Technical Notes

IBM Oracle International Competency Center (ICC)

August 09, 2017 ***email address: ibmoracle@us.ibm.com***

PowerHA SystemMirror 7 and Oracle RAC 12c Compatibility and Requirements

Overview

The note provides a snapshot of the certified versions of PowerHA SystemMirror, AIX, and Oracle RAC 12c. For the most up-to-date certified combination, visit My Oracle Support (MOS) certification tab.

Additionally, requirements are listed which need to be fulfilled to ensure successful installation of the PowerHA, AIX with Oracle RAC 12c.

The Oracle MOS (My Oracle Support) Doc 2293047.1 contains the same information as in this note.

Certification matrix

The following table shows the certified Oracle RAC releases, PowerHA versions and AIX versions.

RAC Database release	PowerHA version	AIX version
12.1.0.2.0	7.2	7.1 & 7.2
12.1.0.2.0	7.1	7.1
12.1.0.1.0	7.1	7.1
11.2.0.4	7.1	7.1

PowerHA, VIOS, and AIX version requirements

The table below lists the minimum version requirement:

AIX min. version	PowerHA min. version	VIOS min. version
7.1 TL4 SP3	7.1.3 SP4	2.2.3 SP3
7.2 TL1 SP1	7.2 SP1	

a. It is required to install the IFIX Bundles according to the specific PowerHA version and AIX version/release.

   IFIX Bundles information and downloads are available from PowerHA SystemMirror IFIX Bundles Information.

b. If the PowerHA cluster consists of more than two nodes, install the fix for APAR IV93378 - *PowerHA 7 cluster startup (also VIOS SSP clustering) may be delayed due to Group Services domain establishment issues*.

c. When using shared-processor pools, it is required to set the AIX tunable parameter vpm_xvcpus to 2 for Oracle RAC to function properly.
   # schedo –p -o vpm_xvcpus=2

Required PowerHA communication configuration

To support high availability and safeguard cross kill exposure (discussed in the next section), the required PowerHA communications configuration includes at least three IP-based networks and two repository disks (primary and backup).

In a typical Oracle RAC environment, there is an Oracle public Ethernet network and at least two Ethernet networks for Oracle private (Interconnect) networks. By default, PowerHA automatically makes use of all IP-based networks for its communications.

When creating PowerHA cluster as part of the initial configuration, a disk has to be configured as the repository disk which also serves as a communication device. The only extra step is to manually add a backup repository disk.

Cross kill exposure

When failures in the Oracle Interconnect paths occur such that partitions are formed and all partitions have the same number of nodes, a "split brain" scenario is created. The "split brain" scenario is characterized by nodes within a partition can communicate with each other over the Interconnect

but not across the partitions. Oracle resolves the "split brain" scenario by keeping the partition which has a lower node number up and evict the nodes in the other partition. If I/Os from the Oracle processes cannot be fenced off on these nodes, instead of reboot less node evictions, they are forced to reboot.

PowerHA "split brain" occurs when communications to all the repository disks fail and the IP-based networks have partial communications whereby the nodes are partitioned with the same number of nodes in each partition and there is no IP-based communications between any two partitions. PowerHA resolves the "split brain" scenario by keeping the nodes in a partition up and reboot the nodes in other partitions.

Cross kill is a scenario where both Oracle and PowerHA detect a "split brain" and each tries to kill a different partition because PowerHA and Oracle RAC use different algorithm to determine which partition to kill. When PowerHA detects a "split brain" scenario, Oracle detects it too because PowerHA communication paths are a superset of Oracle's Interconnect paths. However, when PowerHA is configured as stated in the communication requirements, failure in all PowerHA's disk-based communications and all IP-based communications experience partial failure that have equal number of nodes in the partitions simultaneously is rather rare.

For PowerHA 7.1 and later versions, it is required to configure at least three IP-based networks and two repository disks (the primary and the backup) for heart beating to minimize the cross kill exposure.

Typically, an Oracle RAC environment would have at a minimum of an Oracle public Ethernet network and at least two Oracle private Ethernet networks (Interconnects). By default, PowerHA will automatically use all IP-based networks for communications and thus include at least the three IP-based networks for heart beating. During the configuration step to create a PowerHA cluster, it will prompt for the path of the repository disk. To achieve the required five redundant heartbeat paths, the only extra manual step is to add a backup repository disk.

Oracle bug

Bug 20601073: ADVM Resource Fails to Start After Node Eviction (Doc ID 2135375.1)
Contact My Oracle Support (MOS) to obtain the patch.

Configuring Oracle voting disks

The Oracle voting disks were placed in an ASM diskgroup configured from raw disk devices in the Oracle RAC 12c R2 certification tests. While they can be configured differently, they have not been tested. It is recommended voting disks be placed in an ASM diskgroup backed by raw disk devices.

Ethernet switch settings for VIOS SEA failover

When SEA failover is configured on the VIOS, the Ethernet switch should follow the settings below:

- Spanning tree protocol = enable (set spantree enable port)
- Start port fast = enable (set spantree portfast mod/port enable)
- Delay forwarding = 15 secs (set spantree fwddelay 15 vlans)
- Trunking = off (set trunk mod/port off)
- Etherchannel = off (set port channel mod/port off)

The command, in parenthesis, may be different on your network switch.

Additional information may be found at these websites:

1. My Oracle Support (for the Oracle Certification Matrix)[(1)]: https://support.oracle.com/CSP/ui/flash.html

[(1)] A userid and password are required to access that website. This is available for Oracle customers and Oracle Business Partners.

2. Oracle's website for downloads of the database: http://www.oracle.com/technetwork/database/enterprise-edition/downloads/index.html
3. Oracle's website for Database information: http://www.oracle.com/us/products/database/overview/index.html
4. IBM Power Systems website: http://www-03.ibm.com/systems/power/
5. IBM PowerHA website: https://www-03.ibm.com/systems/power/software/availability/

For further information on this Technical Note , please e-mail the IBM Oracle International Competency Center at ibmoracle@us.ibm.com with any questions you may have on this topic.