



Oracle Networking and High Availability Options (with Linux on System z) & Red Hat/SUSE Oracle Update

Speaker Names: David Simpson & Kathryn Arrell

Speakers Company: IBM

Date of Presentation: **Wednesday, August 14, 2013: (09:30 AM)**

Hynes, Room 308

Session Number: **13523**

Twitter -> @IBMANDOracle



Agenda – Part I

- **Why use Live Guest Relocation with Oracle Database**
- **Oracle Support Policy**
- **Challenges**
- **Factors**
- **Scenarios**
- **Recommendation**

Oracle with LGR on System z: Business Value



■ Availability

- Improve applications availability (load balancing)
- Planned software or hardware maintenance - Offload a physical server to allow replacement, firmware upgrade or disruptive physical upgrade

■ Performance

- Move guest to free resources for other partitions
- Move guest to a more powerful physical server

■ Flexibility

- Move guest to follow business needs
- Move test guests to production servers for real tests
- Free a physical server for planned maintenance
- Migration between z/VM maintenance levels

■ Energy efficiency

- Move partitions from underused physical servers to other servers to increase their resource usage and their energy efficiency
- Shutdown underused physical servers

■ Not designed for unplanned outages

- Live guest relocations are initiated by a manual VMRELOCATE command.
The command is not automatically issued for any guest.
- **LGR is not a high availability solution, and it is not a disaster recovery solution.**
- Think of LGR as a **continuous availability solution** for the Linux guest.

Complete your sessions evaluation online at SHARE.org/BostonEval



Oracle Support for LGR

- **Oracle has a process to certify running an active Oracle DB in virtualized environments to ensure customers will be able to get support when running in that virtualized environment**
- **IBM has to complete a robust set of tests to prove Oracle databases can be running when LGR is executed. (Includes tests on SLES 11 and RH6 and single instance and RAC)**
- **Similar to support for Live Partition Mobility on AIX and Live Guest Migration on Oracle VM (Intel)**

LGR mechanism

Step 0

VMRELOCATE MOVE

Step 1

Eligibility Checks

Step 2

Create skeleton on destination

Step 3

Moves guest memory
Over ISCF channel(s)

Step 4

Quiesce guest

Step 5

Move guest state, I/O device...

Step 6

Final memory move pass

Step 7

Guest resumed on destination

**Oracle
Support
Challenge**

**Quiesce time
must be less
than the RAC
time out or
node will
reboot**

**Environment
must be able
to run with a
stress load
even for single
instance**

**Quiesce
Time**

Factors affecting relocation



Linux guest memory size

Linux guest memory update rate

Channel-to-channel adapter (CTCA) setup
Number and speed

VMRELOCATE options (SYNCH, ASYNCH, IMMEDIATE,
MAXTOTAL, MAXQUIESCE)
(Real Application Cluster time out is 30 seconds)

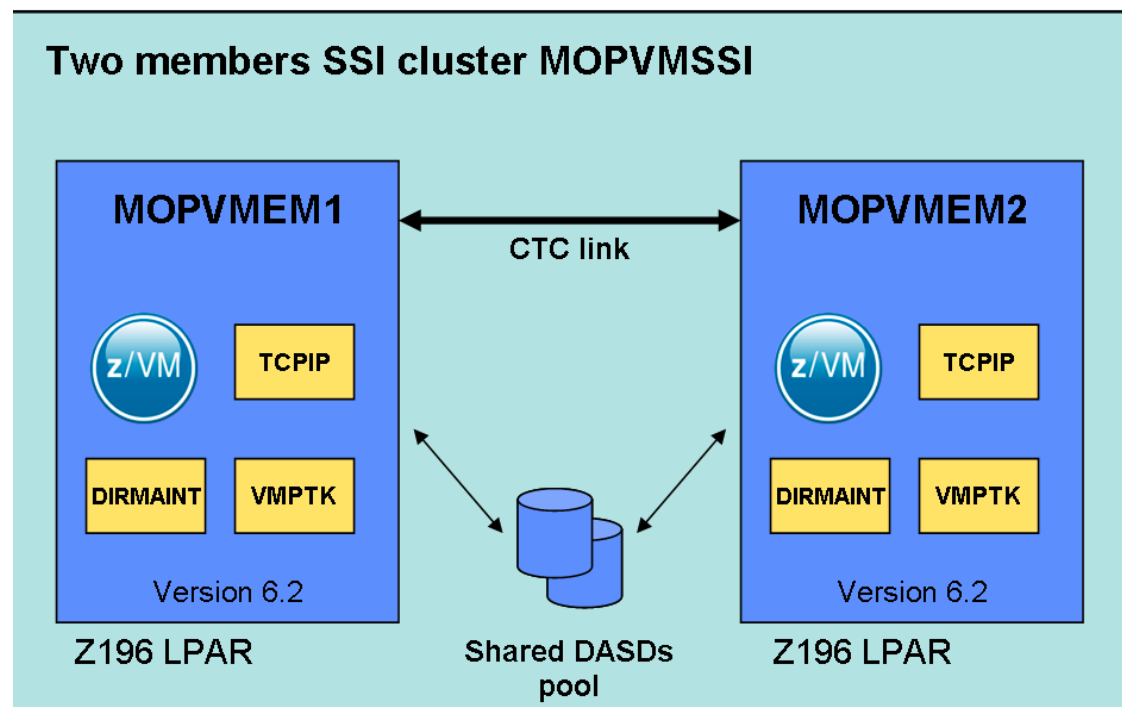
Other workloads on target system

Safe Guest Relocation

- **Eligibility checks** done multiple times throughout the relocation process.
- Checks more than just eligibility to move the virtual machine, but also checks whether it is **safe** to move.
 - Does the virtual machine really have access to all the same resources and functions?
 - Will moving the virtual machine over commit resources to the point of jeopardizing other workload on the destination system?
- **Pacing logic** to minimize impact to other work in more memory constrained environments
- **Timing Logic** that aborts a migration if the quiesce period exceeds a predefined window.

Scenarios tested for Redbook in 2012

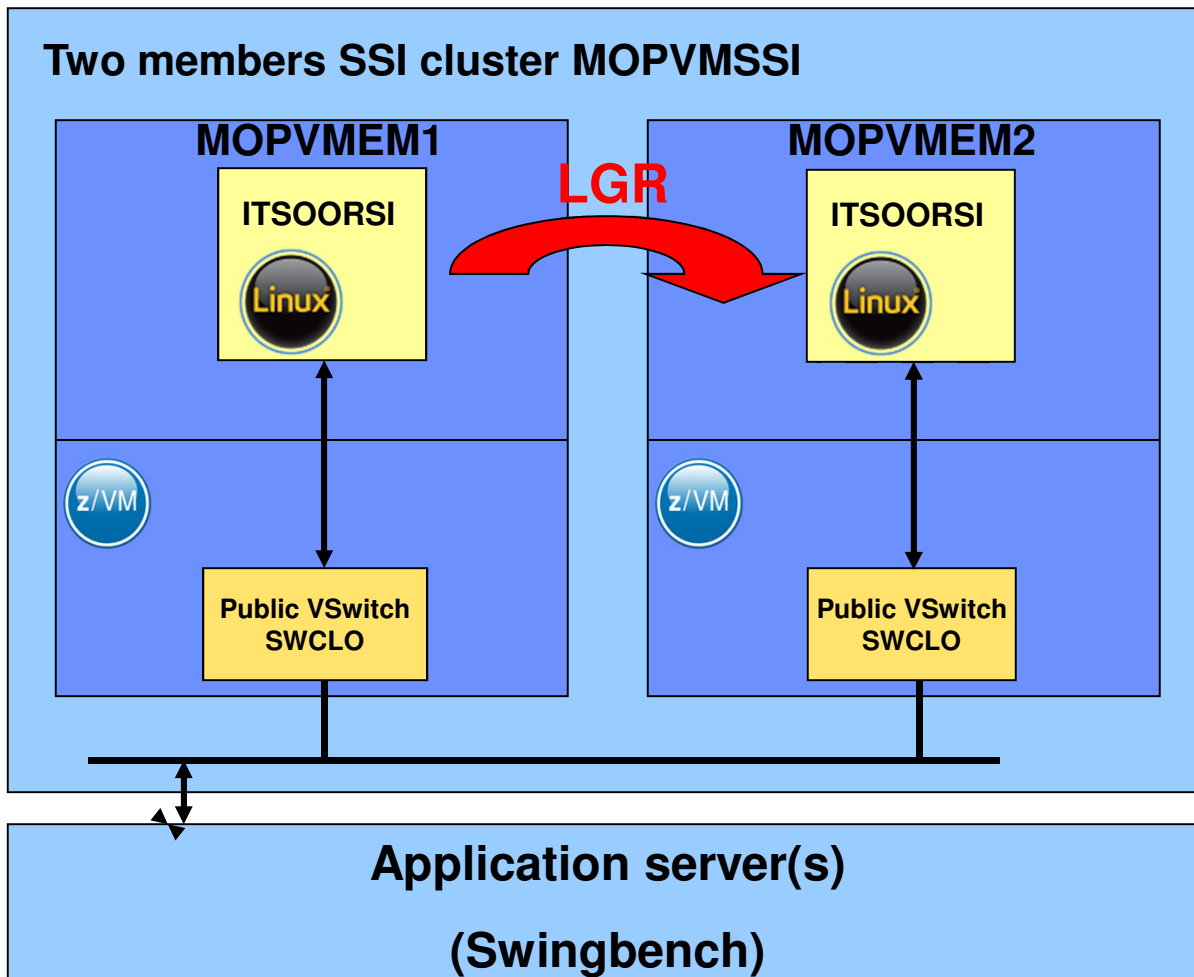
- Environment
 - z/VM 6.2 cluster (2 members)
 - Linux SLES11 SP1
 - Oracle 11.2.0.3
 - Oracle Single Instance
 - Oracle RAC
 - Swingbench



Oracle Live Guest Relocation Scenario

- **3 scenarios:**
 - Oracle Single Instance relocation using LGR
 - Oracle RAC Node relocation with stopping the Node before relocation (recommended while LGR is not certified with Oracle RAC)
 - Oracle RAC Node relocation using Live Guest Relocation on an active node

Scenario 1: Oracle Single Instance relocation



- Application Server (swingbench) simulates users and generates workload to the Oracle DB

- The Oracle DB is run as a Single Instance in a z/VM 6.2

- You need to:
 - perform a maintenance on the left z/VM server partition
 - or run additional workload on the z/VM LPAR
 - or migrate the production to a new z/VM LPAR.

Today the supported way is to bring down the database and move Linux guest

Oracle Single Instance relocation



- The application server (Swingbench) simulates users and generates workload to the Oracle DB.

The screenshot displays the SwingBench 2.3.0.422 interface for an Oracle Single Instance relocation. The window title is "SwingBench 2.3.0.422 : 'Order Entry (PLSQL)'".

Configuration Panel (Left):

- User Details:** Username: soe, Password: ..., Connect String: ITSO, Driver Type: Oracle10g Type II jdbc driver (oci).
- Load:** Number of Users: 200, Min. Delay Between Transactions (ms): 300, Max. Delay Between Transactions (ms): 1,000.
- Logon:** Logon Delay (milliseconds): 1,000, Logon Group: 1, Wait Till All Sessions Log On: false, Logoff Post Transaction: false.
- Benchmark Run Time:** 0:10 (hh:min).
- Recording:** Record Statistics After (hh:min): 0:0, Stop Recording After (hh:min): 0:0.

Transactions Table (Top Right):

Id	Class Name	Short Name	Load Ratio	Activate ?
Customer Registration	com.dom.benchmarking.swingbench.plsqltransactions....	NCR	50	<input checked="" type="checkbox"/>
Browse Products	com.dom.benchmarking.swingbench.plsqltransactions.B...	BP	50	<input checked="" type="checkbox"/>
Order Products	com.dom.benchmarking.swingbench.plsqltransactions....	OP	50	<input checked="" type="checkbox"/>
Process Orders	com.dom.benchmarking.swingbench.plsqltransactions.P...	PO	50	<input checked="" type="checkbox"/>
Process Orders	com.dom.benchmarking.swingbench.plsqltransactions.B...	BO	50	<input checked="" type="checkbox"/>

Overview Chart (Bottom Right):

Overview Chart showing performance metrics over time (Minutes Ago).

- Disk IO:** Block Read (yellow), Block Write (blue).
- CPU:** Usr (green), System (red), I/O Wait (blue).
- Response Time (milliseconds):** Maximum: 18, Average: 9.
- TPS (Transactions Per Second):** Maximum: 73, Average: 33.
- TPM (Transactions Per Minute):** Maximum: 3515, Average: 1543.

Users Logged On : 200



Oracle Single Instance relocation

- After the load of all swingbench users, we moved the Oracle guest

```
QUERY USER ITS0ORSI AT ALL
MOPVMEM1 : ITS0ORSI - DSC
Ready; T=0.01/0.01 16:27:47
VMRELOCATE TEST ITS0ORSI TO MOPVMEM2
User ITS0ORSI is eligible for relocation to MOPVMEM2
Ready; T=0.01/0.01 16:28:17
VMRELOCATE MOVE ITS0ORSI TO MOPVMEM2
Relocation of ITS0ORSI from MOPVMEM1 to MOPVMEM2 started
User ITS0ORSI has been relocated from MOPVMEM1 to MOPVMEM2
Ready; T=0.01/0.01 16:56:06
QUERY USER ITS0ORSI AT ALL
MOPVMEM2 : ITS0ORSI - DSC
Ready; T=0.01/0.01 16:58:09
```

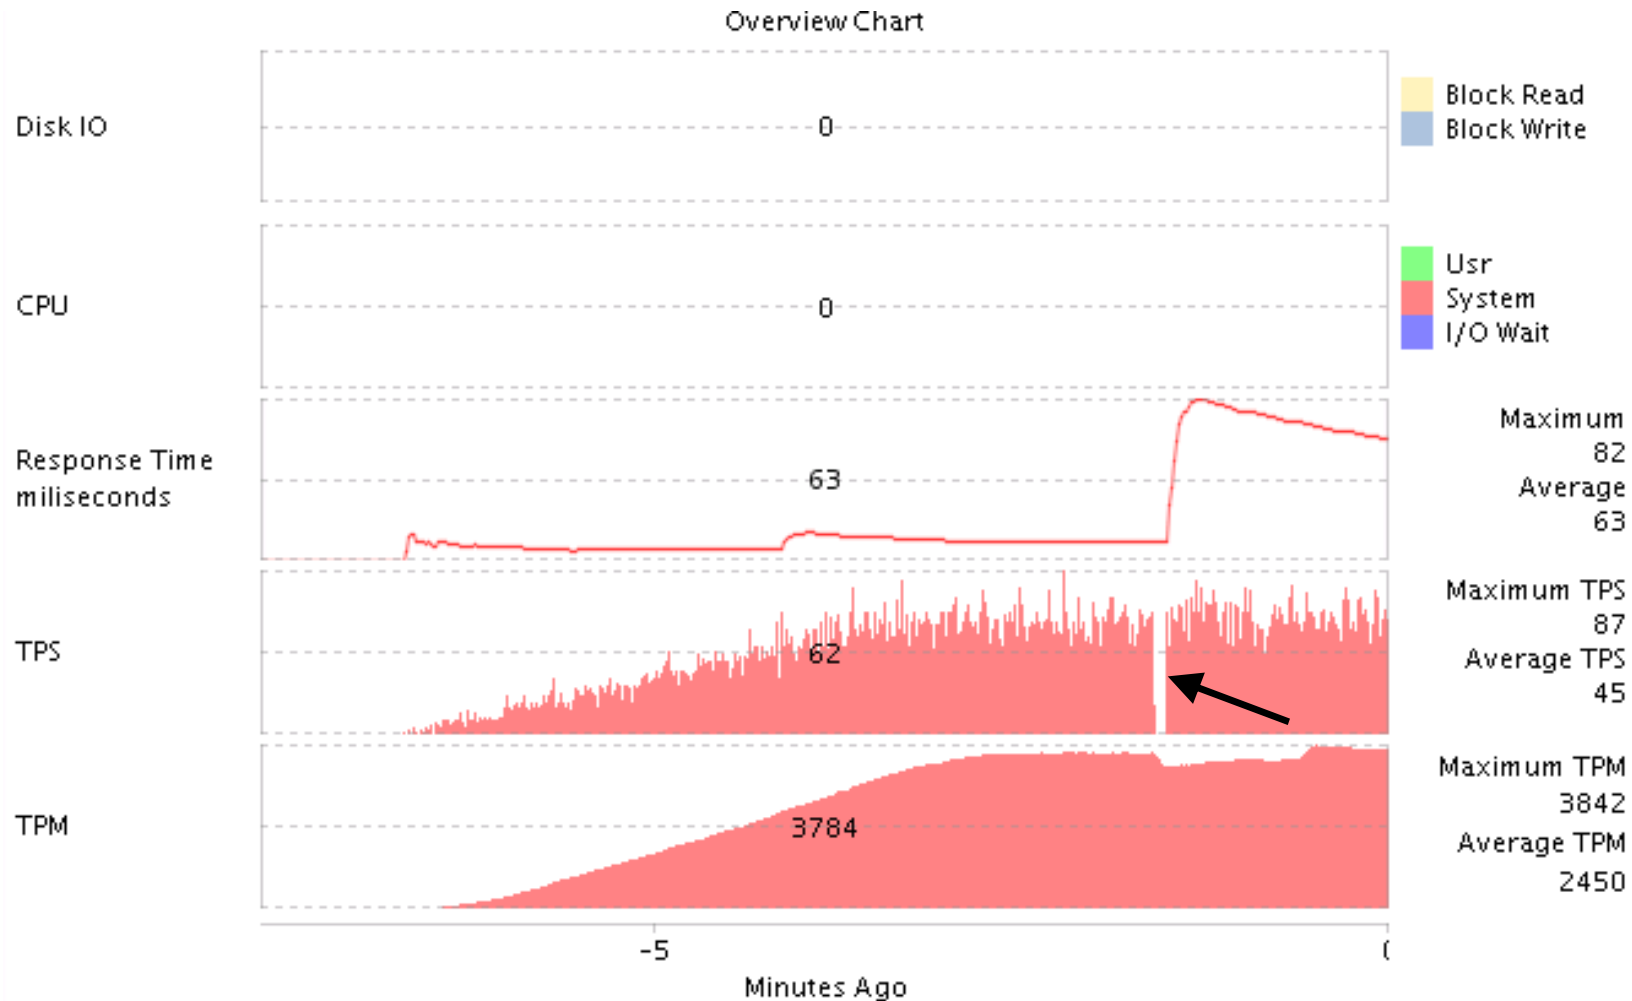
- Using a ping command running each second, we can evaluate the guest *quiesce* time (≈ 5 seconds)

```
64 bytes from 10.3.58.254: icmp_seq=41 ttl=255 time=1.02 ms Tue Nov 20 15:55:57 CET 2012
64 bytes from 10.3.58.254: icmp_seq=42 ttl=255 time=0.766 ms Tue Nov 20 15:55:58 CET 2012
64 bytes from 10.3.58.254: icmp_seq=43 ttl=255 time=0.643 ms Tue Nov 20 15:55:59 CET 2012
64 bytes from 10.3.58.254: icmp_seq=45 ttl=255 time=8.54 ms Tue Nov 20 15:56:06 CET 2012
64 bytes from 10.3.58.254: icmp_seq=46 ttl=255 time=0.436 ms Tue Nov 20 15:56:07 CET 2012
64 bytes from 10.3.58.254: icmp_seq=47 ttl=255 time=0.516 ms Tue Nov 20 15:56:08 CET 2012
```

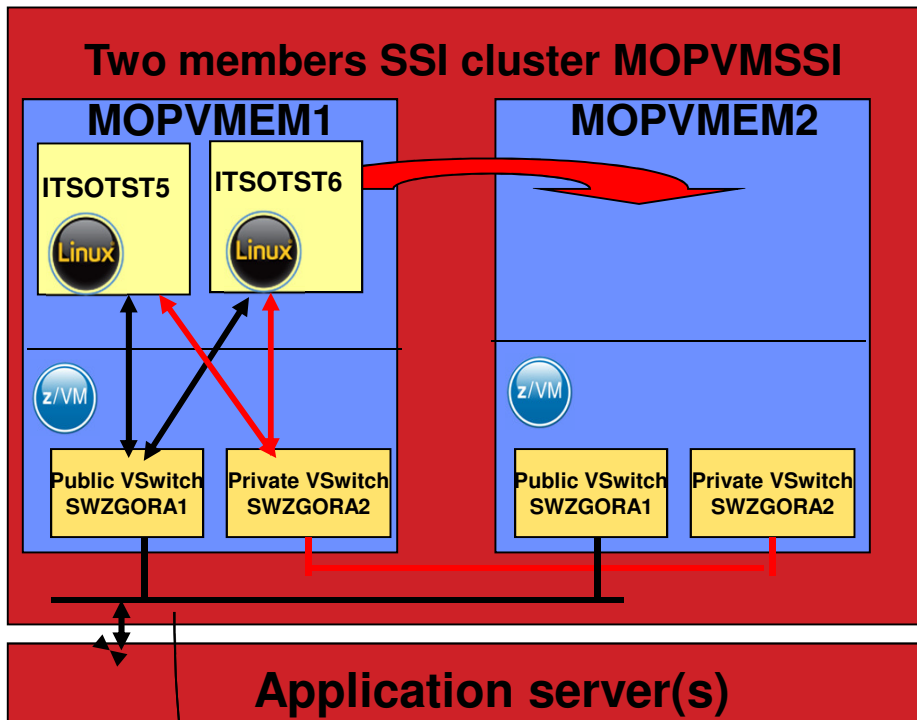
Oracle Single Instance relocation



- During the guest relocation, we can observe a freeze time (transactions are suspended) but all users stayed logged on and no transaction have been lost.



Scenario 2 & 3: Oracle RAC Nodes relocation




vmrelocate



Complete your sessions evaluation online at SHARE.org/BostonEval

- Application Server (swingbench) simulates users and generates workload to the Oracle DB
- The Oracle DB is run as a Real Application Cluster (RAC) with 2 nodes in a single z/VM 6.2 instance .
- You need to:
 - perform a maintenance on the left z/VM server partition
 - or run additional workload on the z/VM LPAR
 - or migrate the production to a new z/VM LPAR.

Scenario 2: Oracle RAC inactive Nodes relocation



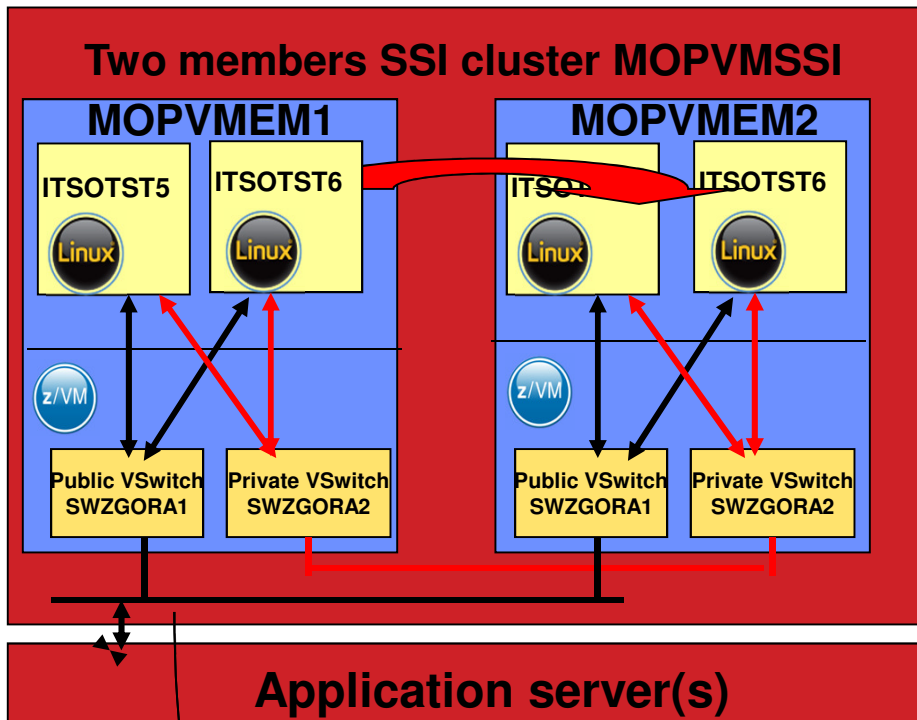
Here is the scenario if you need to stop the cluster first. This is because LGR is not yet certified on Oracle RAC.

- Stop the Oracle cluster on the node on first z/VM LPAR, keep Linux alive.
crsctl stop cluster
- Eventually activate more IFL on second z/VM LPAR
- z/VM relocate the Linux node on second z/VM
- Restart the Oracle cluster on the migrated node
crsctl start cluster

Shutdown Node
vmrelocate
Restart Node



Scenario 3: Oracle RAC active Nodes relocation



vmrelocate



Complete your sessions evaluation online at SHARE.org/BostonEval

In the scenario we relocate the nodes without taking them down

- Eventually activate more IFL on second z/VM LPAR
- z/VM relocate the nodes on second z/VM one after the other
- Must be certified by Oracle

Recommendation

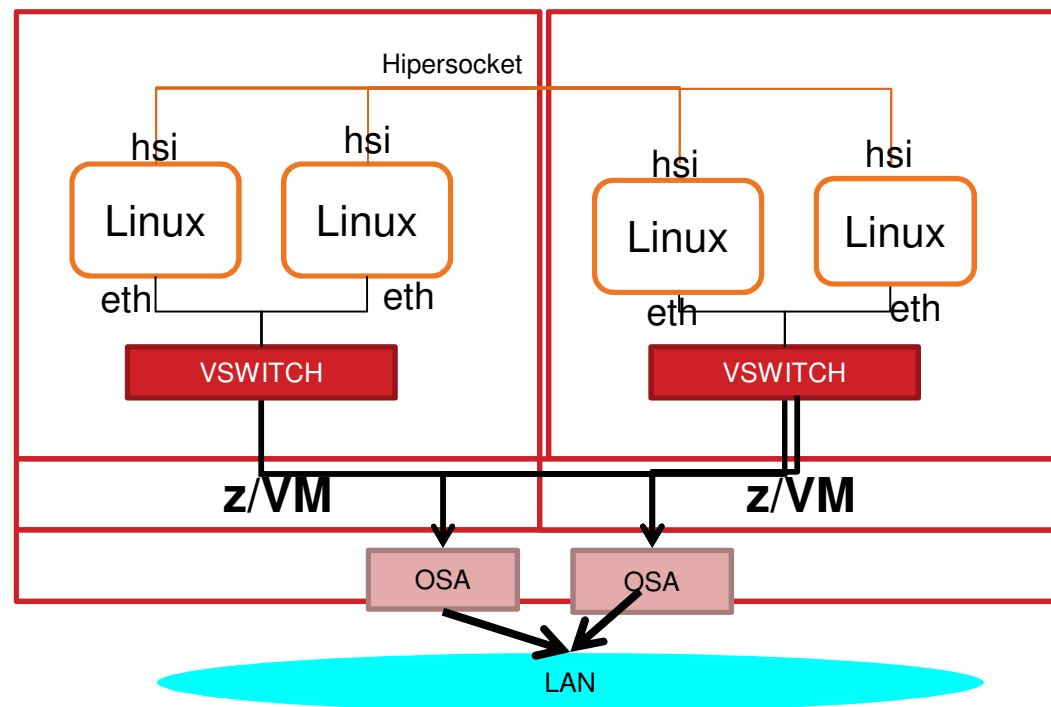
- **For production workloads use only Oracle supported methods Bring down the Oracle instance or node, relocate Linux guest, restart instance or node.**
- **For test workloads try the relocate command with active database in your environments on smaller Linux guests to determine if there is no impact**
- **Avoid very large Linux guests in a stressed environment on source or target VM.**

Agenda – Part II


- Networking alternatives for running Oracle RAC databases and Application servers connecting to Oracle databases. running Linux on System z.
- Oracle updates.
- Oracle 12c Cloud Control – How to deploy the monitoring agent.

Networking Design with Linux on System z

- High Availability Network designs are important not only for **Oracle RAC** but for Configurations with **Applications** connecting to Oracle databases running on System z.



Oracle Virtualization Support Policy:



To Bottom To Bottom

☆ Oracle Linux Support Policies for Virtualization and Emulation [ID 417770.1]

Modified: Jul 3, 2012 Type: REFERENCE Status: PUBLISHED Priority: 1

Comments (0)

Rate this document



Printable Page

This document applies to anyone wishing to use Oracle Linux running under a virtualized or emulated environment.

Note that the scope also limited to the hardware architectures supported by Oracle Linux releases.

Details

Oracle Linux Support Program provides support for:

- Oracle VM Server for x86
- The Oracle Unbreakable Enterprise Kernel running on Oracle Linux 5.5 or higher.
- Operating system support for Oracle Linux 4 (and higher) under the Oracle Linux Support Program on Oracle VM
- Operating system support for Oracle Linux 4 (and higher) under the Oracle Linux Support Program on VMware vSphere (ESX Server).
- Operating system support for Oracle Linux 4 (and higher) under the Oracle Linux Support Program on Citrix XenServer Enterprise Edition
- Xen components as part of Oracle Linux 5 and RHEL5 under the Oracle Linux Support Program. This does not include Oracle Product support on Xen offerings (see below)
- KVM components as part of Oracle Linux 5, Oracle Linux 6, RHEL5 and RHEL6 under the Oracle Linux Support Program. This does not include Oracle Product support on KVM offerings (see below)

Oracle products have been certified to run with Oracle VM. [Document 464754.1](#) lists additional information and exceptions for some Oracle products running under Oracle VM.

Oracle Products are not certified to run on Virtual Machines/guests provided by Xen or KVM offerings by Red Hat, Novell or XenSource.

Oracle Products are not certified to run on VMware vSphere (ESX Server). For more information on Oracle Product Support of Oracle Products running under VMware vSphere - refer to [Document 249212.1](#).

Oracle Products are not certified to run on operating systems on top of Red Hat, Novell SLES or Citrix XenServer Enterprise Edition Xen Hypervisors.

Oracle software stack is certified and supported on certified distributions of Linux (RHEL, SLES) running natively in LPARs or as a guest OS in z/VM Virtual Machines deployed on IBM System z 64-bit servers.

Oracle RAC Recommended Configurations for System z



Architecture	Oracle Private Network (interconnect)	Oracle Public Network
All z/VM Linux guests in one LPAR	<ul style="list-style-type: none"> • Private Layer2 VSwitch Guest LAN OSA recommended • Real layer 2 Hipersocket possible • Guest LAN Hipersocket not supported 	<ul style="list-style-type: none"> • Shared Public VSwitch recommended • Shared or dedicated OSA card is possible
z/VM Linux guests on different LPARs	<ul style="list-style-type: none"> • Real Layer 2 Hipersocket recommended • Private Layer 2 Gigabit OSA card possible 	<ul style="list-style-type: none"> • Shared Public VSwitch recommended • Shared or dedicated OSA card
z/VM Linux guests on different physical machines	<ul style="list-style-type: none"> • Private Layer 2 Gigabit OSA card recommended with physical switch in between (one hop) 	<ul style="list-style-type: none"> • Dedicated OSA card Possible



Great Paper- June 2012 Author: John P. McHugh (Oracle)



Oracle Real Application Clusters (RAC) and Oracle Clusterware Interconnect Virtual Local Area Networks (VLANs) Deployment Considerations

<http://www.oracle.com/technetwork/database/clusterware/overview/interconnect-vlan-06072012-1657506.pdf>

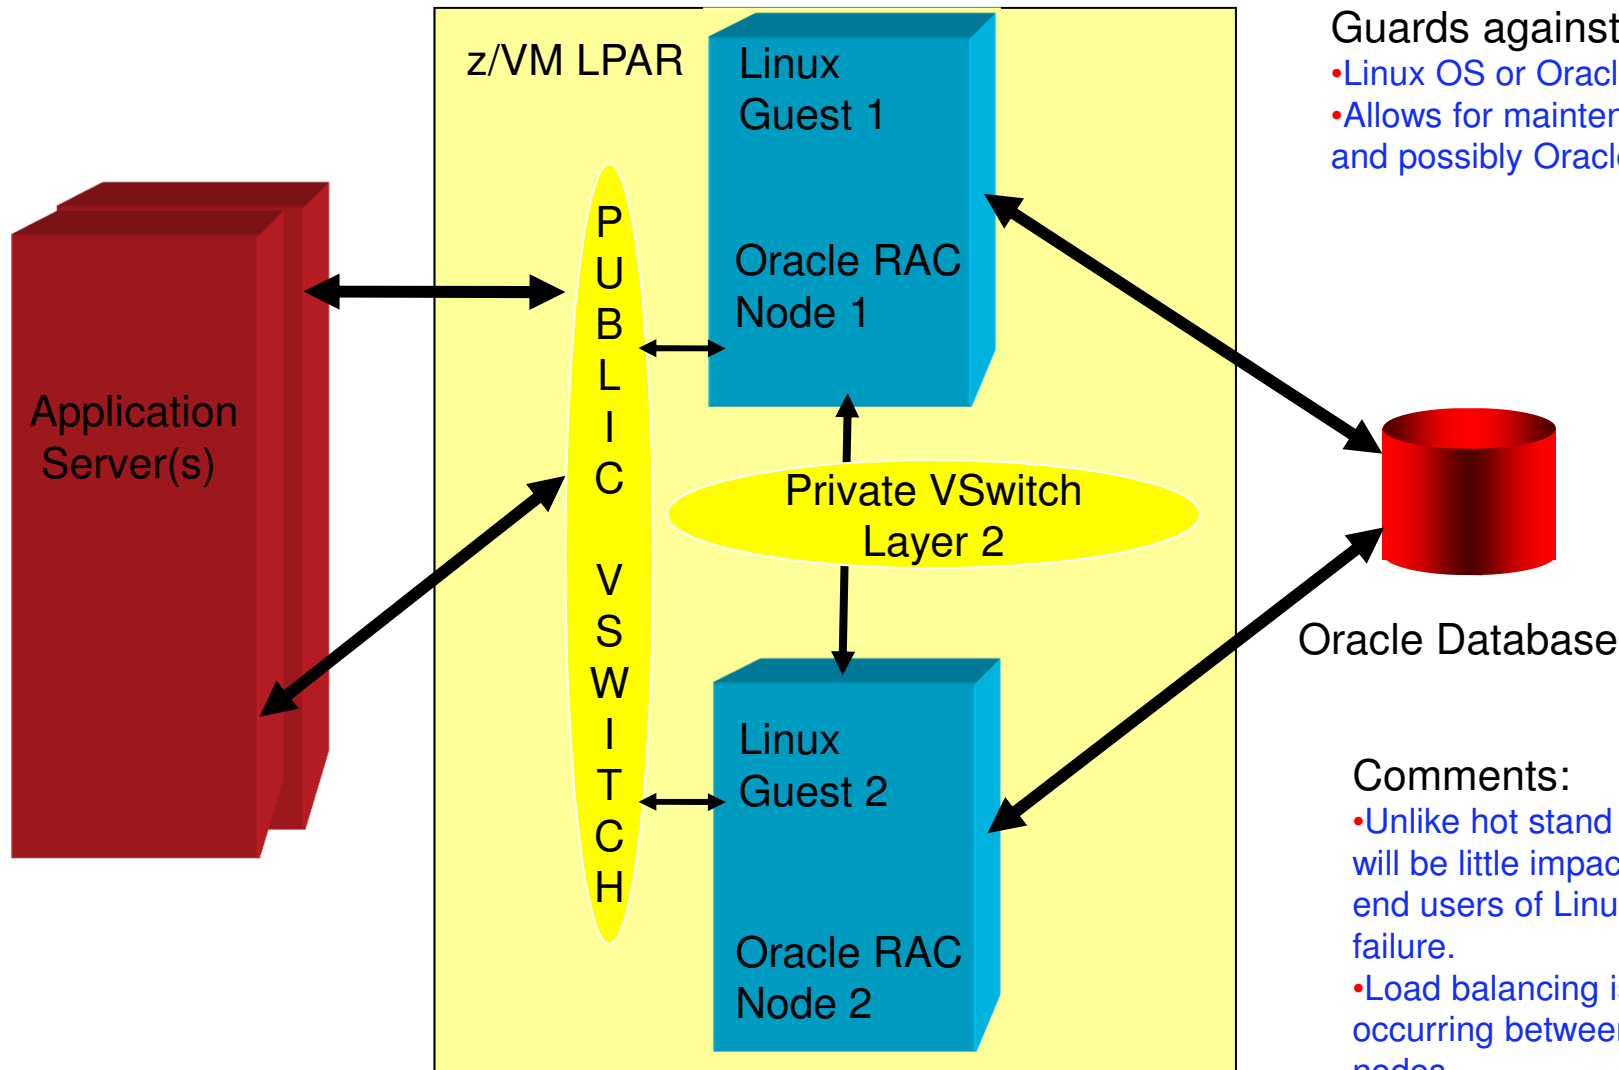
Consolidation and Converged Networks

Consolidation of RAC databases implies that network traffic can be consolidated. For Oracle Clusterware interconnect deployment, this is entirely possible. Oracle supports consolidation of RAC databases and associated private interconnect traffic on dedicated network adapters. A common consolidation is the simultaneous migration of databases to a RAC environment along with the upgrade of 1GbE network interfaces to 10GbE. The consolidated databases in the RAC environment can share the same network interface. The network interface must respect the same interconnect requirements of a dedicated, non-routed subnet. Just as in 1GbE deployments, the consolidated interconnect may be deployed in a single VLAN. If the environment requires segregated networks for the interconnect, tagged VLANs on the interface are supported for network isolation. A common consequence of network consolidation may be a reduction of required IP subnets and supporting VLANs where like-traffic is consolidated from multiple subnets to a single subnet. A common use case is where consolidated RAC databases in a single cabinet may share a single, non-routed subnet mapped to a single static VLAN on the switch.

Complete your sessions evaluation online at SHARE.org/BostonEval



Oracle Database Single LPAR with Oracle RAC



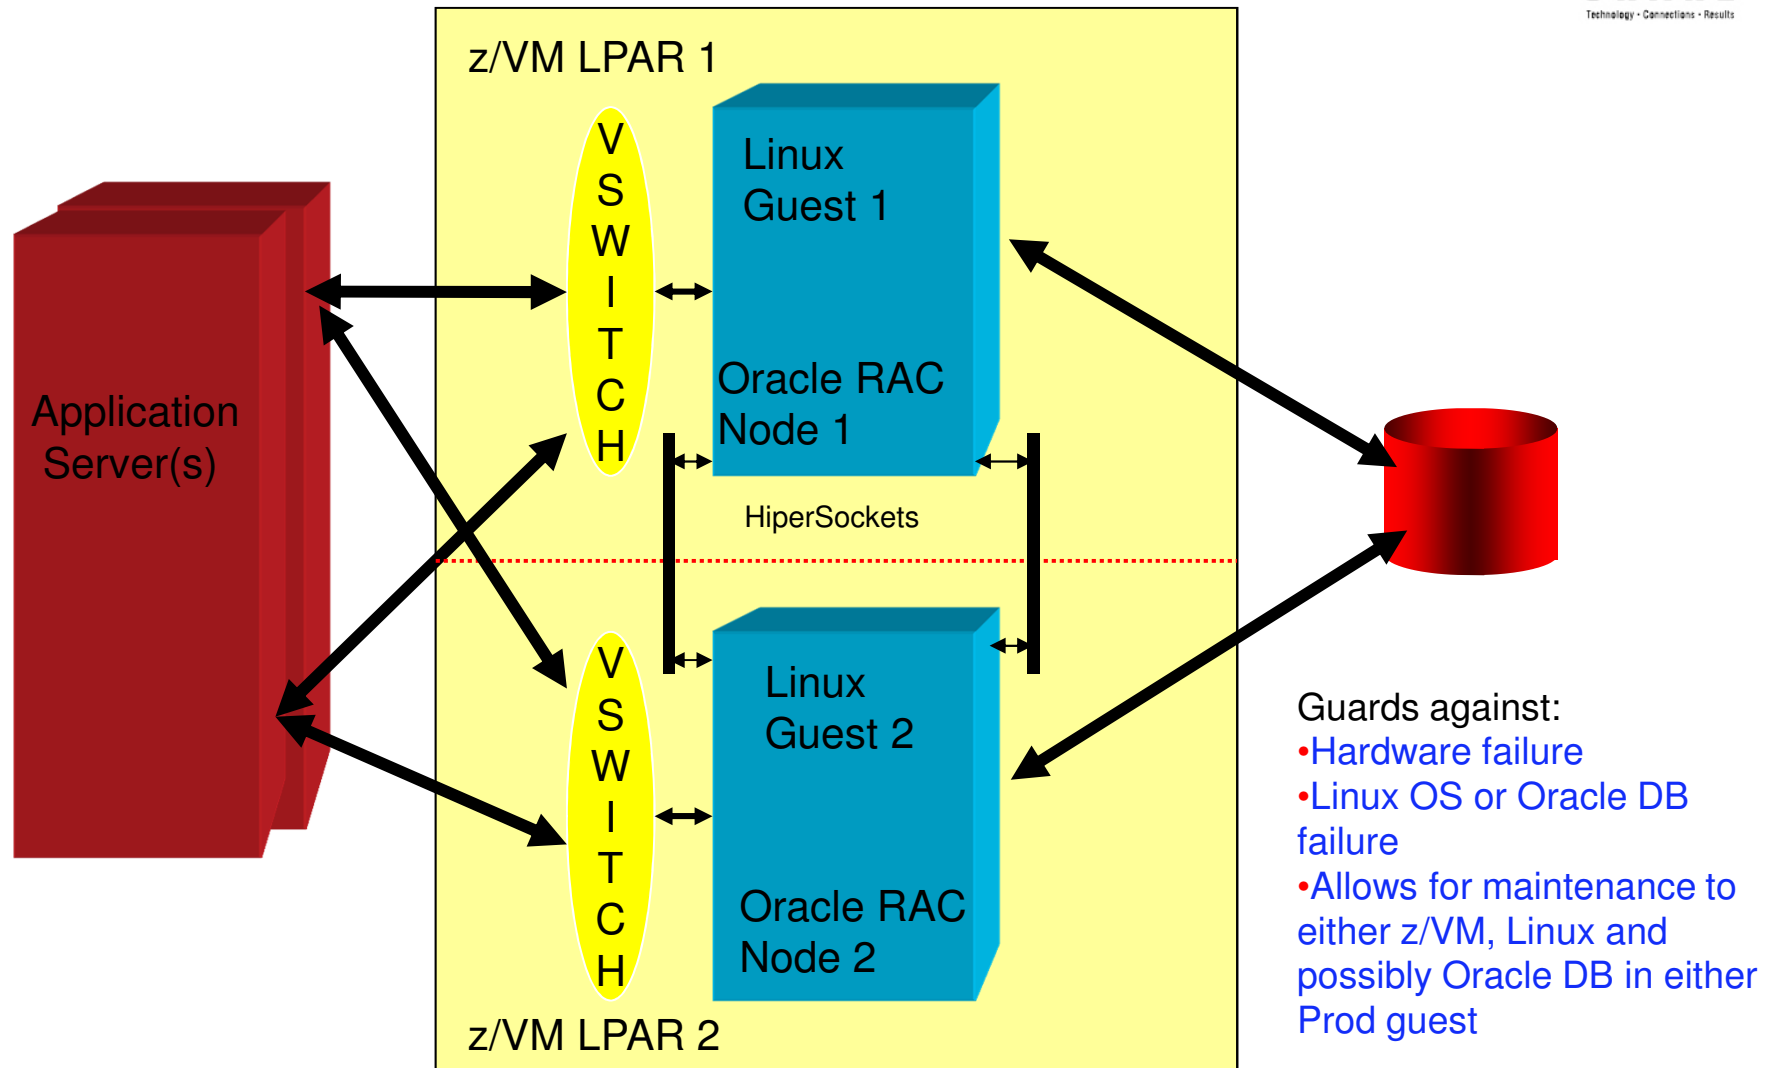
Guards against

- Linux OS or Oracle DB failure
- Allows for maintenance to Linux and possibly Oracle

Comments:

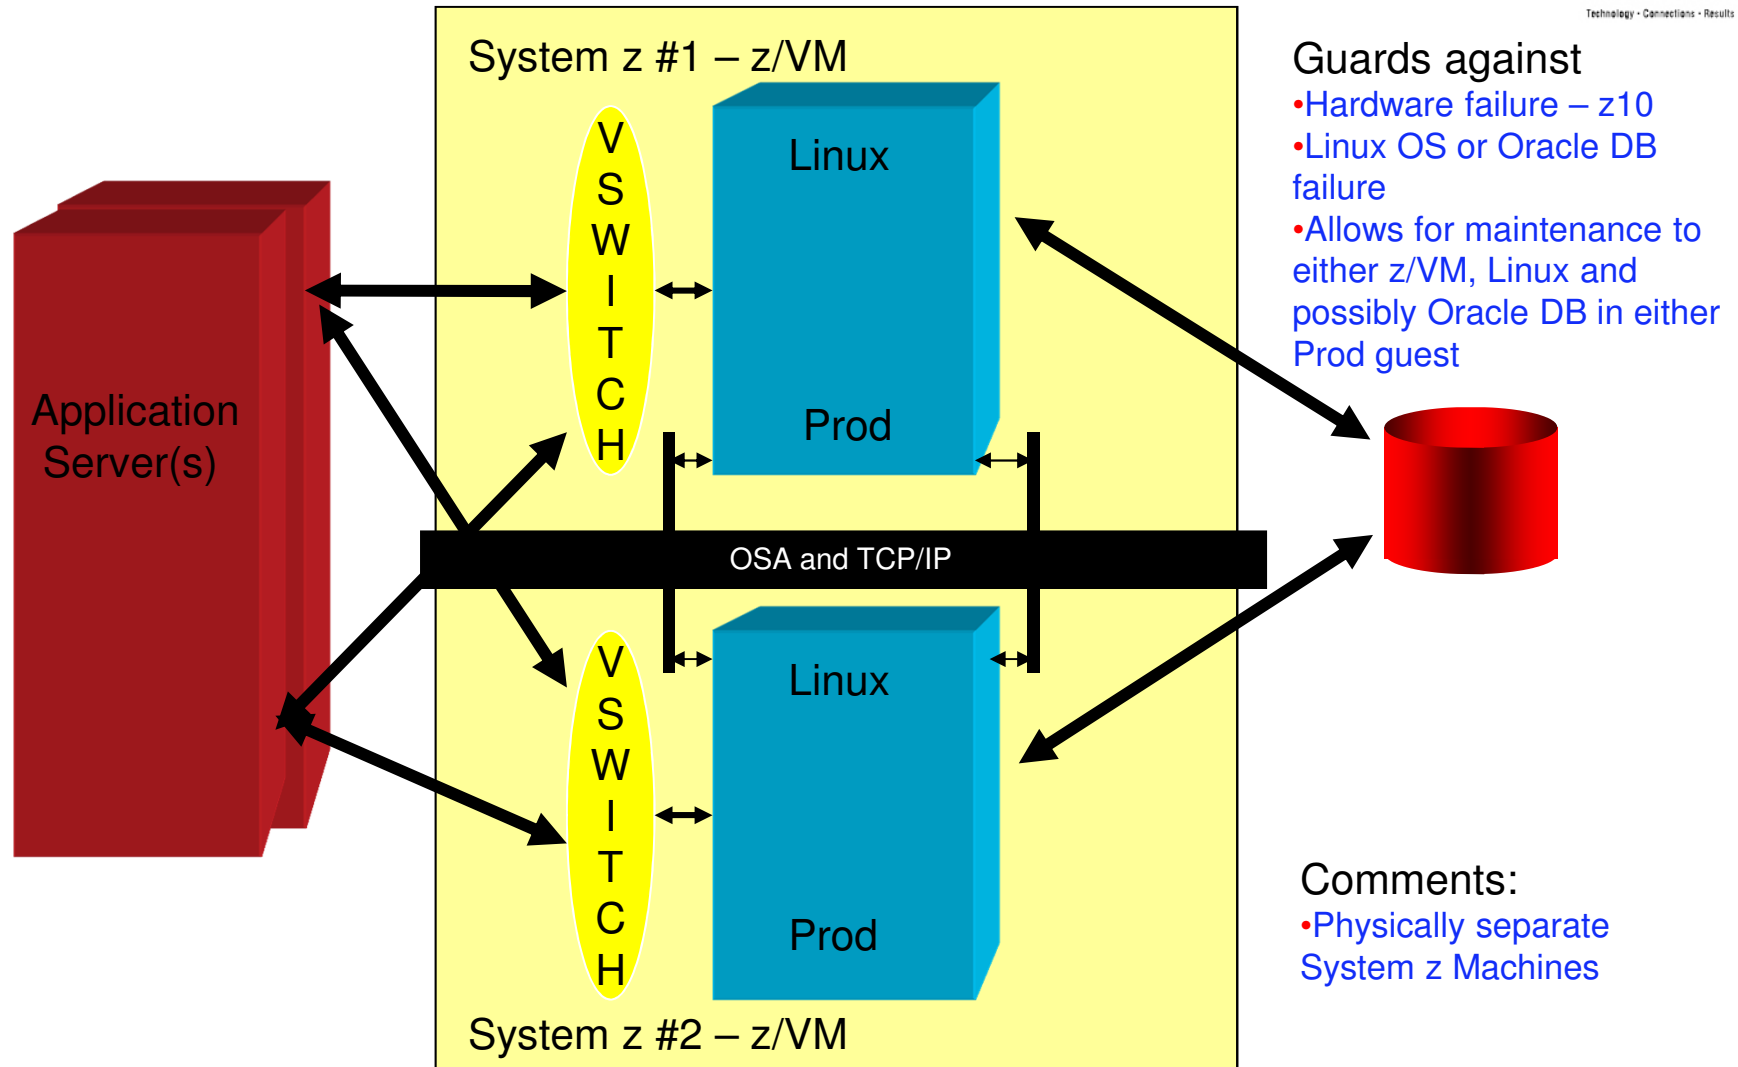
- Unlike hot stand by there will be little impact to the end users of Linux node failure.
- Load balancing is occurring between the RAC nodes.

Oracle Database – Multiple LPAR with Oracle RAC



- Guards against:
- Hardware failure
 - Linux OS or Oracle DB failure
 - Allows for maintenance to either z/VM, Linux and possibly Oracle DB in either Prod guest

Oracle Database Oracle RAC across Multiple System z Machines



Guards against

- Hardware failure – z10
- Linux OS or Oracle DB failure
- Allows for maintenance to either z/VM, Linux and possibly Oracle DB in either Prod guest

Comments:

- Physically separate System z Machines

High Availability Network Options:

- **Virtual Switch** – (Active / Passive) – When one Open System Adapter (OSA) Network port fails, z/VM will move workload to another OSA Card port. Highly available configurations should consider failover time.
- **Link Aggregation** – (Active / Active) Allow up to 8 OSA-Express adapters to be aggregated per virtual switch Each OSA-Express port must be exclusive to the virtual switch (eg. can not be shared).
- **Linux Bonding** – create 2 Linux interfaces – e.g. **eth1** & **eth2** and create a bonded interface **bond0** made up of eth1 and eth2.
- **Oracle HAIP** – New in 11.2.0.2 Oracle can have up to 4 Private interconnect interfaces to load balance interconnect traffic.

Test Plan:



- Tests involved running a single baseline test in each of the clusters, with the various network configuration options (VSwitch, Linux Bonding, and Oracle HAIP).
- Then tests were made concurrently with workload running on multiple clusters at the same time using shared Virtualized infrastructure.
- Each interconnect test consisted of 4 tests READ/READ, READ/WRITE, WRITE/READ and WRITE/WRITE

```
create table stress_ipc
(id number not null,compteur number not null, c2 varchar(30))
partition by range (id)
(
partition stress_ipc1 values less than (1) tablespace USERS,
partition stress_ipc2 values less than (2) tablespace USERS,
partition stress_ipc3 values less than (3) tablespace USERS,
partition stress_ipc4 values less than (4) tablespace USERS,
...
partition stress_ipcmx values less than (MAXVALUE) tablespace USERS);
create unique index pk_stress_ipc on stress_ipc(id,compteur) local tablespace USERS;
```
- Two dedicated 1 Gb OSA Cards were used for the Private InterConnect **shared** among 3 separate RAC Clusters. 10 Gb would be recommended for shared production.

Complete your sessions evaluation online at SHARE.org/BostonEval



Collecting Data from Oracle AWR Reports



- Review the Avg Latencies for 500B and 8K messages from AWR Reports for All Nodes in the Cluster


```
Interconnect Ping Latency Stats          DB/Inst: RACC/raccl  Snaps: 393-394
-> Ping latency of the roundtrip of a message from this instance to -> target in
-> The target instance is identified by an instance number.
-> Average and standard deviation of ping latency is given in miliseconds
-> for message sizes of 500 bytes and 8K.
-> Note that latency of a message from the instance to itself is used as
-> control, since message latency can include wait for CPU
```

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	74	.19	.06	74	.17	.07
2	74	.76	1.97	74	.94	1.94

- Latencies for Instance 1 where this report was ran to be baseline



3 Clusters Read Test Comparison – HAIP – Test18 RR



Red Hat 6.2 - 1 Node sharing Private OSA Baseline

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	9	.31	.05	9	.54	.15
2	9	.08	.00	9	.07	.00

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	9	.18	.00	9	.16	.00
2	9	.29	.05	9	.46	.03

Red Hat 6.2 - 2 Nodes sharing Private OSA Cluster B

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	10	.35	.11	10	.60	.17
2	10	.11	.00	10	.10	.00

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	10	.35	.11	10	.60	.17
2	10	.11	.00	10	.10	.00

Red Hat 6.2 - 3 Nodes sharing Private OSA Cluster C

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	10	.44	.17	10	.67	.18
2	10	.07	.00	10	.07	.00

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	10	.20	.01	10	.18	.00
2	10	.37	.18	10	.52	.13

3 Clusters Intensive Writes Comparisons – Test 20 WW HAIP



Red Hat 6.2 Baseline
1 Node

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	64	.63	1.41	64	.85	1.38
2	64	.13	.06	64	.12	.07

Cluster B Red Hat 6.2
2 Nodes

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	64	.14	.06	64	.13	.05
2	64	.42	.38	64	.64	.36

Cluster C Red Hat 6.2
3 Nodes

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	71	.55	.80	71	.78	.79
2	71	.12	.04	71	.11	.04

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	71	.16	.18	71	.15	.16
2	71	.64	1.06	71	.82	1.02

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	74	.74	1.95	74	.97	1.91
2	74	.08	.02	74	.07	.01

Target Instance	500B Pin Count	Avg Latency 500B msg	Stddev 500B msg	8K Ping Count	Avg Latency 8K msg	Stddev 8K msg
1	74	.19	.06	74	.17	.07
2	74	.76	1.97	74	.94	1.94

Complete your sessions evaluation online at

Choose the Correct Network MTU size


netstat -s of Interconnect	MTU Size of 1492 (default)	MTU Size of 8992 (with 8K DB block size)
Before reassemblies	43,530,572	1,563,179
After reassemblies	54,281,987	1,565,071
Delta assemblies	10,751,415	1,892

Network Queue Length

- The device queue length should be increased from the default size of 1000 to at least 2000 using sysctl:

sysctl -w net.core.netdev_max_backlog =2000

Oracle RAC - Scaling device queue length

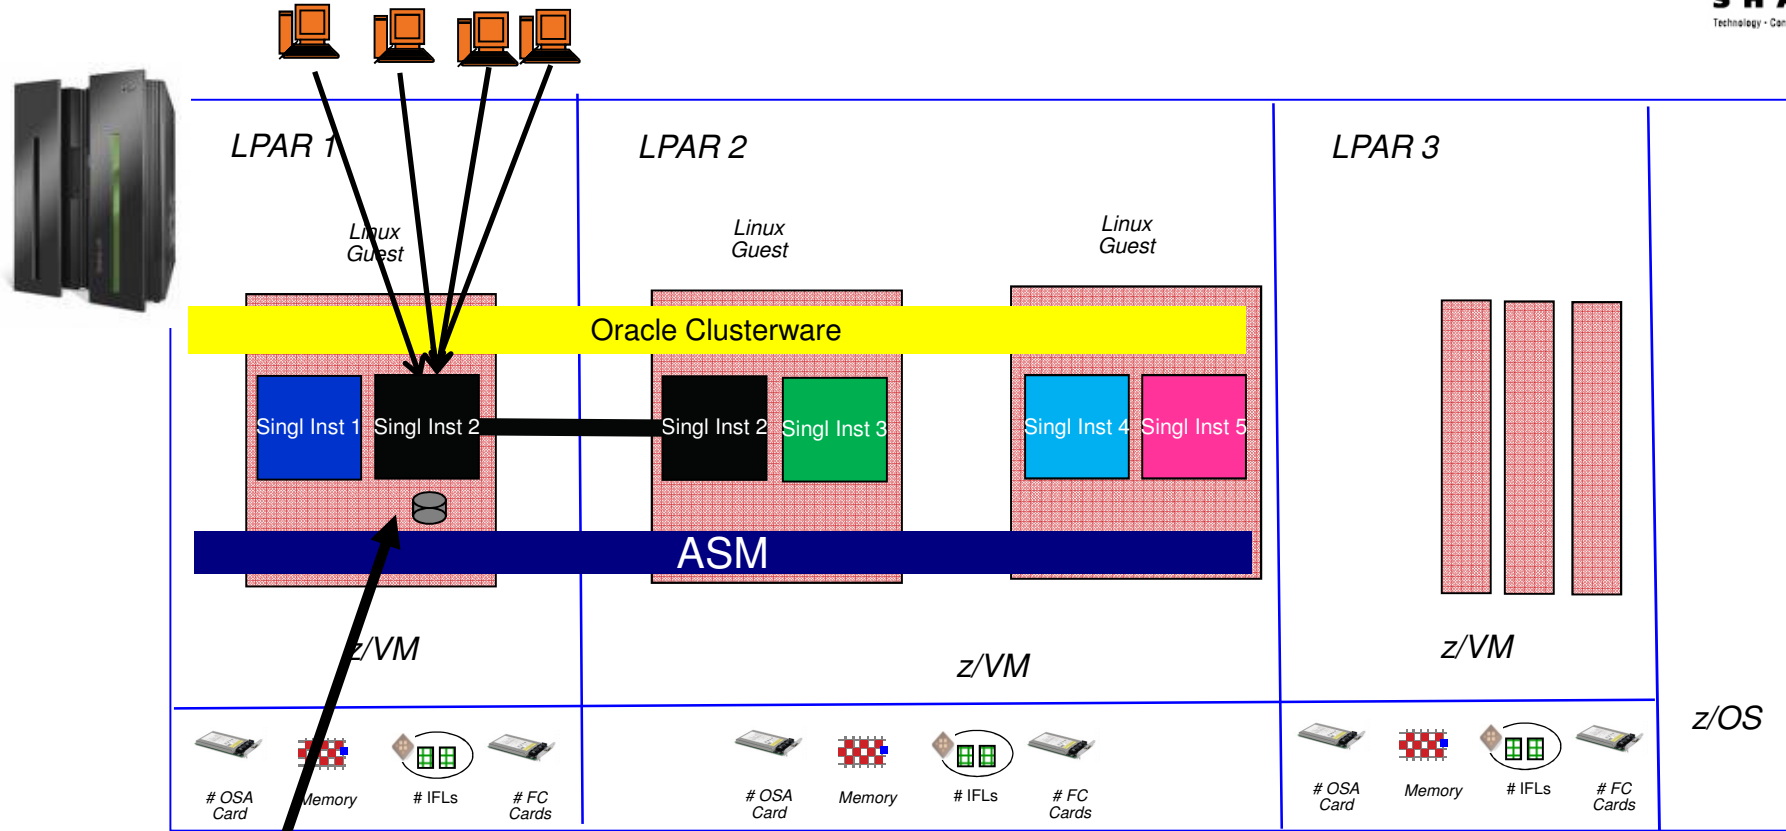


Best Practices for Consolidating on System z




- Dedicated **Multiple** 1Gb or Greater NICs per Oracle cluster for the Private Interconnect.
- Private Interconnect can be shared but watch AWR Ping latencies.
- Isolate Oracle RAC Interconnect traffic from other network traffic.
- Utilize VLANs to segregate Cluster interconnect network traffic. VLANs are fully supported for Oracle Clusterware interconnect deployments on System z.
- Utilize Server Pools - Logical division of the cluster into pools of servers with many Linux Guests utilizing one cluster interconnect.
- Oracle RAC nodes on the same System z can utilize System z Hipersocket(s) for the interconnect traffic defined on layer 2.
- Guest LAN Hipersocket does not support Layer 2, hence is not supported.
- Single LPAR RAC configurations should be used for Test/Development only.
- Configure Network switches so that **VLAN Interconnect traffic is Pruned from Uplink traffic.**

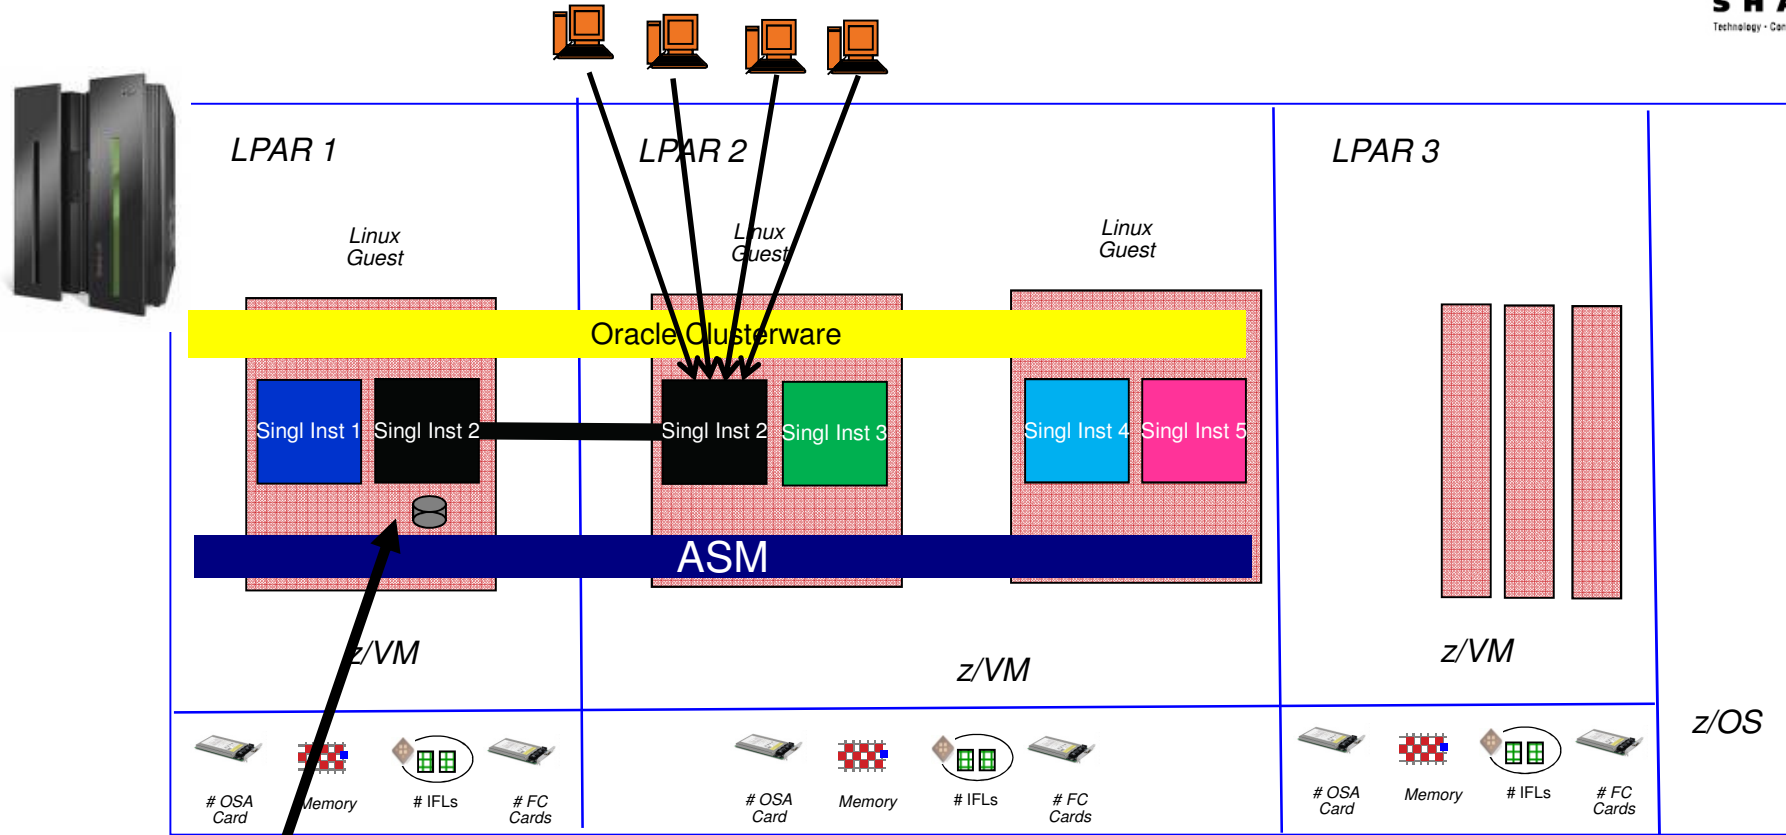
Oracle RAC One Node - deployment Omotion



- Patch Oracle binaries, modify Linux parameters, etc..
- Groups




Oracle RAC One Node - Omotion



- Restart Instance Service

ASM Disk Groups



OCR and Voting

RAC ONE – Transparent Application Failover (TAF)



- My Oracle Support (MOS) Note – **453293.1**
- TAF is required to be configured post install or else failover will not work properly.
- 11gR2 allows for selects to failover seamlessly (READ TAF)
- Database 12c will allow for READ-WRITE TAF where an API can replay on failover instance.

```
srvctl modify service -d <db> -s<service> -P BASIC -e SELECT -z 180 -w 5 -m BASIC -j SHORT
```

MODULE	MACHINE	LOGON_TIM	INST_ID	OSUSER	FAILOVER_TYPE	FAILOVER_M
SQL*Plus	ora-raca-2	01-FEB-13	2	oracle	SELECT	BASIC
SQL*Plus	ora-raca-2	01-FEB-13	4	oracle	SELECT	BASIC

-e :failover type (NONE, SESSION, SELECT)

-z: failover retries

-w: fail over retry time

-m: failover method

-j: connection load balancing

Suggestion: check that your connections show as Failover eligible.

Complete your session's evaluation online at SHARE.org/BostonEval



Best Practices for Consolidating Networks



- Keep Private Interconnect traffic separate from Public network interfaces. (separate network, cards, VSWITCH etc.)
- Private Interconnect's can be shared but watch AWR Ping latencies.
- Utilize VLANs to segregate Cluster interconnect network traffic. VLANs fully supported for Oracle Clusterware interconnect deployments on System z.
- **Utilize Server Pools** - Logical division of the cluster into pools of servers with many Linux Guests/ databases utilizing the one cluster interconnect.
- RAC nodes on the same System z host can utilize System z Hipersocket(s) for interconnect traffic defined on Layer 2. (Guest LAN Hipersocket not supported)
- Single LPAR RAC configurations recommended for Test/Development only to avoid outages for z/VM maintenance.
- Configure Network switches so that **VLAN Interconnect traffic is Pruned from Uplink network traffic.**

Oracle Certification Updates – Linux on System z

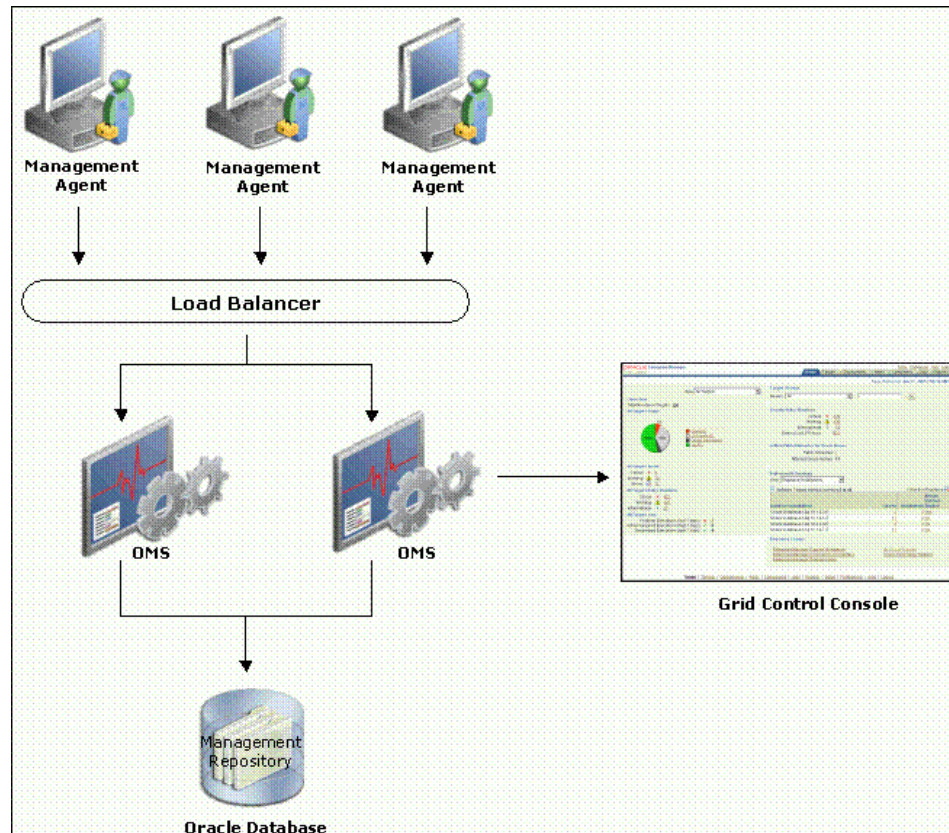


- **Grid Agent 12c - 12.1.0.3** - July 2013
- **WebLogic 12c - 12.1.2** - July 2013
- **Oracle Database Patch Set 11.2.0.3.7** - July 2013
same date as other platforms (patch - 16619182).
- **SuSe 11 SP3** just released and is certified (and tested)
for 11gR2 on System.

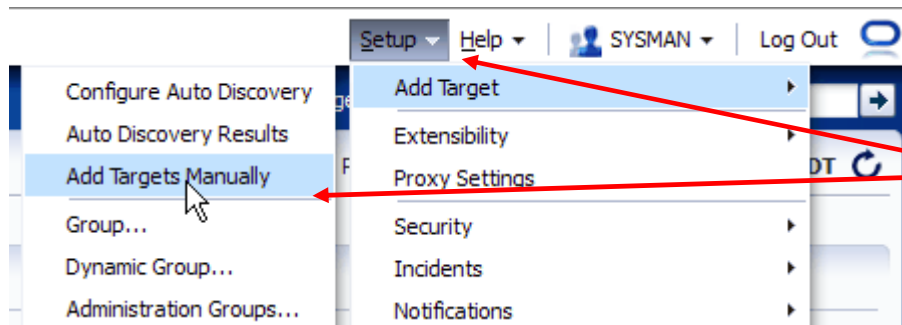
Oracle 12c Cloud Control Monitoring Agent



- **New Oracle 12.1.0.3 grid agent** certified for Linux on System z
- Create repository database on any supported platform (including System z)
- Install Oracle Management Server (OMS) on a supported platform.
- Deploy / “push” Oracle agent to System z databases by simply entering hostname and login credentials.

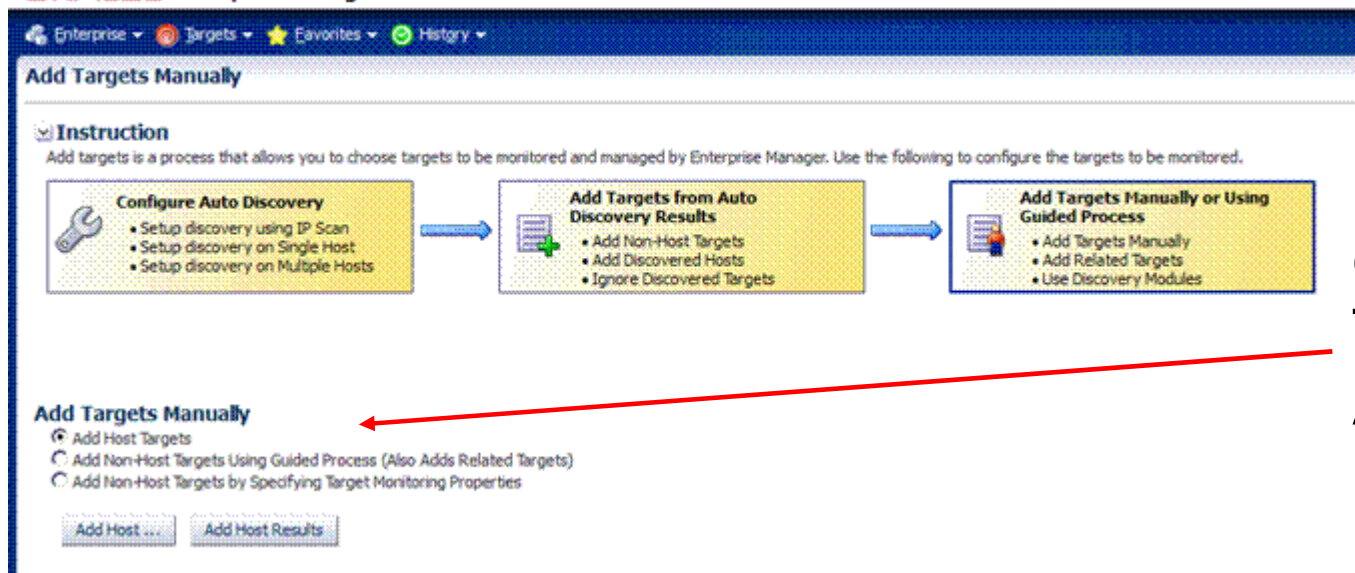


Oracle Cloud Control – Adding a monitoring target (1)



Got to Setup-> Add Targets manually -> Add Target

ORACLE Enterprise Manager Cloud Control 12c



Click Add Host Targets, then Add Hosts button

Oracle Cloud Control – Adding a monitoring target (2)



ORACLE Enterprise Manager Cloud Control 12c

Add Target

Host and Platform Installation Details Review

Add Host Targets : Host and Platform

This wizard enables you to install Management Agents on unmanaged hosts, thereby converting them to managed hosts. Enter a session name

* Session Name

Add Remove | Load from File Add Discovered Hosts | Platform

Host	Platform
<input type="text" value="lnxcd2n1.itso.ibm.com"/>	IBM: Linux on System z

Enter the hostname(s) to monitor then "Next"

TIP The target host's platform is defaulted based on a combination of factors, including hints received from automated discovery and the OMS host. The default is a suggestion, however, we recommend you to check the platform details before processing to the next step.

TIP If the platform name is appended with "Agent Software Unavailable", then download the software for that platform using Self Update.

Complete your sessions evaluation online at SHARE.org/BostonEval



Oracle Cloud Control – Adding a monitoring target (4)



ORACLE Enterprise Manager Cloud Control 12c

Add Target

Host and Platform **Installation Details** Review

Add Host Targets : Installation Details

On this screen, select each row from the following table and provide the installation details in the Installation Details section.

Deployment Type
Select the type of deployment you want to perform.

- Fresh Agent Install
- Clone Existing Agent
- Add Host to Shared Agent

Platform	Agent Software Version	Hosts
IBM: Linux on System z	12.1.0.2.0	lnxcd2n1.itso.ibm.com

IBM: Linux on System z : Installation Details

* Installation Base Directory

* Instance Directory

* Named Credential

Privileged Delegation Setting

Port

Optional Details

Enter your Installation details, logon credentials etc.

Oracle Cloud Control – Adding a monitoring target (5)



ORACLE Enterprise Manager Cloud Control 12c

Enterprise ▾ Targets ▾ Favorites ▾ History ▾

Add Host Status

Remote Prerequisite Ch

Agent Deployment Summary : ADD_HOST_SYSMAN_Oct_23_2012_1:47:52_PM_EDT

Platform	Host
IBM: Linux on System z	lnxcd2n1.itso.ibm.com

Agent Deployment Details : lnxcd2n1.itso.ibm.com

Initialization Details

OMS Log Location [zs2l20-11.itso.ibm.com:/orade/mw/gc_inst/em/EMGC_OMS1/sysman/agentpush//2012-10-23_13-47-52-PM/applo](#)

Show only warnings and failures

Initialization Phase Name	Status	Error	Cause
Remote Validations	✓		
Transferring Agent Software to Destination Host	✓		

Remote Prerequisite Check Details

OMS Log Location [zs2l20-11.itso.ibm.com:/orade/mw/gc_inst/em/EMGC_OMS1/sysman/agentpush//2012-10-23_13-47-52-PM/applo](#)

Show only warnings and failures

Prereq Check Name	Status	Error	Cause
No Data Found			

Agent Deployment begins with a bunch of checks.

Complete your sessions evaluation online at SHARE.org/BostonEval



Oracle Cloud Control – Adding a monitoring target (6)



ORACLE Enterprise Manager Cloud Control 12c

Enterprise ▾ Targets ▾ Favorites ▾ History ▾

Add Host Status

Remote Prerequisite C

Agent Deployment Summary : ADD_HOST_SYSMAN_Oct_23_2012_1:47:52_PM_EDT

Platform	Host
IBM: Linux on System z	lnxcd2n1.itso.ibm.com

Agent Deployment Details : lnxcd2n1.itso.ibm.com

already registered with the inventory?	✓	
Can the host communicate with the OMS using HTTP(S)?	✓	
Does the Privilege Delegation tool exist on the remote host?	✓	
Do you have the privileges to run as root using the Privilege Delegation tool?	⚠	The "visiblepw" is not set the sudoers file and as a result, the user will not be able to run sudo over ssh.

May see warning about "root" privileges, select *continue all hosts* option

Oracle Cloud Control – Adding a monitoring target (7)



ORACLE Enterprise Manager Cloud Control 12c

Enterprise Targets Favorites History

Add Host Status

Agent Deployment Summary : ADD_HOST_SYSMAN_Oct_23_2012_1:47:52_PM_EDT

Platform	Host
IBM: Linux on System z	lnxd2n1.itso.ibm.com

Agent Deployment Details : lnxd2n1.itso.ibm.com

Agent Deployment Details

OMS Log Location: zs2l20-11.itso.ibm.com:/oracle/mw/gc_inst/em/EMGC_OMS1/sysman/agentpush//2012-10-23_13-47-52-PM/applog

Show only warnings and failures

Deployment Phase Name	Status	Error	Cause
Installation and Configuration	✓		
Secure Agent	✓		
Root.sh	●	The root.sh script was not run because the user did not have the privilege to run as root using the Privilege Delegation tool.	
Collect Log	✓		
Clean up	✓		

Oracle 12c agent then continues and is now installed monitoring your System z Oracle database.



Latest IBM/Oracle Red Book – SG24-8104 Q1 2013



- Collaboration:
IBM / Oracle / Velocity Software
- Now Available!

Draft Document for Review January 16, 2013 5:19 am

Experiences with Oracle 11gR2 on Linux for System z

Installing Oracle 11gR2 on Linux on
System z

Managing an Oracle environment

Provisioning an Oracle
environment



Sam Amsavelu
Kathryn Arrell
Gaylan Braselton
Armelle Chevé
Ivan Doboš
Hélène Grosch
Romain Pochard
David Simpson
Damian Gallagher
Michael Macisaac
Barton Robinson
Richard Smrcina

ibm.com/redbooks



Complete your sessions evaluation online at SHARE.org/BostonEval

References - Key Oracle & IBM Whitepapers:



- **Oracle Real Application Clusters (RAC) and Oracle Clusterware Interconnect Virtual Local Area Networks (VLANs) Deployment Considerations**

<http://www.oracle.com/technetwork/database/clusterware/overview/interconnect-vlan-06072012-1657506.pdf>

- **Oracle Real Application Clusters in Oracle VM Environments**

<http://www.oracle.com/technetwork/products/clustering/oracle-rac-in-oracle-vm-environment-131948.pdf>

- **Oracle Real Application Clusters on Linux on IBM System z: Set up and network performance tuning – Dr. Juergen Doelle & Margaret Phillips**

<http://public.dhe.ibm.com/software/dw/linux390/perf/ZSW03185-USEN-02.PDF>

References – Key Oracle Notes



Note 1306465.1 Getting Started 11gR2 Grid Infrastructure, Single Instance ASM and DB IBM:Linux on System z

Note 1470834.1 Requirements for Installing Oracle 11gR2 on RHEL 6 on IBM: Linux on System z (s390x)

Note 1290644.1 Requirements for Installing Oracle 11gR2 on SLES11 on IBM: Linux on System z (s390x)

Note:1476511.1 OHASD fails to start on SuSE 11 SP2 on IBM: Linux on System z

Note 1308859.1 Requirements for Installing Oracle 11gR2 on SLES 10 on IBM: Linux on System z (s390x)

Note 1306889.1 Requirements for Installing Oracle 11gR2 on RHEL 5 on IBM: Linux on System z (s390x)

Note 1086769.1 Ensure you have prerequisite rpms to install Oracle Database & AS10g IBM:Linux on System z

Note 1377392.1 How to Manually Configure Disk Storage devices for use with Oracle ASM 11.2 IBM:Linux on System z

Note 1400185.1 How to Upgrade Oracle Restart i.e. Single Node Grid Infrastructure/ASM from 11.2.0.2 to 11.2.0.3

Note 1276058.1 Oracle GoldenGate Best Practices: Instantiation from an Oracle Source Database

Note 1413787.1 How to completely remove 11.2 Grid Infrastructure, CRS and/or Oracle Restart

Note 259301.1 CRS and 10g Real Application Clusters

Note 268937.1 Repairing or Restoring an Inconsistent OCR in RAC

Note 239998.1 10g RAC How to clean up after a failed CRS Install

Note 220970.1 RAC Frequently Asked Questions Topic

Note 1082253 Requirements for Installing Oracle 10gR2 RDBMS on SLES 10 zLinux (s390x)

Note 741646.1 Requirements for Installing Oracle 10gR2 RDBMS on RHEL 5 on zLinux (s390x).

Note 415182.1 DB Install Requirements Quick Reference - zSeries based Linux .

Note 741146.1 Installing Standalone Agent 10.2 on Linux on z



Oracle Networking and High Availability Options (with Linux on System z) & Red Hat/SUSE Oracle Update

Speaker Names: David Simpson & Kathryn Arrell

Speakers Company: IBM

Date of Presentation: **Wednesday, August 14, 2013: (09:30 AM)**

Hynes, Room 308

Session Number: **13523**

Twitter -> @IBMANDOracle

