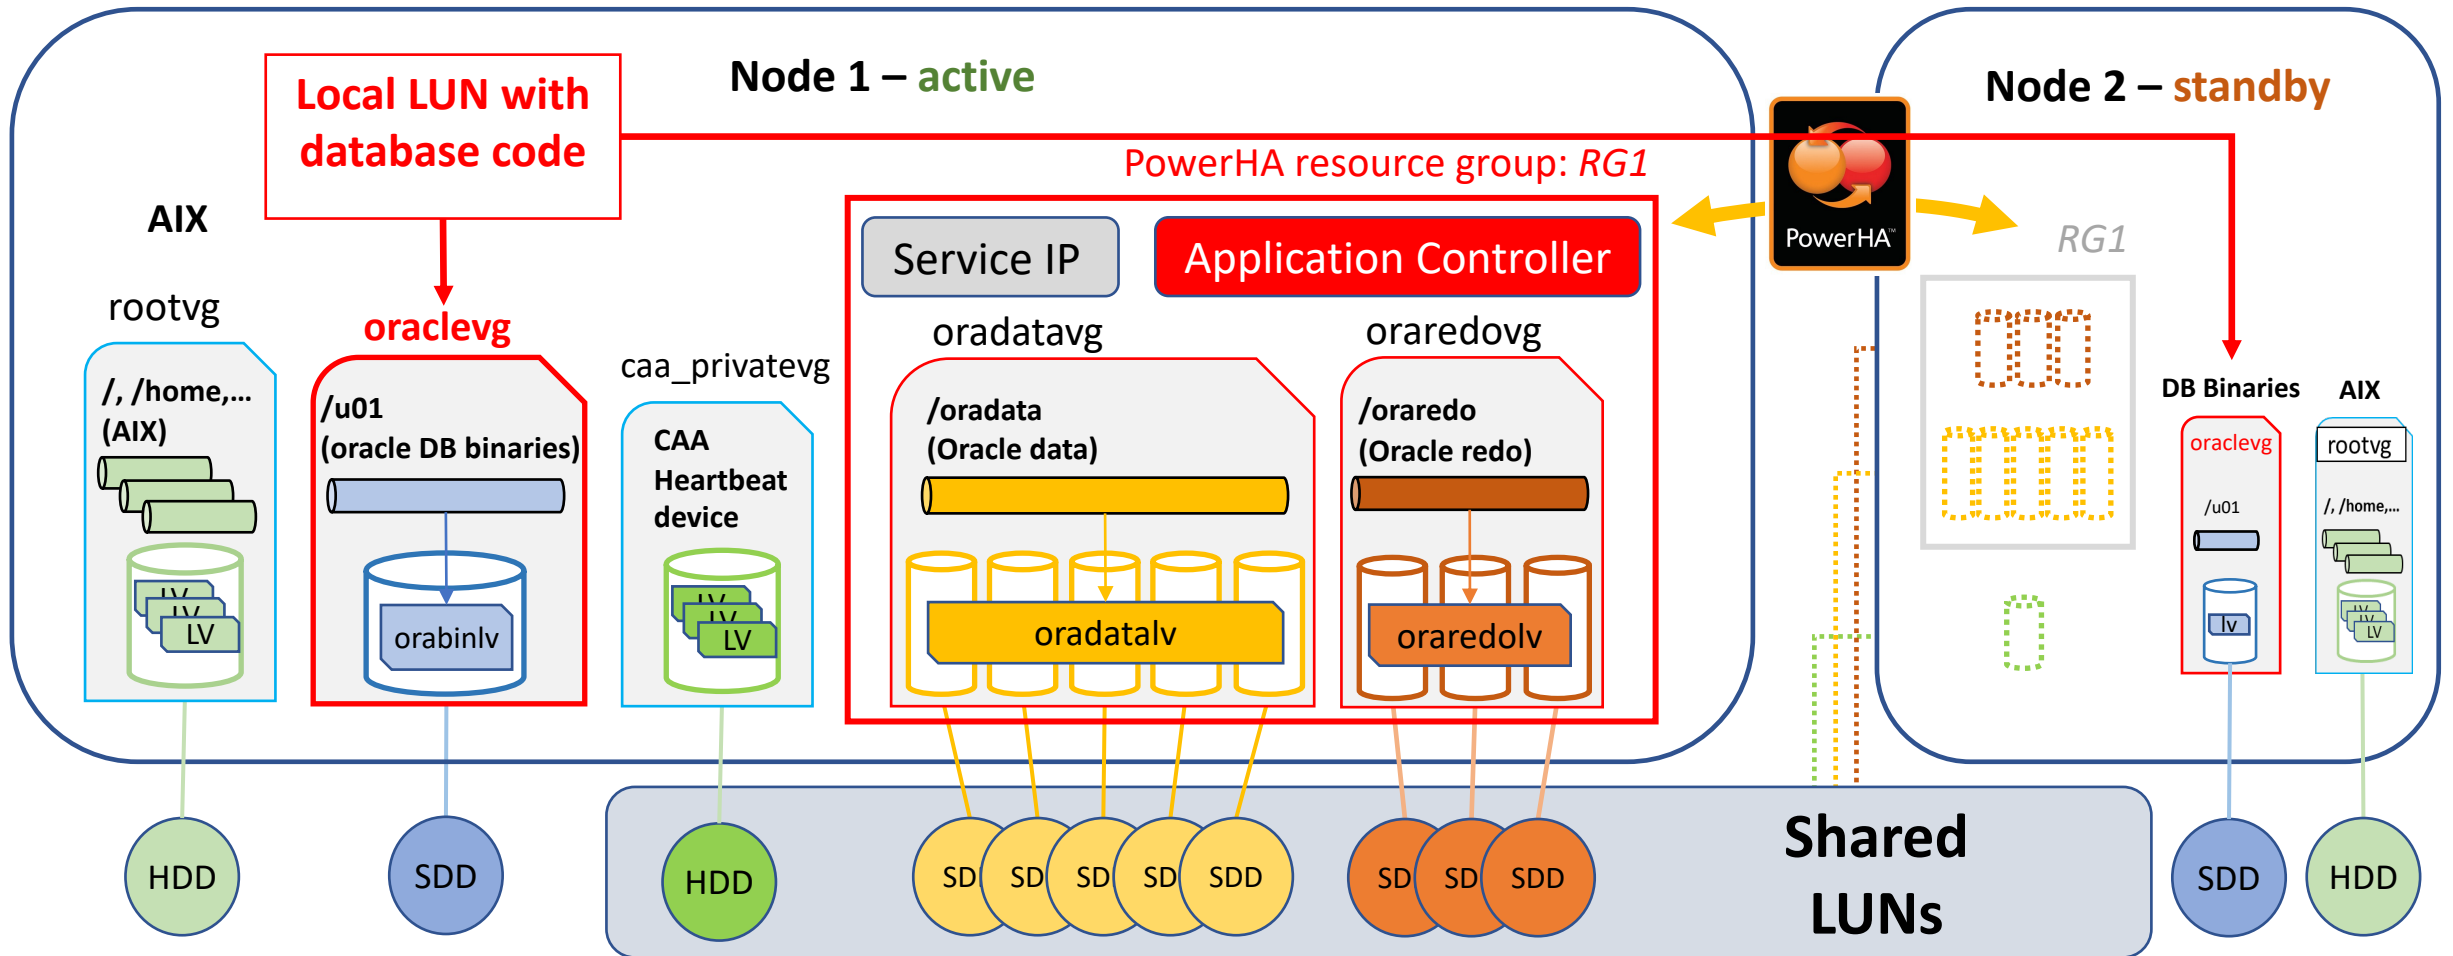
# **ORACLE DB** Deployment Options

---



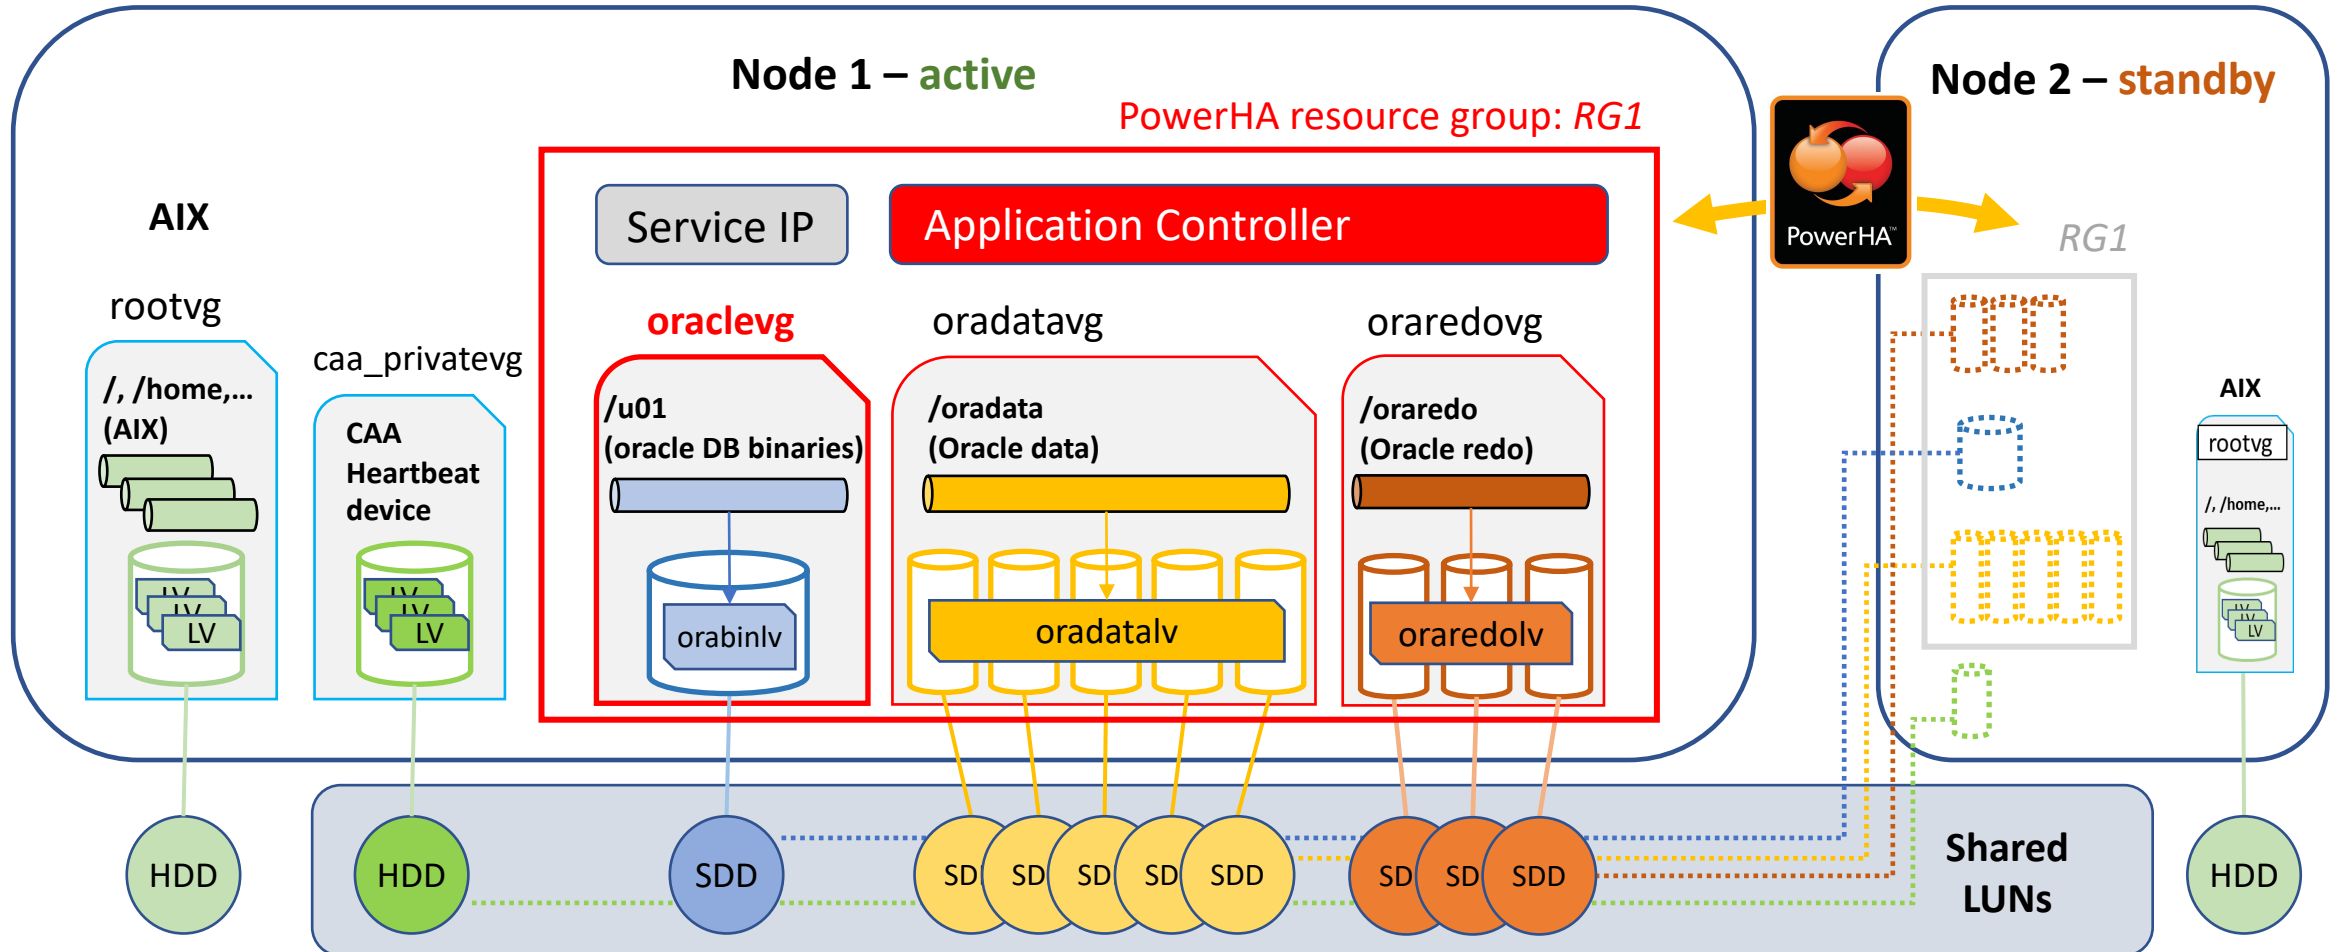
# Clustered ORACLE DB Deployment Model (active/passive)

[ Local Application Binaries ]



# Clustered ORACLE DB Installation (active/passive)

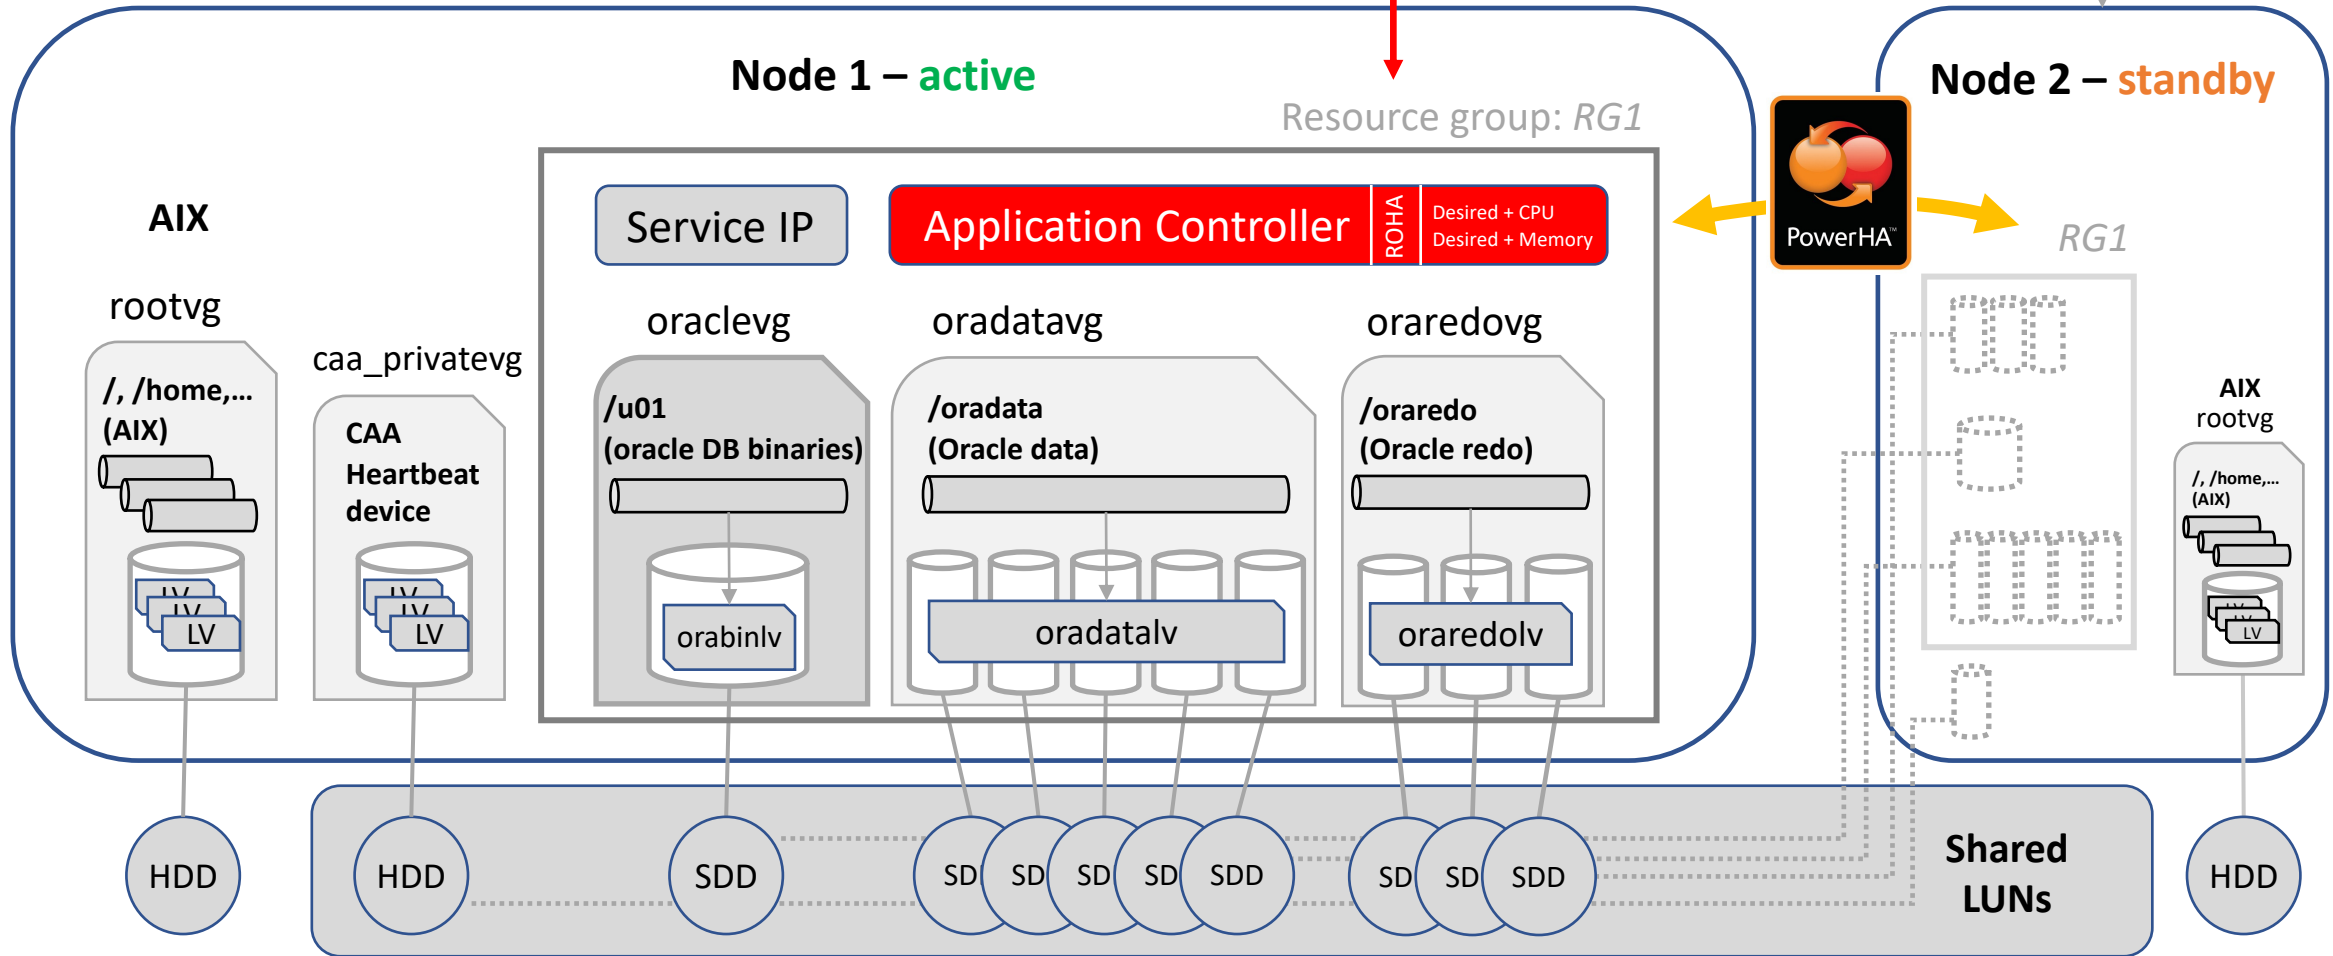
[ Shared Application Binaries ]

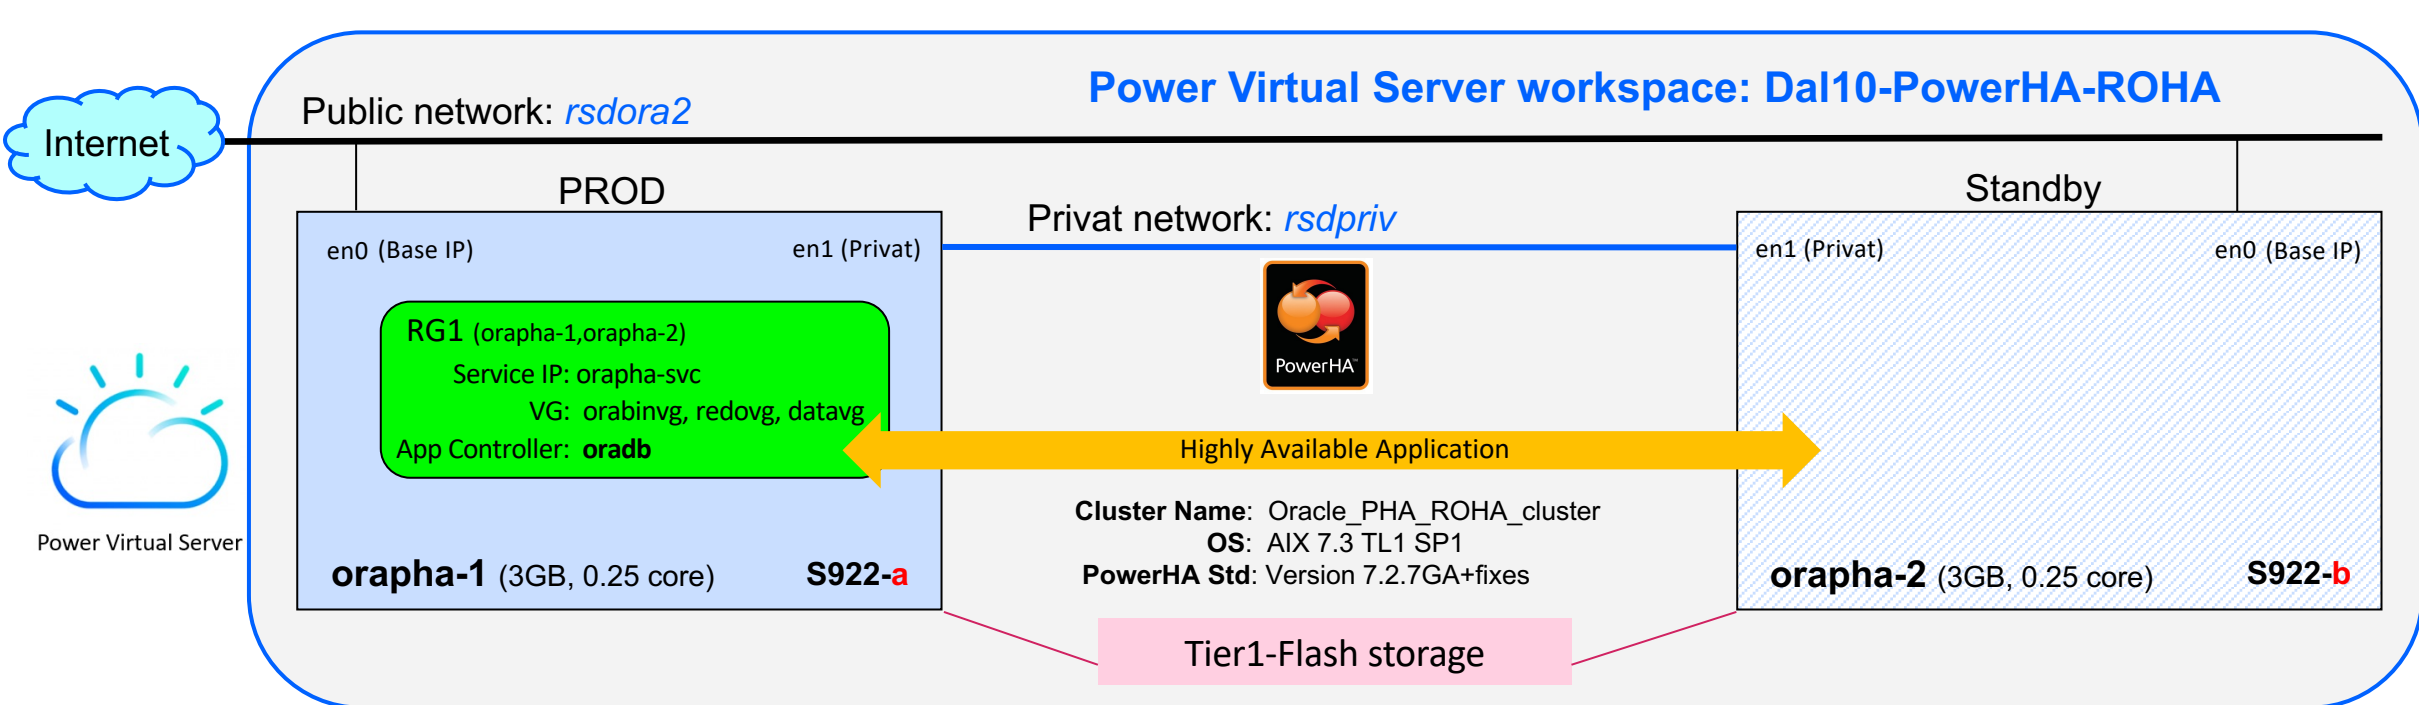


# Clustered Shared Binary ORACLE Installation

## [ PowerHA SystemMirror **ROHA** Functionality]

**ROHA** integration will modify CPU & Memory resources during application acquire | release





Create SSH key **rsdkey**

Performed in GUI

Find AIX 7.3 image

```
ibmcloud pi imglc | grep 7300
24083f6f-6ab7-4b59-bbd1-3ccb9e24a8db 7300-00-01 /pcloud/v1/images/...
```

Create an image copy

```
ibmcloud pi imgc 7300-00-01
```

Get the ID for the image copy

```
ibmcloud pi imgs
ID Name Address
5e218463-31b3-44d3-a860-79fb249d3937 7300-00-01 /pcloud/v1/cloud-instances/...
```

Create networks

(see next slide)

Create the 2 LPARs

```
ibmcloud pi inc orapha --replicants 2 --replicant-scheme suffix --image 5e218463-31b3-44d3-a860-79fb249d3937
--replicant-affinity-policy anti-affinity --network rsdora2 --network rsdpriv --memory 3 --processors 0.25
--processor-type shared --sys-type s922 --pin-policy hard --key-name rsdkey --storage-type tier1
```

**Cloud Shell  
CLI  
Operations**

# Network Details

# Cloud Shell CLI Operations

```
ibmcloud pi netcpu rsdora2 --dns-servers "8.8.8.8 9.9.9.9"  
ibmcloud pi netcpr rsdpriv --cidr-block 10.10.0.0/29 --ip-range "10.10.0.2-10.10.0.6"
```

Dal10-PowerHA-ROHA / Subnets [Learn more](#)

Search Refresh Create subnet +

Name	Type	IP ranges	VLAN ID
rsdora2	Public	192.168.230.10 – 192.168.230.14	2001
rsdpriv	Private	10.10.0.2 – 10.10.0.6	232

Items per page: 10 1–3 of 3 items 1 of 1 page

### Network interfaces

At least one interface, public or private, is required.

**Public networks**  
 On  
Extra steps required.

## Instance orapha-1

Name	IP address	External IP	Gateway	MAC address	VLAN ID	CIDR
rsdora2	192.168.230.12	<IP>	192.168.230.9	fa:35:ea:3d:6d:20	2001	192.168.230.8/29

Private networks

Search Attach existing network

Name	IP address	Gateway	MAC address	VLAN ID	CIDR
rsdpriv	10.10.0.3	10.10.0.1	fa:35:ea:3d:6d:21	232	10.10.0.0/29

[Detach](#)

# Server Placement Group

IBM Cloud

Search resources and products...

Dal10-PowerHA-ROHA / Virtual server instances

Servers **Server placement group**

Colocation policy All Search Create group +

Name	Number of VMs	Colocation policy
orapha-servergroup-1677248434	2	Different server

Dal10-PowerHA-ROHA / Server placement groups / orapha-servergroup-1677248434

Details

Name: orapha-servergroup-1677248434 Placement group ID: d37e066c-cd38-41f2-aafa-2dcf130d7bcb

Colocation policy: Different server

Virtual server instances in placement group

Name	IPs	Remove
orapha-1	192.168.230.12, 10.10.0.3	Remove ×
orapha-2	192.168.230.14, 10.10.0.5	Remove ×



# Storage Details

## Cloud Shell CLI Operations

```
ibmcloud pi in orapha-1 | grep "Storage Pool "
```

```
Storage Pool Tier1-Flash-2
```

```
ibmcloud pi volc --size 20 --storage-pool Tier1-Flash-2 stage
```

```
ibmcloud pi volat stage --instance orapha-1
```

```
ibmcloud pi volc --size 1 --shareable --storage-pool Tier1-Flash-2 CAA_HB
```

```
ibmcloud pi volat CAA_HB --instance orapha-1
```

```
ibmcloud pi volat CAA_HB --instance orapha-2
```

```
...
```

### Attached volumes

Volume affinity setting ⓘ  
Same storage pool

Search

Attach volume

Create volume +

<input type="checkbox"/>	Name	Size	Tier	Storage pool	Shareable	Bootable	
<input type="checkbox"/>	oradata-5	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oradata-4	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oradata-3	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oradata-2	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oradata-1	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oraredo-3	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oraredo-2	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	oraredo-1	5 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	orabin	75 GB	Tier 1	Tier1-Flash-2	On	Off	Detach 🗑️
<input type="checkbox"/>	stage	20 GB	Tier 1	Tier1-Flash-2	Off	Off	Detach 🗑️
<input type="checkbox"/>	CAA_HB	hdisk1 in each node	1 GB	Tier 1	Off	Off	Detach 🗑️
<input type="checkbox"/>	orapha-1-e7d43bdd-0000211c-boot-0	25 GB	Tier 1		Off	On	Detach 🗑️

**ORACLE DB  
Volumes**

**CAA Heartbeat Device**

**AIX Boot**

# VM Pinning

Power Systems Virtual Server

Dal10-PowerHA-ROHA / Virtual server instances / orapha-1

Active [Refresh] [Delete] VM actions

**Virtual server instance details** [Edit details](#)

Name orapha-1	Machine type s922	Core type Uncapped shared processor	Placement group <a href="#">orapha-servergroup-1677248434</a>
ID e7d43bdd-1acb-4300-807c-0affdd05ab90	Boot image 7300-00-01	Size 0.25 cores   5 GiB	
Date created February 24, 2023 at 8:20:36 AM	Virtual server pinning Hard	Shared processor pool None	

Dal10-PowerHA-ROHA / Virtual server instances / orapha-1

**Virtual server instance details**

Name orapha-1	Machine type s922	Core type Uncapped shared processor
ID e7d43bdd-1acb-4300-807c-0affdd05ab90	Boot image 7300-00-01	Size 0.25 cores   5 GiB
Date created February 24, 2023 at 8:20:36 AM	Virtual server pinning Hard	Shared processor pool None

System reference code

SRC	Timestamp
00000000	April 6, 2023 at 9:07:31 AM

### Edit virtual server instance details

Changes will be reflected in your monthly billing.

Name  
orapha-1

Virtual server pinning ⓘ

VM pinning allows you to control the movement of VMs during disasters and other restart events.

Hard pinning is recommended for applications with serial number based licenses.

[Learn more](#)

Shared uncapped  
 Shared capped  
 Dedicated

Cores ⓘ

0.25 - | +

---

# IBM Cloud Shell

- Fast, efficient and can be automated! -
-

```
orapha-1:/usr/local/hascripts# cat jfs2db_start.ksh
```

```
#!/bin/ksh
```

```
ORA_HOME=/u01/app/oracle/product/19c
```

```
ORA_OWNER=oracle
```

```
wall 'STARTING application oradb'
```

```
# Start the Oracle databases:
```

```
# This also starts the listener!
```

```
# The following command assumes that the oracle login
```

```
# will not prompt the user for any values
```

```
# Will start any DB listed in /etc/oratab with "Y" as last entry in line
```

```
su - $ORA_OWNER -c "$ORA_HOME/bin/dbstart $ORA_HOME"
```

```
wall 'Oracle database started'
```

```
exit 0
```

## ORACLE Database Start Logic

```
orapha-1:/usr/local/hascripts# cat jfs2db_stop.ksh
```

```
#!/bin/ksh
```

```
ORA_HOME=/u01/app/oracle/product/19c
```

```
ORA_OWNER=oracle
```

```
wall 'STOPPING application oradb'
```

```
# Stop the Oracle databases:
```

```
# This also stops the listener!
```

```
# The following command assumes that the oracle login
```

```
# will not prompt the user for any values
```

```
# /etc/oratab line with "Y" for DB instances to be started!
```

```
su - $ORA_OWNER -c "$ORA_HOME/bin/dbshut $ORA_HOME"
```

```
wall 'Oracle DB stopped'
```

```
exit 0
```

## ORACLE Database Stop Logic

---

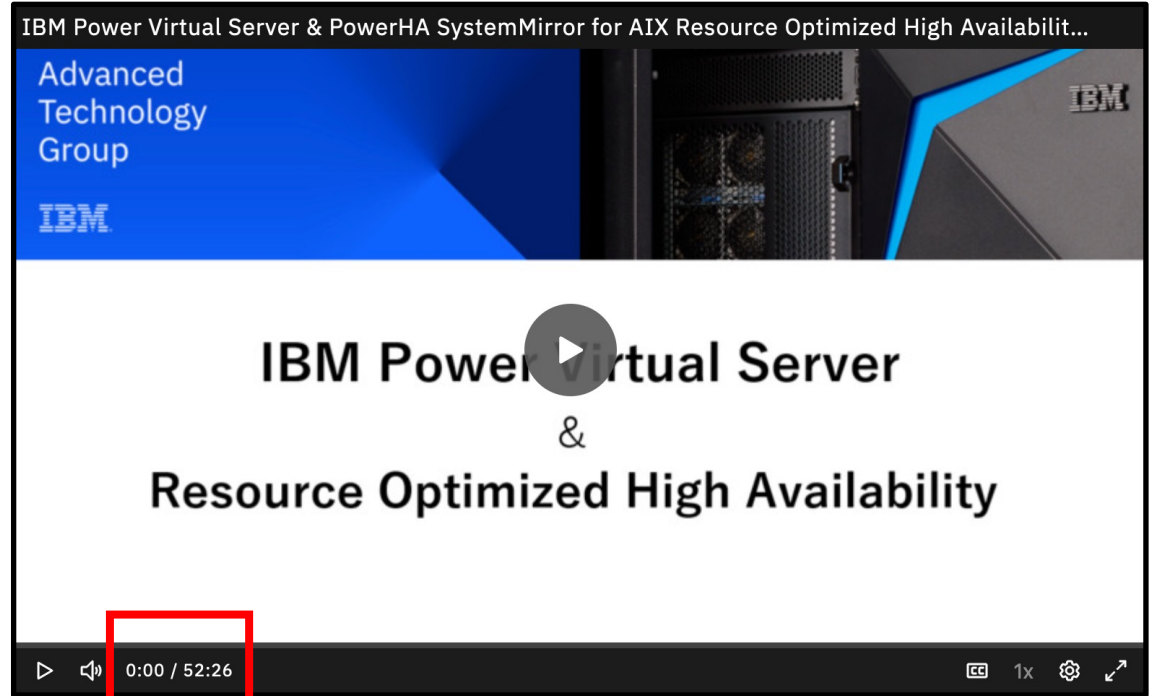
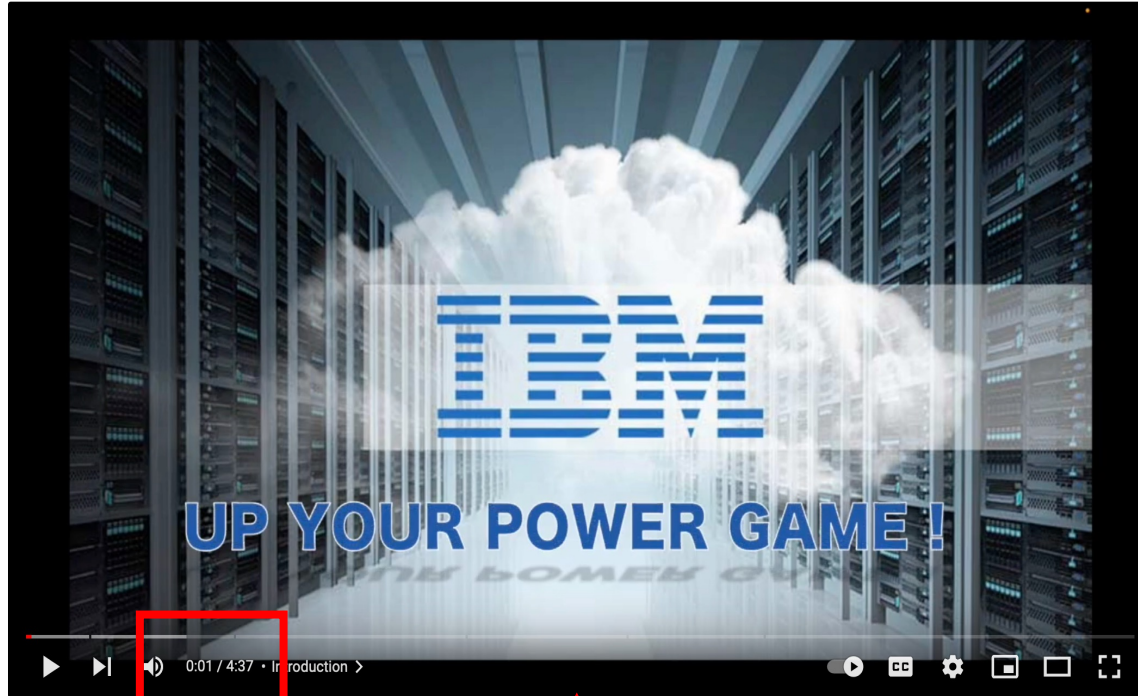
# **PowerHA SystemMirror & ROHA**

---

# Location of Video Intellectual Capital

**IBM Media Center Channel – ISV on Power** [Internal & External Access]

<https://mediacenter.ibm.com/category/ISV+on+Power/248395963>



**YouTube Channel:**

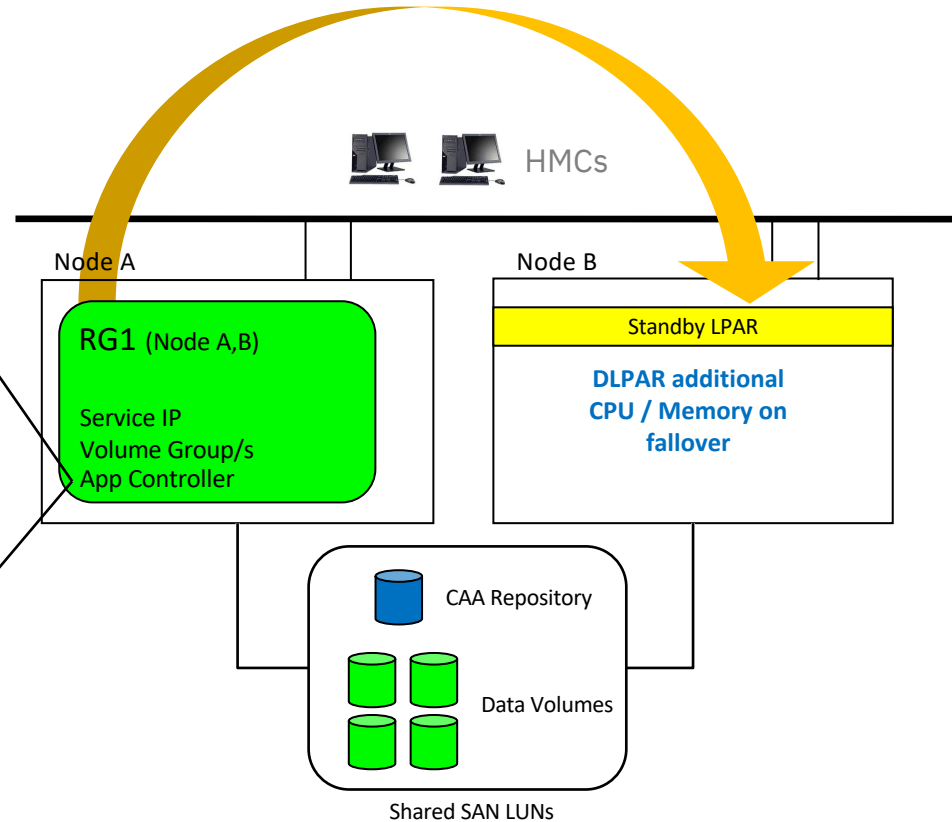
<https://www.youtube.com/@MrPowerHA/videos>



# Resource Optimized High Availability (ROHA)


## Local Integrated DLPAR Capability

```
HACMPserver:  
name = "App1"  
start = "/usr/local/hascrpts/startapp1.sh"  
stop = "/usr/local/hascrpts/stopapp1.sh"  
min_cpu = 0  
desired_cpu = 0  
min_mem = 0  
desired_mem = 0  
use_cod = 0  
min_procs = 0  
min_procs_frac = 0  
desired_procs = 0  
desired_procs_frac = 0  
start_mode = ""  
use_desired = 0  
optimal_mem = 0  
optimal_mem_frac = 0  
optimal_procs = 0  
optimal_proc_units = 0  
optimal_proc_units_frac = 0  
optimal_vprocs = 0  
cpu_usage_monitor = "yes"  
process_to_monitor_cpu = ""  
cpu_usage_monitor_interval = 10
```



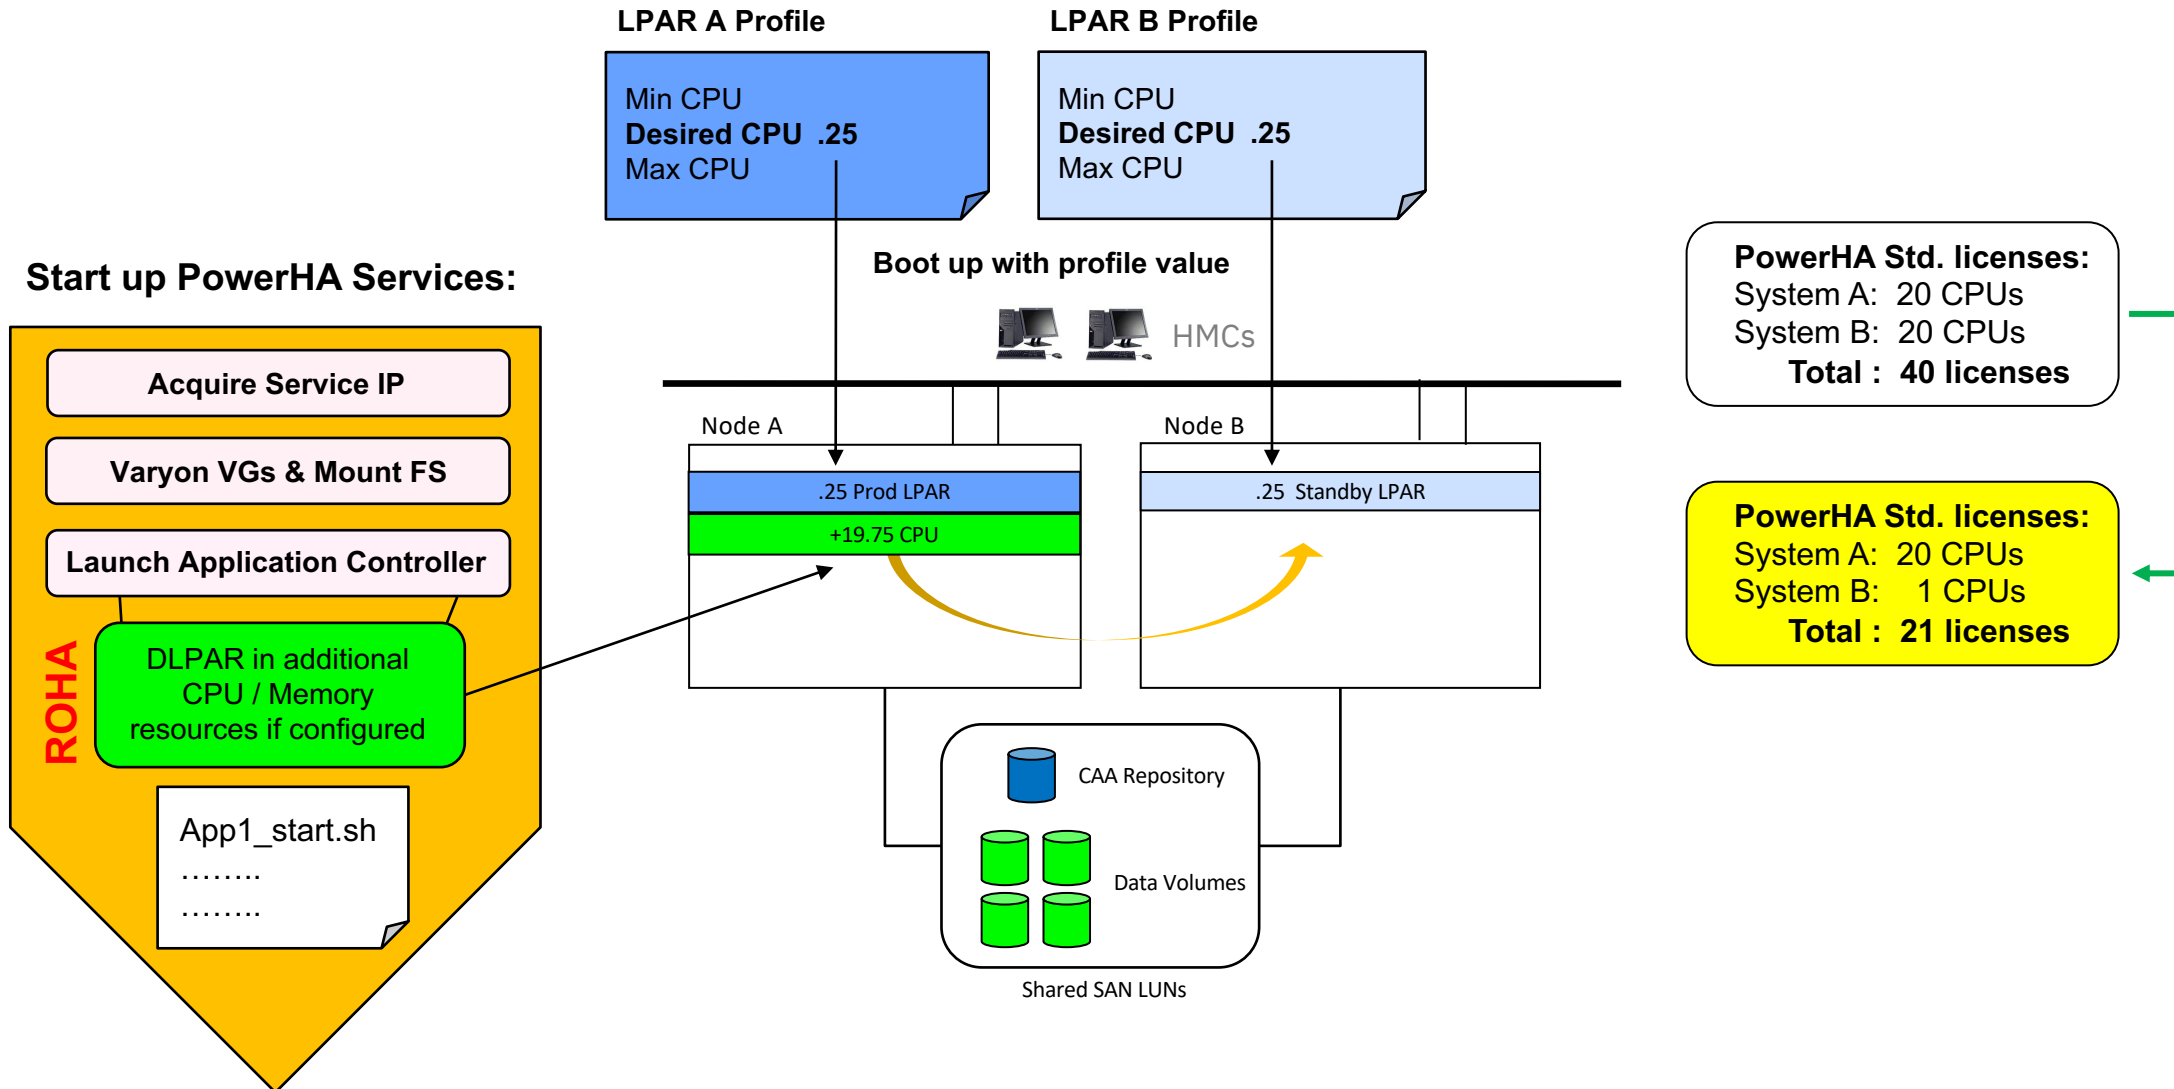
### CPU & Memory Resources

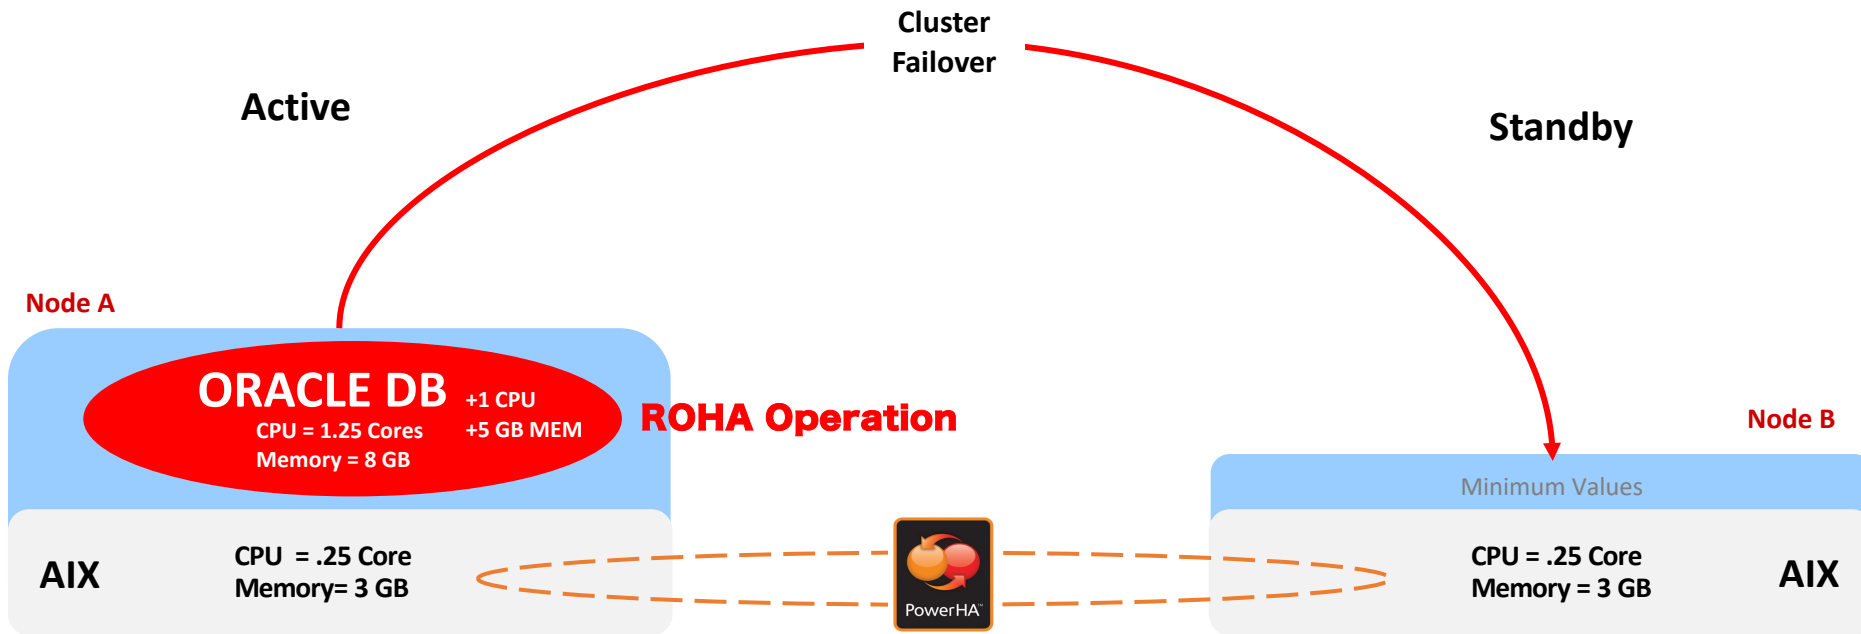
1. Free Shared Processor Pool
2. CoD Resources
3. PEP 1.0 mobile Activations

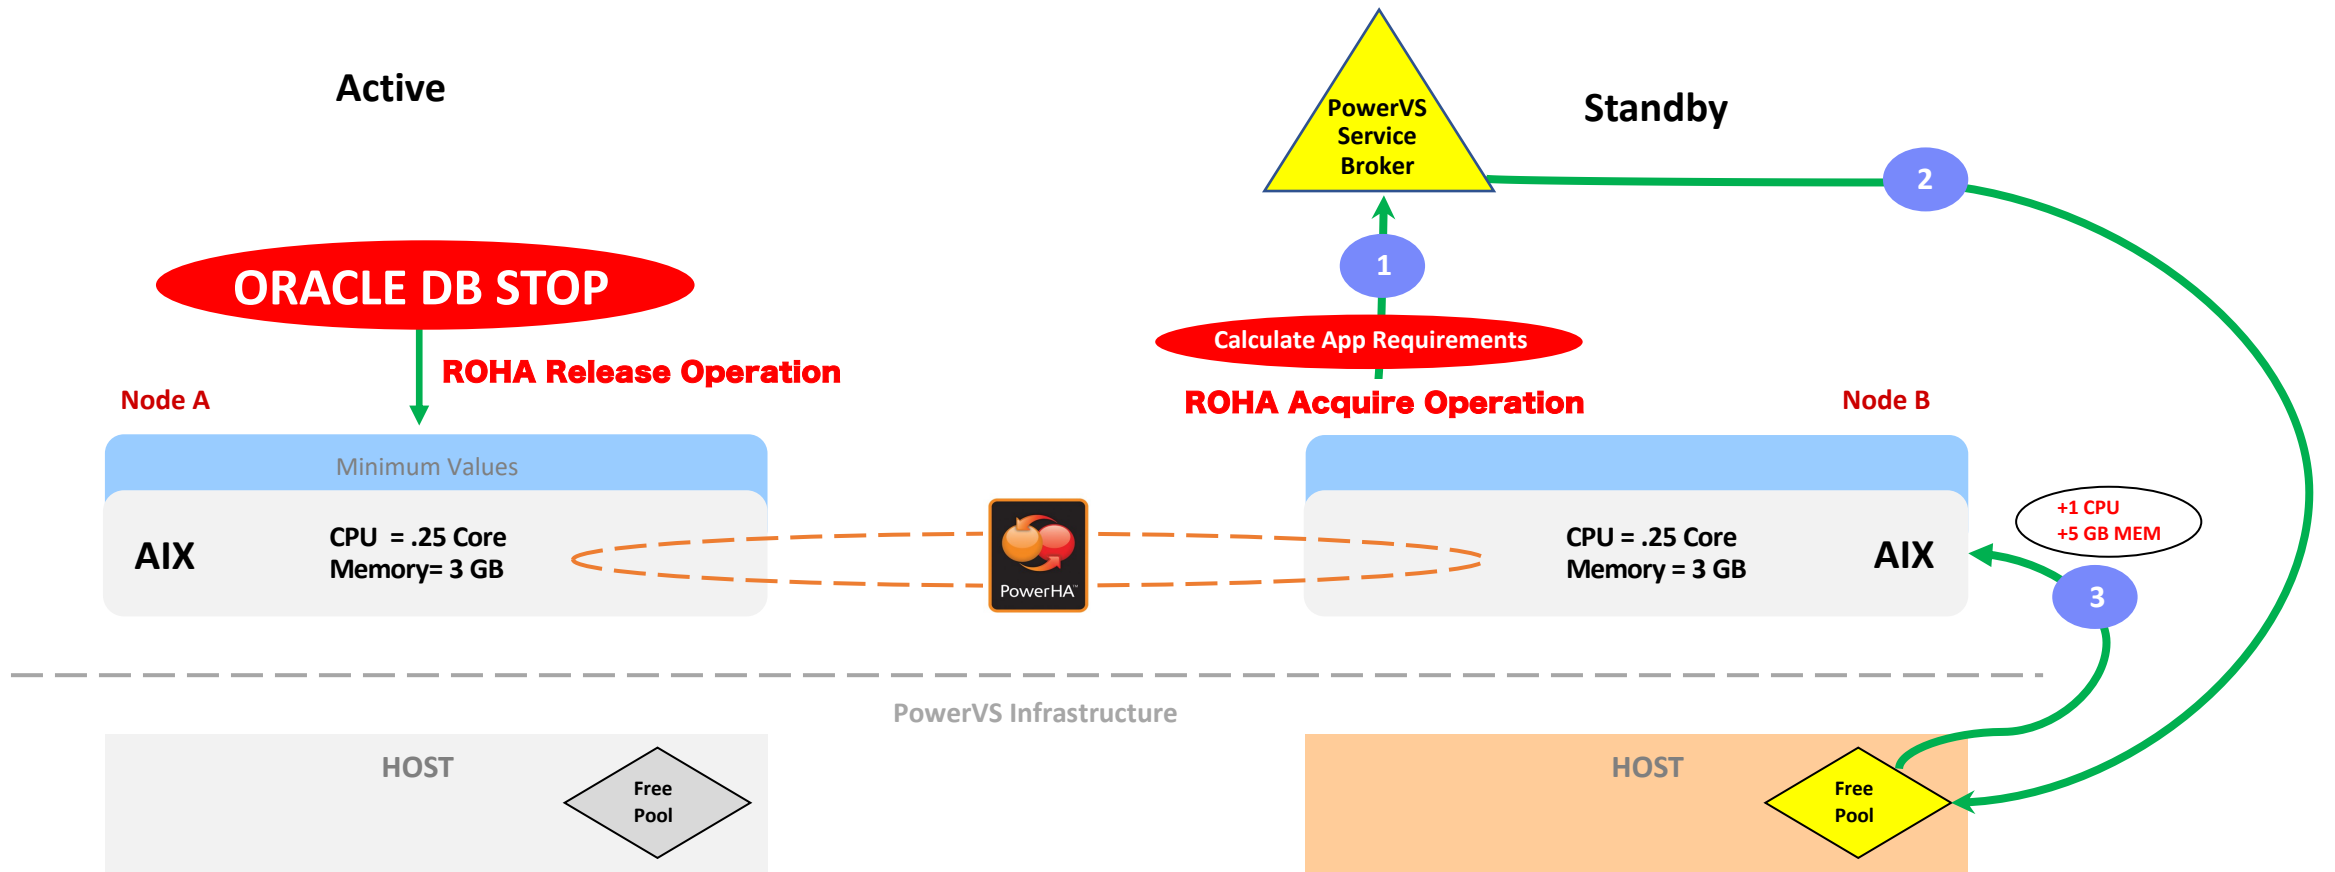


# PowerHA SystemMirror - Integrated DLPAR Scenario

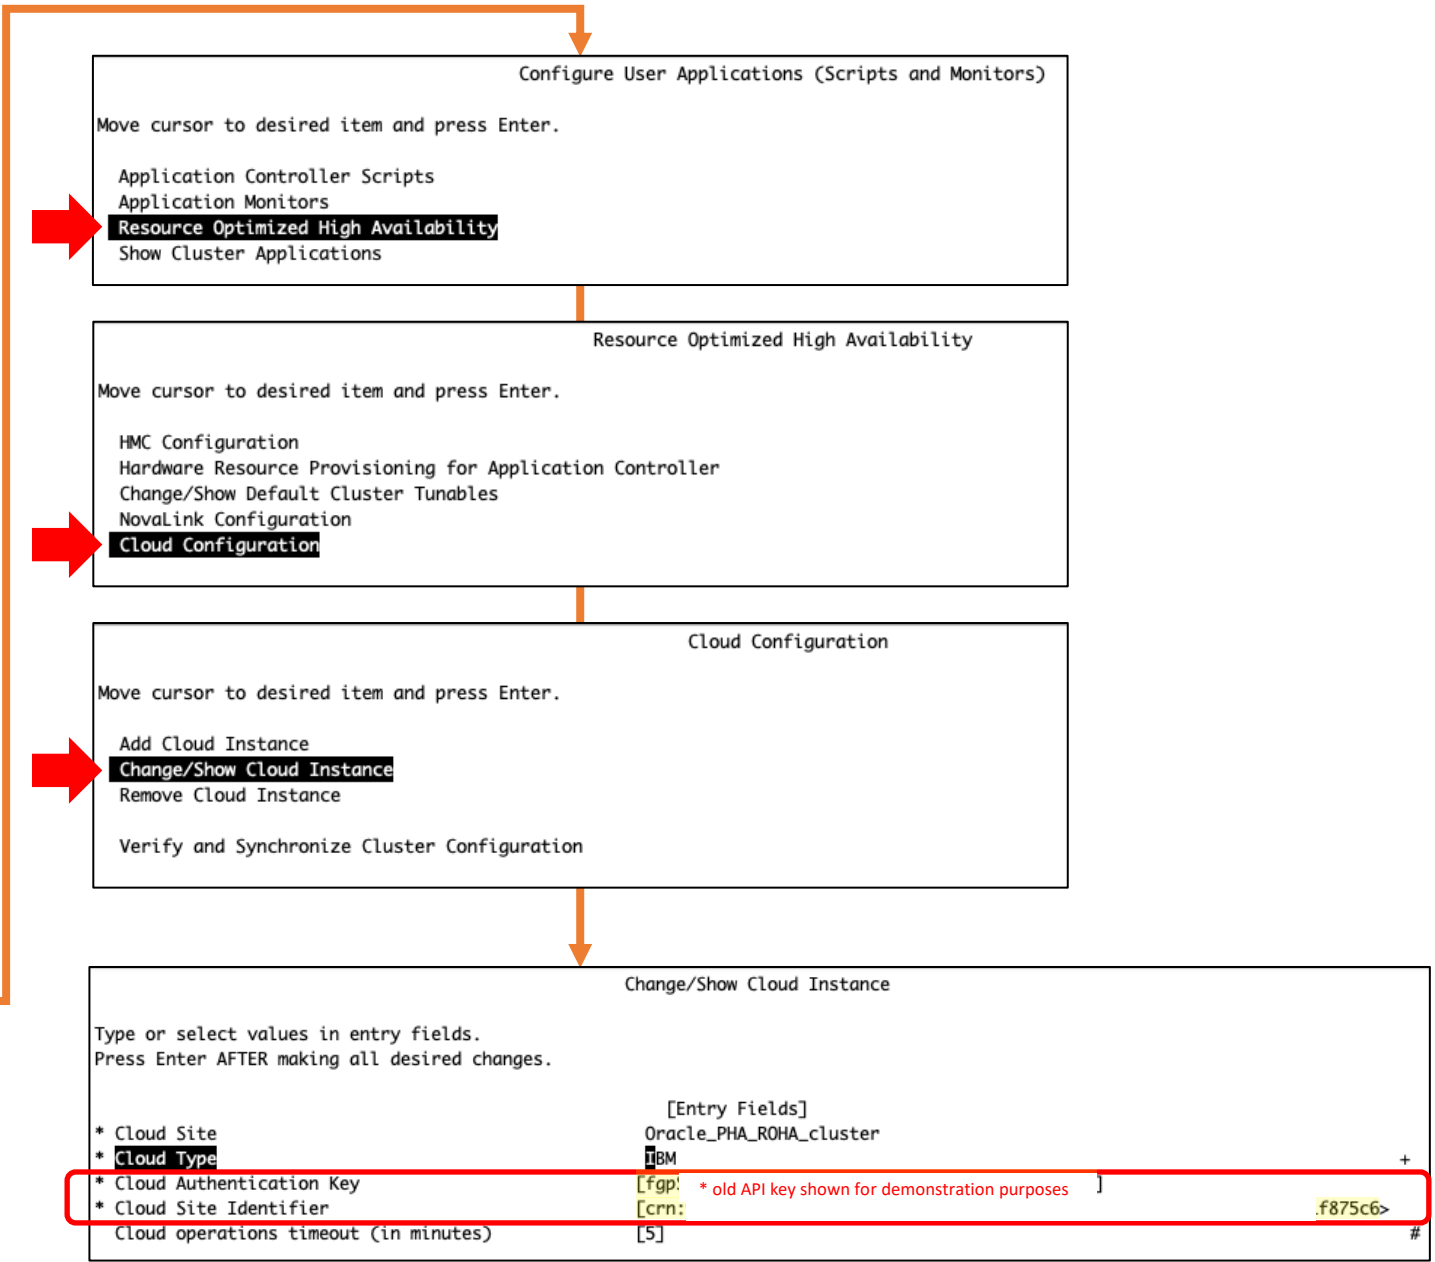
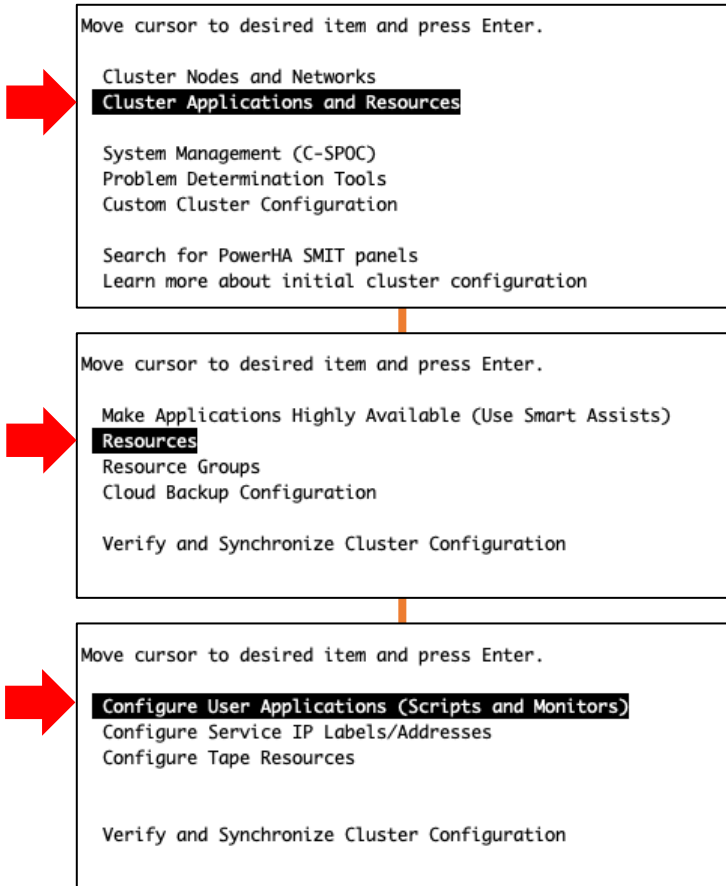
i.e. Highly Available Application that requires 20 CPUs







# # smitty sysmirror



IBM Cloud Search resources and products...

## Workspaces

A workspace is a free working environment tied to a single data center where compute, networking and storage resources are deployed. Resources cannot be moved or shared between workspaces.

All locations  Create +

**Limited permissions**  
Looks like you don't have permission to create, update, or delete a workspace in this account. Contact the account owner to be assigned an Editor or Administrator role on a resource group in the account. [Learn more](#)

Workspace	Location	Date created	Created by	CRN
<b>Dal10-PowerHA-ROHA</b> (Active)	Dallas 10	2/23/2023	[Redacted]	...-f3d451718e41
HMCS_PVS_OCP (Active)	Montreal 01	5/13/2022	[Redacted]	...-995f95401d24
Power Systems Virtual Server-la... (Active)	London 06	2/10/2021	[Redacted]	...-7d63cb531bf9
Power Systems Virtual Server-ct... (Active)	Dallas	12/8/2020	[Redacted]	...-71a4448a5db9

**1) Store the "CRN"  
Cloud Resource Name**

IBM Cloud Search resources and products...

## API keys

Create, view, and work with API keys that you have access to manage. IBM Cloud API keys are associated with a user identity and can be used to access cloud platform and classic infrastructure APIs, depending on the access that is assigned to the user. The following table displays the API keys created in this account. [Learn more](#).

Looking for more options to manage API Keys? Try [IBM Cloud® Secrets Manager](#) for creating and leasing API keys locally and storing them securely in your own dedicated instance.

View: My IBM Cloud API keys

API keys associated with a user's identity have the same access that the user is assigned across all accounts. To update the access for an API key, assign or remove access for the user.

Status	Name	Description	Date Created
🔒	ROHA_API_key	ROHA Testing API key	2023-02-27 15:30 GMT

API key successfully created

Copy the API key or click download to save it. You won't be able to see this API key again, so you can't retrieve it later. The API key is no longer displayed after 267 seconds.

API key  
.....

Copy Download

**2) Store the API Key**

JSON	Raw Data	Headers
Save Copy Collapse All Expand All Filter JSON		
name:	"ROHA_API_key"	
description:	"ROHA Testing API key"	
createdAt:	"2023-02-27T15:30+0000"	
apikey:	"wCWA	tgS"

\* old API key shown for demonstration purposes





---

# **ROHA Demonstration**

---

```

topas_nmon -g=disk-Groups -Host=orapha-1 -Refresh=2 secs -14:19.24
Resources
OS has 8 PowerPC_POWER9 (64 bit) CPUs with 8 CPUs active SMT=8
CPU Speed 2300.0 MHz SerialNumber=784B381 MachineType=IBM,9009-22G
Logical partition=Dynamic HMC-LPAR-Number&Name=4,orapha-1-e7d43bdd-0000211c
AIX Version=7.3.1.1 TL01 Kernel=64 bit Multi-Processor
Power Saving=Dynamic-Performance CPU-Current-Speed 3450.0 MHz
Hardware-Type(NIM)=CHRP=Common H/W Reference Platform Bus-Type=PCI
CPU Architecture =PowerPC Implementation=POWER9 SubProcessor Mode=Unknown
CPU Level 1 Cache is Combined Instruction=32768 bytes & Data=32768 bytes
Level 2 Cache size=not available Node=orapha-1
Event= 0 --- --- SerialNo Old=--- Current=C4B381 When=---

```

#1

## Node 1 - NMON Outputs

```

topas_nmon -3=Top-by-CPU-use -Host=orapha-2 -Refresh=2 secs -14:24.47
Shared-CPU-Logical-Partition
Partition:Number=4 "orapha-2-c857277c-0000211f"
Flags:LPARed DRable SMT Shared UnCapped Migratable Not-Donating AMSable.
Summary: Entitled= 0.25 Used 0.01 ( 2.8%) 0.0% of CPUs in System
- No Shared CPU Pool Util. Authority/Allow Perf. Info. Collect
CPU-Stats----- Capacity----- ID-Memory-----
max Phys in sys 24 Cap. Processor Min 0.25 SPLPAR Group:Pool 32772:0
Phys CPU in sys 24 Cap. Processor Max 2.00 Memory(MB) Min:Max 3072:49152
Virtual Online 1 Cap. Increment 0.01 Memory(MB) Online 3072
Logical Online 8 Cap. Unallocated 0.00 Memory Region LMB 256MB min
Physical pool 18 Cap. Entitled 0.25 Time-----Seconds
SMT threads/CPU 8 -MinReqVirtualCPU 0.05 Time Dispatch Wheel 0.0100
CPU-----Min-Max Weight-----MaxDispatch Latency
Virtual 1 5 Weight Variable 128 Time Pool Idle
Logical 1 40 Weight Unallocated 0 Time Total Dispatch

Event= 0 --- --- SerialNo Old=--- Current=C4BCD1 When=---

Shared_Pools MaxPoolCapacity=18.00 No Shared CPU Pool Authority
SharedCPU=18 EntPoolCapacity= 0.75 No Shared CPU Pool Authority

```

#2

## Node 2 - NMON Outputs

```

PVS - orapha1 (ROHA VM1)
-----
NODE orapha-1.pvs.roha.com
-----
Current state: ST_INIT
-----
NODE orapha-2.pvs.roha.com
-----
Current state: ST_INIT
orapha-1:/#

```

#3

## Node1 - Command Window

```

PVS - orapha2 (ROHA VM2)
-----
orapha-2:/#

```

#4

## Node2 - Command Window

```

PVS - orapha1 (ROHA VM1)
* Welcome to AIX Version 7.3!
*
* Please see the README file in /usr/lpp/bos for information pertinent to
* this release of the AIX Operating System.
*
*****
orapha-1:/# clear
orapha-1:/# tailha
<LAT>|2023-04-05T16:07:35|29896|EVENT COMPLETED: node_down_complete orapha-1 0|</LAT>

+ clycle clavailability.log
+ 1> /dev/null 2>& 1
+ cltime
+ DATE=2023-04-05T16:07:35.474303
+ echo '<EVENT:NODE:DOWN_COMPLETE:END>|2023-04-05T16:07:35.474303|INFO: node_down_complete|orapha-1|0'
+ 1>> /var/hacmp/availability/clavailability.log
clxkit.rc : Normal termination of clstrmgrES. Restart now.
0513-059 The clstrmgrES Subsystem has been started. Subsystem PID is 29163914.

```

#5

## Node1 - hacmp.out log tail

```

PVS - orapha2 (ROHA VM2)
*
* Please see the README file in /usr/lpp/bos for information pertinent to
* this release of the AIX Operating System.
*
*****
orapha-2:/# clear
orapha-2:/# clear
orapha-2:/# tailha
<LAT>|2023-04-05T16:07:27|29899|EVENT COMPLETED: node_down_complete orapha-2 0|</LAT>

+ clycle clavailability.log
+ 1> /dev/null 2>& 1
+ cltime
+ DATE=2023-04-05T16:07:27.582888
+ echo '<EVENT:NODE:DOWN_COMPLETE:END>|2023-04-05T16:07:27.582888|INFO: node_down_complete|orapha-2|0'
+ 1>> /var/hacmp/availability/clavailability.log
clxkit.rc : Normal termination of clstrmgrES. Restart now.
0513-059 The clstrmgrES Subsystem has been started. Subsystem PID is 28049712.

```

#6

## Node2 - hacmp.out log tail

```
Shared-CPU-Logical-Partition
Broadcast message from root@orapha-1-e7d43bdd-0000211c" 7 ...
Flags:LPARed DRable SMT Shared UnCapped Migratable Not-Donating AMSable.
Oracle database started.dbUsed 0.02 ( 1.4%) 0.1% of CPUs in System
- No Shared CPU Pool Util. Authority/Allow Perf. Info. Collect
CPU-Stats-----q=Quit---CI/Ocity----- ID-Memory-----7
max Phys in sys 24 Cap. Processor Min 0.25 SPLPAR Group:Pool 32772:0
Phys CPU in sys 24 Cap. Processor Max 2.00 Memory(MB) Min:Max 3072:49152
Virtual Online 2 Cap. Increment 0.01 Memory(MB) Online 8192
Logical Online 16 Cap. Unallocated 0.00 Memory Region LMB 256MB min
Physical pool 18 Cap. Entitled 1.25 Time-----Seconds
SMT threads/CPU 8 -MinReqVirtualCPU 0.05 Time Dispatch Wheel 0.0100
CPU-----Min-Max Weight----- MaxDispatch Latency 0.0076
Virtual 1 5 Weight Variable 128 Time Pool Idle 0.0000
Logical 1 40 Weight Unallocated 0 Time Total Dispatch 0.0178

Event=30 post add cpu SerialNo Old=C48381 Current=C48381 When=06:39:17

Shared_Pools MaxPoolCapacity=18.00 No Shared CPU Pool Authority
SharedCPU=18 EntPoolCapacity= 1.25 No Shared CPU Pool Authority
```

```
-----
NODE orapha-1.pvs.roha.com
-----
Current state: ST_STABLE
-----
NODE orapha-2.pvs.roha.com
-----
Current state: ST_STABLE
orapha-1:/#
```

```
-----
orapha-1 ONLINE
orapha-2 OFFLINE
:node_up_complete[277] (( 0 == 0 ))
:node_up_complete[279] [[ orapha-2 != orapha-1 ]]
:node_up_complete[281] grep -w In_progress_file /var/hacmp/cl_dr.state
:node_up_complete[281] 2> /dev/null
:node_up_complete[281] cut '-d' -f2
:node_up_complete[281] lpm_in_progress_file=""
:node_up_complete[282] ls '/var/hacmp/.lpm_in_progress/lpm_*'
:node_up_complete[282] 2> /dev/null
:node_up_complete[282] lpm_in_progress_prefix=""
:node_up_complete[283] [[ -n '' ]]
:node_up_complete[300] exit 0
Apr 7 2023 06:41:29 EVENT COMPLETED: node_up_complete orapha-2 0

<LAT>|2023-04-07T06:41:29|31523|EVENT COMPLETED: node_up_complete orapha-2 0|</LAT>

+ clcycle clavailability.log
+ 1> /dev/null 2>& 1
+ cltime
+ DATE=2023-04-07T06:41:29.893660
+ echo '<EVENT:NODE:UP_COMPLETE:END>|2023-04-07T06:41:29.893660|INFO: node_up_complete|orapha-2|0'
+ 1> /var/hacmp/availability/clavailability.log
```

```
Shared-CPU-Logical-Partition
Broadcast message from root@orapha-2 (tty) at 06:34:53 ...
Flags:LPARed DRable SMT Shared UnCapped Migratable Not-Donating AMSable.
Starting Event Manager (clevmgrdES) s01 ( 2.7%) 0.0%a-2 CPUs in System
- No Shared CPU Pool Util. Authority/Allow Perf. Info. Collect
CPU-Stats----- Capacity----- ID-Memory-----
max Phys in sd-Disk-Graph Processor Min 0.25 SPLPAR Group:Pool 32772:0 5
Phys CPU in sys 24 Cap. Processor Max 2.00 Memory(MB) Min:Max 3072:49152
Virtual Online 1 Cap. Increment 0.01 Memory(MB) Online 3072
Logical Online 8 Cap. Unallocated 0.00 Memory Region LMB 256MB min
Physical pool 18 Cap. Entitled 0.25 Time-----Seconds
SMT threads/CPU 8 -MinReqVirtualCPU 0.05 Time Dispatch Wheel 0.0100
CPU-----Min-Max Weight----- MaxDispatch Latency 0.0150
Virtual 1 5 Weight Variable 128 Time Pool Idle 0.0000
Logical 1 40 Weight Unallocated 0 Time Total Dispatch 0.0068

Event= 0 --- --- SerialNo Old=--- Current=C48CD1 When=---

Shared_Pools MaxPoolCapacity=18.00 No Shared CPU Pool Authority
SharedCPU=18 EntPoolCapacity= 0.75 No Shared CPU Pool Authority
```

```
orapha-2:/#
Broadcast message from root@orapha-2 (tty) at 06:34:53 ...
Starting Event Manager (clevmgrdES) subsystem on orapha-2

[]
```

```
get secondary state info for state 4
getPrimaryStateStr: using primary_table => primary_state_table

Cluster Name: Oracle_PHA_ROHA_cluster


Resource Group Name: RGI
Node Timers Group State Delayed
-----
orapha-1 ONLINE
orapha-2 OFFLINE
:node_up_complete[277] (( 0 == 0 ))
:node_up_complete[279] [[ orapha-2 != orapha-2 ]]
:node_up_complete[300] exit 0
Apr 7 2023 06:41:29 EVENT COMPLETED: node_up_complete orapha-2 0

<LAT>|2023-04-07T06:41:29|31523|EVENT COMPLETED: node_up_complete orapha-2 0|</LAT>

+ clcycle clavailability.log
+ 1> /dev/null 2>& 1
+ cltime
+ DATE=2023-04-07T06:41:29.897363
+ echo '<EVENT:NODE:UP_COMPLETE:END>|2023-04-07T06:41:29.897363|INFO: node_up_complete|orapha-2|0'
+ 1> /var/hacmp/availability/clavailability.log
```


# orapha-1

### Physical CPU vs Entitlement - orapha-1 3/31/2023



# orapha-2


### Physical CPU vs Entitlement - orapha-2 3/31/2023



# orapha-1

### 4K - 64K Page Dynamics - orapha-1 3/31/2023


4KB\_used MB 4KB\_free MB 64KB\_used MB 64KB\_free MB



# orapha-2

### 4K - 64K Page Dynamics - orapha-2 3/31/2023

4KB\_used MB 4KB\_free MB 64KB\_used MB 64KB\_free MB





/var/hacmp/log/hacmp.out

ROHALOG:30999000:(10.957)	OPTIMAL APPS	Use Desired	Memory (GB)	CPU(s)	PU(s)/VP(s)
ROHALOG:30999000:(10.962)	oradb	0	5.0000	1	0.0000/ 0
ROHALOG:30999000:(10.971)	test_app2	0	4.0000	0	0.0000/ 0
ROHALOG:30999000:(10.980)	Total	0	9.0000	0	1.0000/ 0
ROHALOG:30999000:(10.984)					
ROHALOG:30999000:(10.989)					

```

[ROHALOG:33030548:(22.193)] ===== Compute ROHA [on Cloud] CPU(s) =====
[ROHALOG:33030548:(22.197)] == Status Reminder ==
[ROHALOG:33030548:(22.203)] Current resources      : 0.2500
[ROHALOG:33030548:(22.207)] == Raw computation from AC settings ==
[ROHALOG:33030548:(22.212)] LPAR profile minimum : 0.2500
[ROHALOG:33030548:(22.216)] APPs running         : 0.0000
[ROHALOG:33030548:(22.220)] APPs to start (optimal) : 1.0000
[ROHALOG:33030548:(22.225)] Total raw            : 1.2500
[ROHALOG:33030548:(22.229)] Delta acquisition raw : 1.0000
[ROHALOG:33030548:(22.234)] == Maximum limits (fit into if exceeding and always_start_rg=1, otherwise fail).
[ROHALOG:33030548:(22.238)] LPAR profile maximum : 2.0000
[ROHALOG:33030548:(22.241)] == Minimum limits (adjust if under) ==
[ROHALOG:33030548:(22.244)] Delta minimum        : 0.00
[ROHALOG:33030548:(22.247)] == Final computation, limits considered ==
[ROHALOG:33030548:(22.250)] Total                 : 1.2500
[ROHALOG:33030548:(22.252)] Delta acquisition     : 1.0000
[ROHALOG:33030548:(22.255)] ===== End =====

```

```

[ROHALOG:33030548:(22.125)] ===== Compute ROHA [on Cloud] Memory =====
[ROHALOG:33030548:(22.132)] == Status Reminder ==
[ROHALOG:33030548:(22.135)] Current resources      : 3.0000 GB
[ROHALOG:33030548:(22.142)] == Raw computation from AC settings ==
[ROHALOG:33030548:(22.145)] LPAR profile minimum : 3.0000 GB
[ROHALOG:33030548:(22.151)] APPs running         : 0.0000 GB
[ROHALOG:33030548:(22.155)] APPs to start (optimal) : 9.0000 GB
[ROHALOG:33030548:(22.161)] Total raw            : 12.0000 GB
[ROHALOG:33030548:(22.165)] Delta acquisition raw : 9.0000 GB
[ROHALOG:33030548:(22.167)] == Maximum limits (fit into if exceeding and always_start_rg=1, otherwise fail).
[ROHALOG:33030548:(22.170)] LPAR profile maximum : 48.0000 GB
[ROHALOG:33030548:(22.173)] == Minimum limits (adjust if under) ==
[ROHALOG:33030548:(22.176)] Delta minimum        : 0.00 GB
[ROHALOG:33030548:(22.179)] == Final computation, limits considered ==
[ROHALOG:33030548:(22.182)] Total                 : 12.0000 GB
[ROHALOG:33030548:(22.185)] Delta acquisition     : 9.0000 GB
[ROHALOG:33030548:(22.187)] ===== End =====

```

`clicloudroha -o acquire -m 12.0000 -p 1.2500`

EVENT START: start\_server oradb test\_app2

**AIX .25 CPU Instance 3GB Mem**

## Example Acquiring (2) Application controllers

**ROHA Values**

- oradb + 1 CPU + 5GB Mem**
- test\_app2 + 4GB Mem**

**Result: 12GB Memory & 1.25 CPU**

**ROHA Command PowerHA invokes to PowerVS Service broker**

```
orapha-1:/# clmgr view report roha
Cluster: 'Oracle_PHA_ROHA_cluster' of No Site Cluster type
```

```
Cluster tunables
Dynamic LPAR
Always Start Resource Groups: '0'
Adjust Shared Processor Pool size if required: '0'
Force synchronous release of DLPAR resources: '0'
```

```
Cloud Site: Oracle_PHA_ROHA_cluster
```

```
Cloud Node: orapha-1
```

```
Cloud Type: IBM
Cloud Instance ID: e7d43bdd-1acb-4300-807c-0affdd05ab90
Cloud Operations Timeout: 5 minutes
Cloud LPAR Profile:
```

Status: ACTIVE	Health: OK	HostID: 103			
Memory (in GB):	Minimum=3	Desired=NA	Current=3	Maximum=48	Available=803
Processing units:	Minimum=0.25	Desired=NA	Current=0.25	Maximum=2	Available=12.75
Virtual processors:	Minimum=1	Desired=NA	Current=1	Maximum=5	

```
Processing mode: shared
Shared processor pool: NA
```

```
ROHA provisioning for resource groups
```

```
Resource group: 'RG1' Application controller: 'oradb'
Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'
```

```
Cloud Node: orapha-2
```

```
Cloud Type: IBM
Cloud Instance ID: c857277c-73cb-4c58-ba5f-011399d34814
Cloud Operations Timeout: 5 minutes
Cloud LPAR Profile:
```

Status: ACTIVE	Health: OK	HostID: 100			
Memory (in GB):	Minimum=3	Desired=NA	Current=3	Maximum=48	Available=799
Processing units:	Minimum=0.25	Desired=NA	Current=0.25	Maximum=2	Available=12.25
Virtual processors:	Minimum=1	Desired=NA	Current=1	Maximum=5	

```
Processing mode: shared
Shared processor pool: NA
```

```
ROHA provisioning for resource groups
```

```
Resource group: 'RG1' Application controller: 'oradb'
Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'
```

Instance #1

Instance #2

CPU & Memory  
ROHA Resources  
Requested



```
orapha-1:/# clmgr view report roha
Cluster: 'Oracle_PHA_ROHA_cluster' of No Site Cluster type
Cluster tunables
  Dynamic LPAR
    Always Start Resource Groups: '0'
    Adjust Shared Processor Pool size if required: '0'
    Force synchronous release of DLPAR resources: '0'
Cloud Site: Oracle_PHA_ROHA_cluster
```

## Instance #1

Cloud Node: orapha-1  
 Cloud Type: IBM  
 Cloud Instance ID: e7d43bdd-1acb-4300-807c-0affdd05ab90  
 Cloud Operations Timeout: 5 minutes  
 Cloud LPAR Profile:

Status: ACTIVE	Health: OK	HostID: 103	Desired=NA	Current=3	Maximum=48	Available=803
Memory (in GB):	Minimum=3	Desired=NA	Current=NA	Current=0.25	Maximum=2	Available=12.75
Processing units:	Minimum=0.25	Desired=NA	Current=NA	Current=1	Maximum=5	
Virtual processors:	Minimum=1	Desired=NA	Current=NA			
Processing mode:	shared					
Shared processor pool:	NA					

**Currently neither Instance is expanded**

**Available Resources on Backend Servers**

```
ROHA provisioning for resource groups
  Resource group: 'RG1' Application controller: 'oradb'
    Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
  Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'
```

## Instance #2

Cloud Node: orapha-2  
 Cloud Type: IBM  
 Cloud Instance ID: c857277c-73cb-4c58-ba5f-011399d34814  
 Cloud Operations Timeout: 5 minutes  
 Cloud LPAR Profile:

Status: ACTIVE	Health: OK	HostID: 100	Desired=NA	Current=3	Maximum=48	Available=799
Memory (in GB):	Minimum=3	Desired=NA	Current=NA	Current=0.25	Maximum=2	Available=12.25
Processing units:	Minimum=0.25	Desired=NA	Current=NA	Current=1	Maximum=5	
Virtual processors:	Minimum=1	Desired=NA	Current=NA			
Processing mode:	shared					
Shared processor pool:	NA					

```
ROHA provisioning for resource groups
  Resource group: 'RG1' Application controller: 'oradb'
    Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
  Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'
```

```

orapha-1:/# clmgr view report roha
Cluster: 'Oracle_PHA_ROHA_cluster' of No Site Cluster type
Cluster tunables
  Dynamic LPAR
    Always Start Resource Groups: '0'
    Adjust Shared Processor Pool size if required: '0'
    Force synchronous release of DLPAR resources: '0'
Cloud Site: Oracle_PHA_ROHA_cluster

```

## Instance #1

```

Cloud Node: orapha-1
Cloud Type: IBM
Cloud Instance ID: e7d43bdd-1acb-4300-807c-0affdd05ab90
Cloud Operations Timeout: 5 minutes
Cloud LPAR Profile:
  Status: ACTIVE      Health: OK      HostID: 103
  Memory (in GB):    Minimum=3      Desired=NA
  Processing units:  Minimum=0.25  Desired=NA
  Virtual processors: Minimum=1        Desired=NA
  Processing mode: shared
  Shared processor pool: NA

```

**Oracle DB  
currently  
hosted on this  
Instance**

Current=8  
Current=1.25  
Current=2

Maximum=48  
Maximum=2  
Maximum=5

Available=798  
Available=11.75

**Backend  
Server  
Resources  
Reduced**

```

ROHA provisioning for 'ONLINE' resource groups
  Resource group 'RG1' Application controller: 'oradb'
    Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
  Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'

```

## Instance #2

```

Cloud Node: orapha-2
Cloud Type: IBM
Cloud Instance ID: c857277c-73cb-4c58-ba5f-011399d34814
Cloud Operations Timeout: 5 minutes
Cloud LPAR Profile:
  Status: ACTIVE      Health: OK      HostID: 100
  Memory (in GB):    Minimum=3      Desired=NA
  Processing units:  Minimum=0.25  Desired=NA
  Virtual processors: Minimum=1        Desired=NA
  Processing mode: shared
  Shared processor pool: NA

```

```

ROHA provisioning for 'OFFLINE' resource groups
  Resource group 'RG1' Application controller: 'oradb'
    Memory='5.0' Processors='NA' Processing units='1.00' Virtual Processors='NA'
  Total: Use desired number='0' Memory='5' Processors='NA' Processing units='1' Virtual Processors='NA'

```

---

**Observation: Resource Allocation**

---

```
orapha-1:/usr/local/hascripts# cat jfs2db_start.ksh
```

```
#!/bin/ksh
```

```
ORA_HOME=/u01/app/oracle/product/19c
```

```
ORA_OWNER=oracle
```

```
wall 'STARTING application oradb'
```

```
# Start the Oracle databases:
```

```
# This also starts the listener!
```

```
# The following command assumes that the oracle login
```

```
# will not prompt the user for any values
```

```
# Will start any DB listed in /etc/oratab with "Y" as last entry in line
```

```
su - $ORA_OWNER -c "$ORA_HOME/bin/dbstart $ORA_HOME"
```

```
wall 'Oracle database started'
```

```
exit 0
```

## ORACLE Database Start Logic

```
orapha-1:/usr/local/hascripts# cat jfs2db_stop.ksh
```

```
#!/bin/ksh
```

```
ORA_HOME=/u01/app/oracle/product/19c
```

```
ORA_OWNER=oracle
```

```
wall 'STOPPING application oradb'
```

```
# Stop the Oracle databases:
```

```
# This also stops the listener!
```

```
# The following command assumes that the oracle login
```

```
# will not prompt the user for any values
```

```
# /etc/oratab line with "Y" for DB instances to be started!
```

```
su - $ORA_OWNER -c "$ORA_HOME/bin/dbshut $ORA_HOME"
```

```
wall 'Oracle DB stopped'
```

```
exit 0
```

## ORACLE Database Stop Logic

## Resource Optimized High Availability

Move cursor to desired item and press Enter.

HMC Configuration

Hardware Resource Provisioning for Application Controller

**Change/Show Default Cluster Tunables**

NovaLink Configuration

Cloud Configuration

## Change/Show Default Cluster Tunables

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

Dynamic LPAR

**Always Start Resource Groups**

Adjust Shared Processor Pool size if required

Force synchronous release of DLPAR resources

Enterprise Pool

Resource Allocation order

On/Off CoD

I agree to use On/Off CoD and be billed for extra costs No

Number of activating days for On/Off CoD requests [30]

[Entry Fields]

No

No

No

Free Pool before Enterprise pool

### CONTEXTUAL HELP

Press Enter or Cancel to return to the application.

**Enter 'Yes' to have PowerHA SystemMirror start Resource Groups even if resources are insufficient. It may occur when the total requested resources exceeds the LPAR profile maximum or the combined available resources. Thus the best-can-do allocation is performed.**

Enter 'No' to prevent starting Resources Groups with insufficient resources. Resource Groups may move to error state if resources are insufficient.

Default is 'Yes'.

```
orapha-1:/usr/local/hascripts# cat jfs2db_start.ksh
```

```
#!/bin/ksh  
# set -x  
ORA_HOME=/u01/app/oracle/product/19c  
ORA_OWNER=oracle
```

```
echo "`date`: Entered jfs2db_start.ksh ..." > /tmp/DBstartLog.txt
```

LOG File

```
minmem=`/usr/es/sbin/cluster/utilities/clcloudroha -o query -a minmem | sed '/minmem/d`  
meminc=`clmgr q roha oradb|awk -F\" '/OPTIMAL_MEM/ {print $2;}`  
let reqmem=${minmem}+${meminc}  
currmem=`lparstat -i |awk '/^Online Memory/ { print $4/1024;}`
```

```
let count=60
```

60 second wait

```
echo "`date`: Starting wait for memory ..." >> /tmp/DBstartLog.txt
```

```
while [ $count -gt 0 -a $currmem -lt $reqmem ]  
do  
    sleep 1  
    currmem=`lparstat -i |awk '/^Online Memory/ { print $4/1024;}`  
    echo "`date`: $count seconds left : Memory not active yet - need: $reqmem and have: $currmem." >>  
    /tmp/DBstartLog.txt  
    let count=count-1  
done
```

Memory check every second up to 60s

```
if [ $count -eq 0 ]  
then  
    echo "`date`: Did not get sufficient memory within 60 seconds - aborting DB start!" >> /tmp/DBstartLog.txt  
    exit 1  
fi
```

Abort after 60s if memory is not there

```
wall 'STARTING application oradb'  
# This also starts the listener!  
# Will start any DB listed in /etc/oratab with "Y" as last entry in line
```

```
su - $ORA_OWNER -c "$ORA_HOME/bin/dbstart $ORA_HOME"  
echo "`date`: DBs started if marked in /etc/oratab with 'Y.'" >> /tmp/DBstartLog.txt  
wall 'Oracle database started'  
exit 0
```

```
orapha-1:/tmp# cat DBstartLog.txt  
Wed Apr 5 15:47:48 CDT 2023: Entered jfs2db_start.ksh ...  
Wed Apr 5 15:47:55 CDT 2023: Starting wait for memory ...  
Wed Apr 5 15:47:56 CDT 2023: 60 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:47:57 CDT 2023: 59 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:47:58 CDT 2023: 58 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:47:59 CDT 2023: 57 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:00 CDT 2023: 56 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:01 CDT 2023: 55 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:02 CDT 2023: 54 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:03 CDT 2023: 53 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:04 CDT 2023: 52 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:05 CDT 2023: 51 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:06 CDT 2023: 50 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:07 CDT 2023: 49 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:08 CDT 2023: 48 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:09 CDT 2023: 47 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:10 CDT 2023: 46 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:11 CDT 2023: 45 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:12 CDT 2023: 44 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:13 CDT 2023: 43 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:14 CDT 2023: 42 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:15 CDT 2023: 41 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:16 CDT 2023: 40 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:17 CDT 2023: 39 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:18 CDT 2023: 38 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:19 CDT 2023: 37 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:20 CDT 2023: 36 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:21 CDT 2023: 35 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:22 CDT 2023: 34 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:23 CDT 2023: 33 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:24 CDT 2023: 32 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:25 CDT 2023: 31 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:26 CDT 2023: 30 seconds left : Memory not active yet - need: 8 and have: 3.  
Wed Apr 5 15:48:58 CDT 2023 : DBs started if marked in /etc/oratab with 'Y'.
```

## ORACLE Database Start Logic with Validation of Memory expansion

PVS - orapha2 (ROHA VM2)

☺

PVS - orapha2 (ROHA VM2)

☺

topas\_rmon N=NFS Host=orapha-2 Refresh=2 secs 07:34.12

**Shared-CPU-Logical-Partition**

Partition: Number=4 "orapha-2-c857277c-0000211f"

Flags: LPARed DRable SMT Shared UnCapped Migratable Not-Donating AMSable.

Summary: Entitled= 0.25 Used 0.03 ( 12.8%) 0.1% of CPUs in System

- No Shared CPU Pool Util. Authority/Allow Perf. Info. Collect

CPU-Stats-----		Capacity-----		ID-Memory-----	
max Phys in sys	24	Cap. Processor Min	0.25	SPLPAR Group:Pool	32772:0
Phys CPU in sys	24	Cap. Processor Max	2.00	Memory(MB) Min:Max	3072:49152
Virtual Online	1	Cap. Increment	0.01	Memory(MB) Online	3072
Logical Online	8	Cap. Unallocated	0.00	Memory Region LMB	256MB min
Physical pool	18	Cap. Entitled	0.25	Time-----Seconds	
SMT threads/CPU	8	-MinReqVirtualCPU	0.05	Time Dispatch Wheel	0.0100
CPU-----Min-Max		Weight-----		MaxDispatch Latency	0.0150
Virtual	1 5	Weight Variable	128	Time Pool Idle	0.0000
Logical	1 40	Weight Unallocated	0	Time Total Dispatch	0.0319

Event= 0 --- SerialNo Old=--- Current=C4BCD1 When=---

Shared\_Pools MaxPoolCapacity=18.00 No Shared CPU Pool Authority

SharedCPU=18 EntPoolCapacity= 0.75 No Shared CPU Pool Authority

orapha-2:/# r tail  
tail -f /tmp/DBstartLog.txt

---

**Observation: I/O Throttling**

---



Attached volumes			
Volume affinity setting ⓘ Same storage pool			
Search			
<input type="checkbox"/>	Name	Size	Tier
<input type="checkbox"/>	oradata-5	5 GB	Tier 1
<input type="checkbox"/>	oradata-4	5 GB	Tier 1
<input type="checkbox"/>	oradata-3	5 GB	Tier 1
<input type="checkbox"/>	oradata-2	5 GB	Tier 1
<input type="checkbox"/>	oradata-1	5 GB	Tier 1
<b>ORACLE Data VG</b>			
<input type="checkbox"/>	oraredo-3	5 GB	Tier 1
<input type="checkbox"/>	oraredo-2	5 GB	Tier 1
<input type="checkbox"/>	oraredo-1	5 GB	Tier 1
<b>ORACLE Redo VG</b>			
<input type="checkbox"/>	orabin	75 GB	Tier 1
<b>ORACLE Binary VG</b>			
<input type="checkbox"/>	stage	20 GB	Tier 1
<input type="checkbox"/>	CAA_HB	hdisk1 in each node	1 GB Tier 1
<input type="checkbox"/>	orapha-1-e7d43bdd-0000211c-boot-0	25 GB	Tier 1

Cluster Operation	Total Time
<b>1) Start of Custer Services</b>	<b>10 min (stabilized cluster)</b>
• Oracle DB available	8 minutes
<b>2) RG move operation</b>	<b>6 minutes (stabilized cluster)</b>
• CPU & Memory release on node 1	within 3 minutes
• CPU & Memory allocation on node 2	within 5 minutes

Original LUN sizes



Attached volumes

Volume affinity setting ⓘ  
Same storage pool

Search

<input type="checkbox"/>	Name	Size	Tier	
<input type="checkbox"/>	oradata-5	5 GB	Tier 1	
<input type="checkbox"/>	oradata-4	5 GB	Tier 1	
<input type="checkbox"/>	oradata-3	5 GB	Tier 1	
<input type="checkbox"/>	oradata-2	5 GB	Tier 1	
<input type="checkbox"/>	oradata-1	75GB	Tier 1	
<input type="checkbox"/>	oraredo-3	5 GB	Tier 1	
<input type="checkbox"/>	oraredo-2	5 GB	Tier 1	
<input type="checkbox"/>	oraredo-1	75GB	Tier 1	
<input type="checkbox"/>	orabin	75 GB	Tier 1	
<input type="checkbox"/>	stage	20 GB	Tier 1	
<input type="checkbox"/>	CAA_HB	hdisk1 in each node	1 GB	Tier 1
<input type="checkbox"/>	orapha-1-e7d43bdd-0000211c-boot-0	25 GB	Tier 1	

**ORACLE Data VG**

**ORACLE Redo VG**

**ORACLE Binary VG**

## Original LUN sizes Times:

Cluster Operation	Total Time
<b>1) Start of Custer Services</b>	<b>10 min (stabilized cluster)</b>
• Oracle DB available	8 minutes
<b>2) RG move operation</b>	<b>6 minutes (stabilized cluster)</b>
• CPU & Memory release on node 1	within 3 minutes
• CPU & Memory allocation on node 2	within 5 minutes

## After LUN Resize (Increased 1<sup>st</sup> disk in each volume group):

Cluster Operation	Total Time
<b>1) Start of Custer Services</b>	<b>5 min (stabilized cluster)</b>
• Oracle DB available	4 minutes
<b>2) RG move operation</b>	<b>~3 to 4 minutes (stabilized cluster)</b>
• CPU & Memory release on node 1	within 1 minutes
• CPU & Memory allocation on node 2	within 2 minutes

LESSONS

LEARNED



✓ **Use PowerHA SystemMirror V727 SP1**

✓ **Default PowerVS Instance Minimum sizes [ .25 CPU & 2GB Memory ]**

→ support ticket to change | Instance needs to be offline. (Ideally PHA deploy should have gone with 4GB of mem)

✓ **ROHA Tunable to “Always Start Resource Groups” needs update** (doesn't guarantee resources)

(Fallover works, but Implemented script work around)  
5 min default timeout (operations occur within a min+)

Change/Show Default Cluster Tunables

Type or select values in entry fields.  
Press Enter AFTER making all desired changes.

[Entry Fields]

Dynamic LPAR	
<b>Always Start Resource Groups</b>	<b>No</b>
Adjust Shared Processor Pool size if required	No
Force synchronous release of DLPAR resources	No

✓ **PowerVS IOPs Storage Limitations** (Slower Start / Takeover)

✓ **Expectations with multiple Application Controllers using ROHA**

✓ **Implication of using Large Pages (16MB) for Oracle SGA**

# KEY TAKEAWAYS



**New Easy way for clustered AIX instances on IBM Power Virtual Server infrastructure to integrate Automatic Resource Adjustments of CPU & Memory**



**Feature enables OPEX reduction & minimizes resource utilization**  
( Note that it is NOT limited to **ORACLE** )



**Consider resource availability in a Multi-tenant environment**  
( Option to reserve capacity with Shared Processor Pools )

# IBM Advanced Technology Group

Our mission is to provide **technical pre-sales** assistance to IBM and Business Partner hardware sellers within North America for **IBM Power and IBM Storage** opportunities

## What ATG offers:

- Competitive engagements (Server and storage)
- Proof of Concepts, Proof of Technologies, performance benchmarks
- Complex solution architecture and design
- Articulate IBM value propositions
- Establish credibility through best practices guidance
- Subject Matter Expertise to improve client experience on IBM platforms
- Customer presentations, Lunch & Learns, canned & live demonstrations
- System Health Check, best practices review
- Tools, collateral & cookbooks

## ATG Technical Leads:

### ▪ **ISV on Power:**

- Oracle: Ralf Schmidt-Dannert ([dannert@us.ibm.com](mailto:dannert@us.ibm.com))
- HA & DR: Mike Herrera ([mherrera@us.ibm.com](mailto:mherrera@us.ibm.com))
- SAP: Louis Lamprinakos ([lnlampri@us.ibm.com](mailto:lnlampri@us.ibm.com))
- SAS: Harry Seifert ([seifert@us.ibm.com](mailto:seifert@us.ibm.com))

### ▪ **IBM Tape:**

- Roy Peek ([rhpeek@us.ibm.com](mailto:rhpeek@us.ibm.com))

### ▪ **IBM Storage:**

- Carl Brown ([carltb@us.ibm.com](mailto:carltb@us.ibm.com))

## ATG team engaged via:

- IBM seller
- IBM Business partner