#SHAREorg

SHARE
Technology · Connections · Results

Oracle RAC Networking Alternatives on Linux on System z and Red Hat 6 Oracle DB Support

Speaker Names: David Simpson & Kathryn Arrell
Speakers Company: IBM

Date of Presentation: **Wednesday, February 6, 2013 (9:30am)**
Franciscan D, Ballroom Level

Session Number: **12758**
Twitter -> @IBMandOracle

SHARE in San Francisco 2013

Abstract / Agenda

- This session will discuss the different networking alternatives for running Oracle RAC databases on Linux on System z.
- Alternatives include VSWITCH both in active and passive mode with Link aggregation.
- Linux bonding and Oracle's HAIP Redundant Interconnect
- Support of the Oracle Database on Red Hat 6

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.**

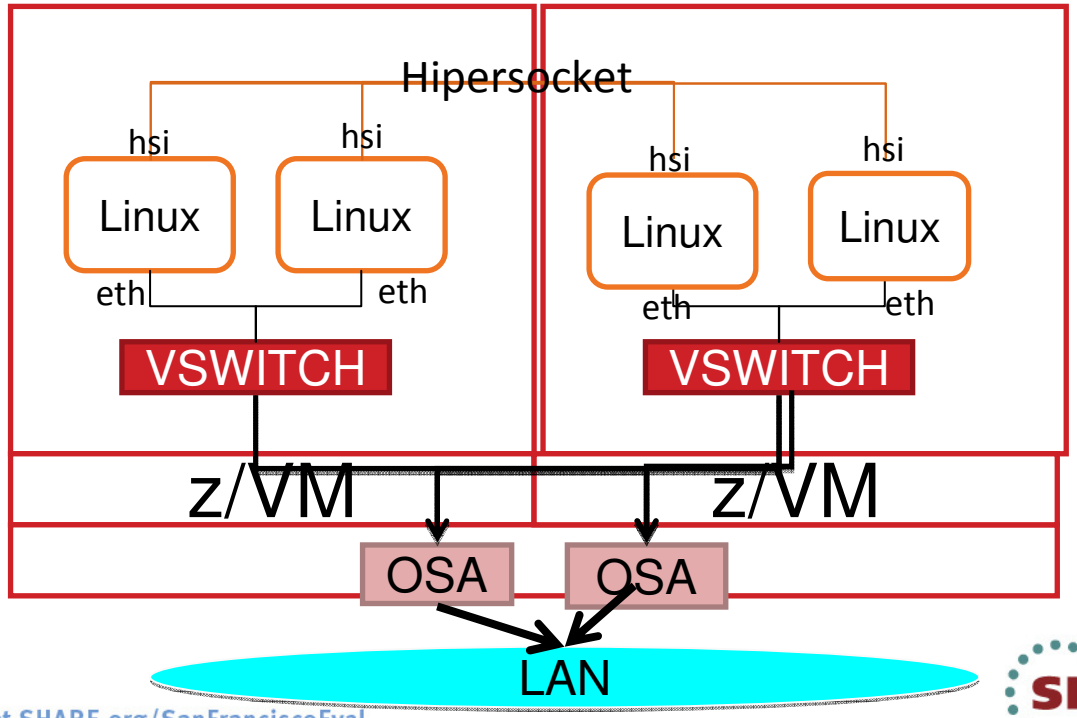

Supported Oracle RAC Configurations Linux on System z

Platform	Technology Categories	Technology	Notes
IBM zSeries Linux	Server Processor Architecture	IBM System z	Certified and supported on certified distributions of Linux running natively in LPAR or as a guest OS in z/VM virtual machines, deployed on IBM System z 64-bit servers
	Network Interconnect Technologies	•VLAN within one System z Ethernet over Gigabit OSA card for two System z •HyperSockets	

Source: http://www.oracle.com/technetwork/database/clustering/tech-generic-linux-new-086754.html

Oracle Virtualization Support Policy:

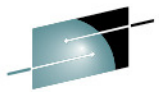

Oracle Linux Support Policies for Virtualization and Emulation [ID 417770.1]

To Bottom To Bottom

Modified: Jul 3, 2012 Type: REFERENCE Status: PUBLISHED Priority: 1 Comments (0) Rate this document Printable Pag

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.

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.

Oracle Database Single LPAR with Oracle RAC

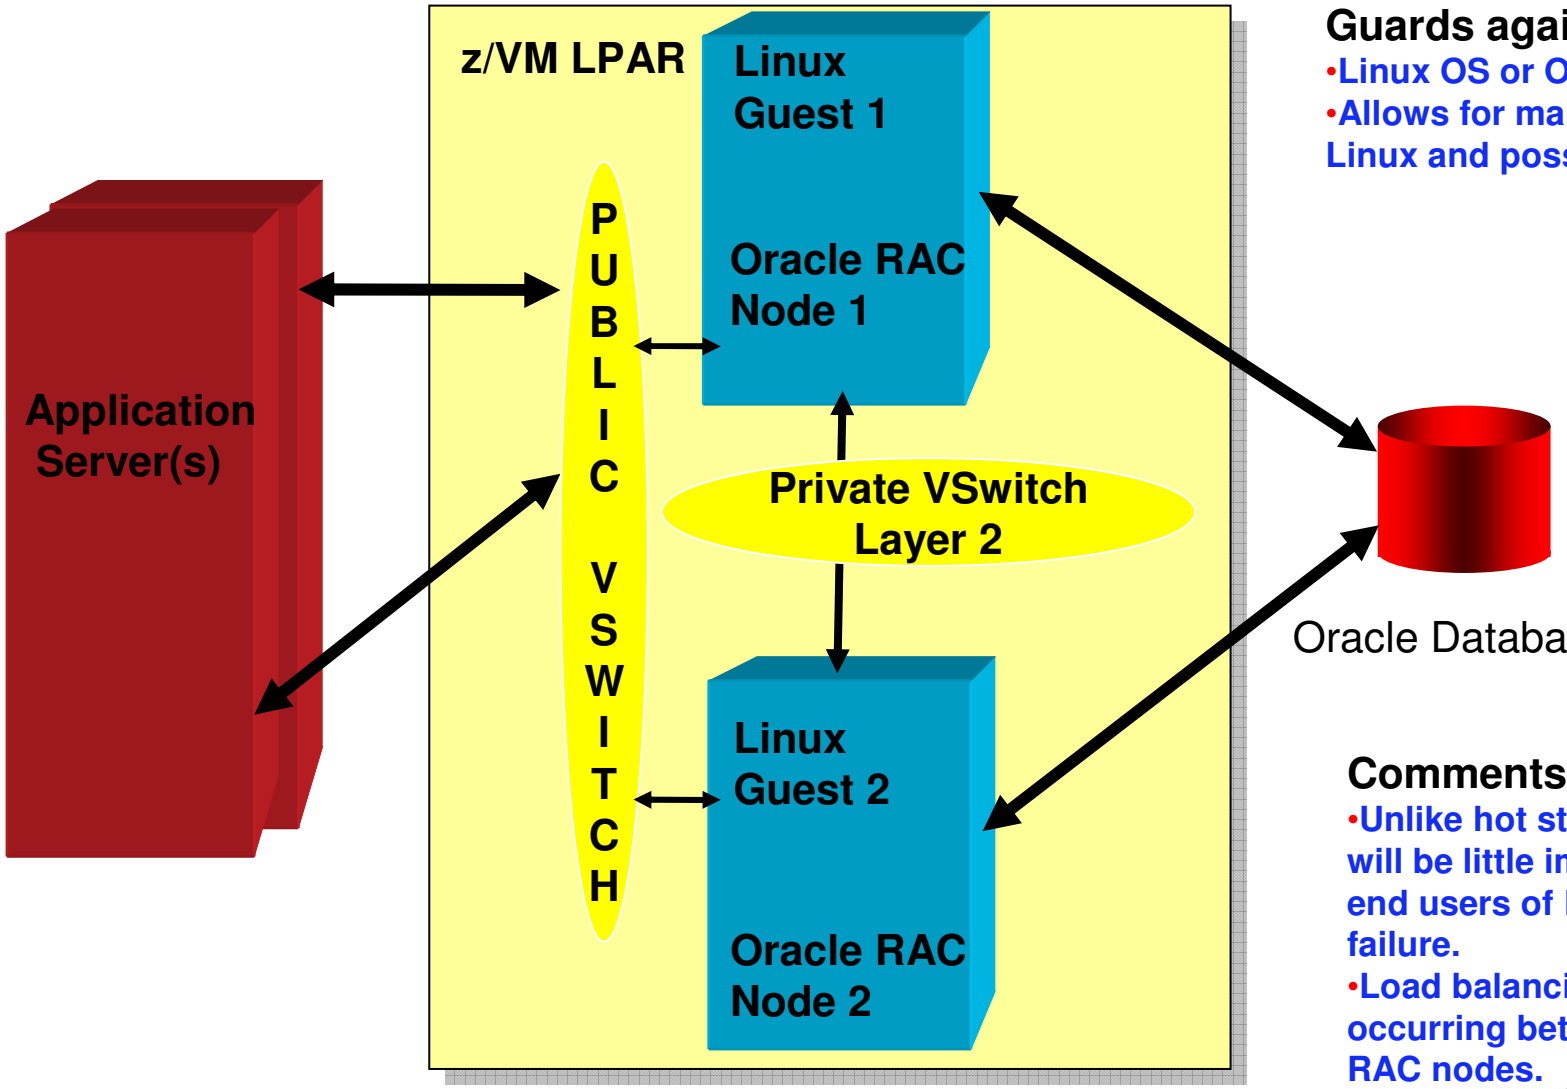

Guards against

- 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

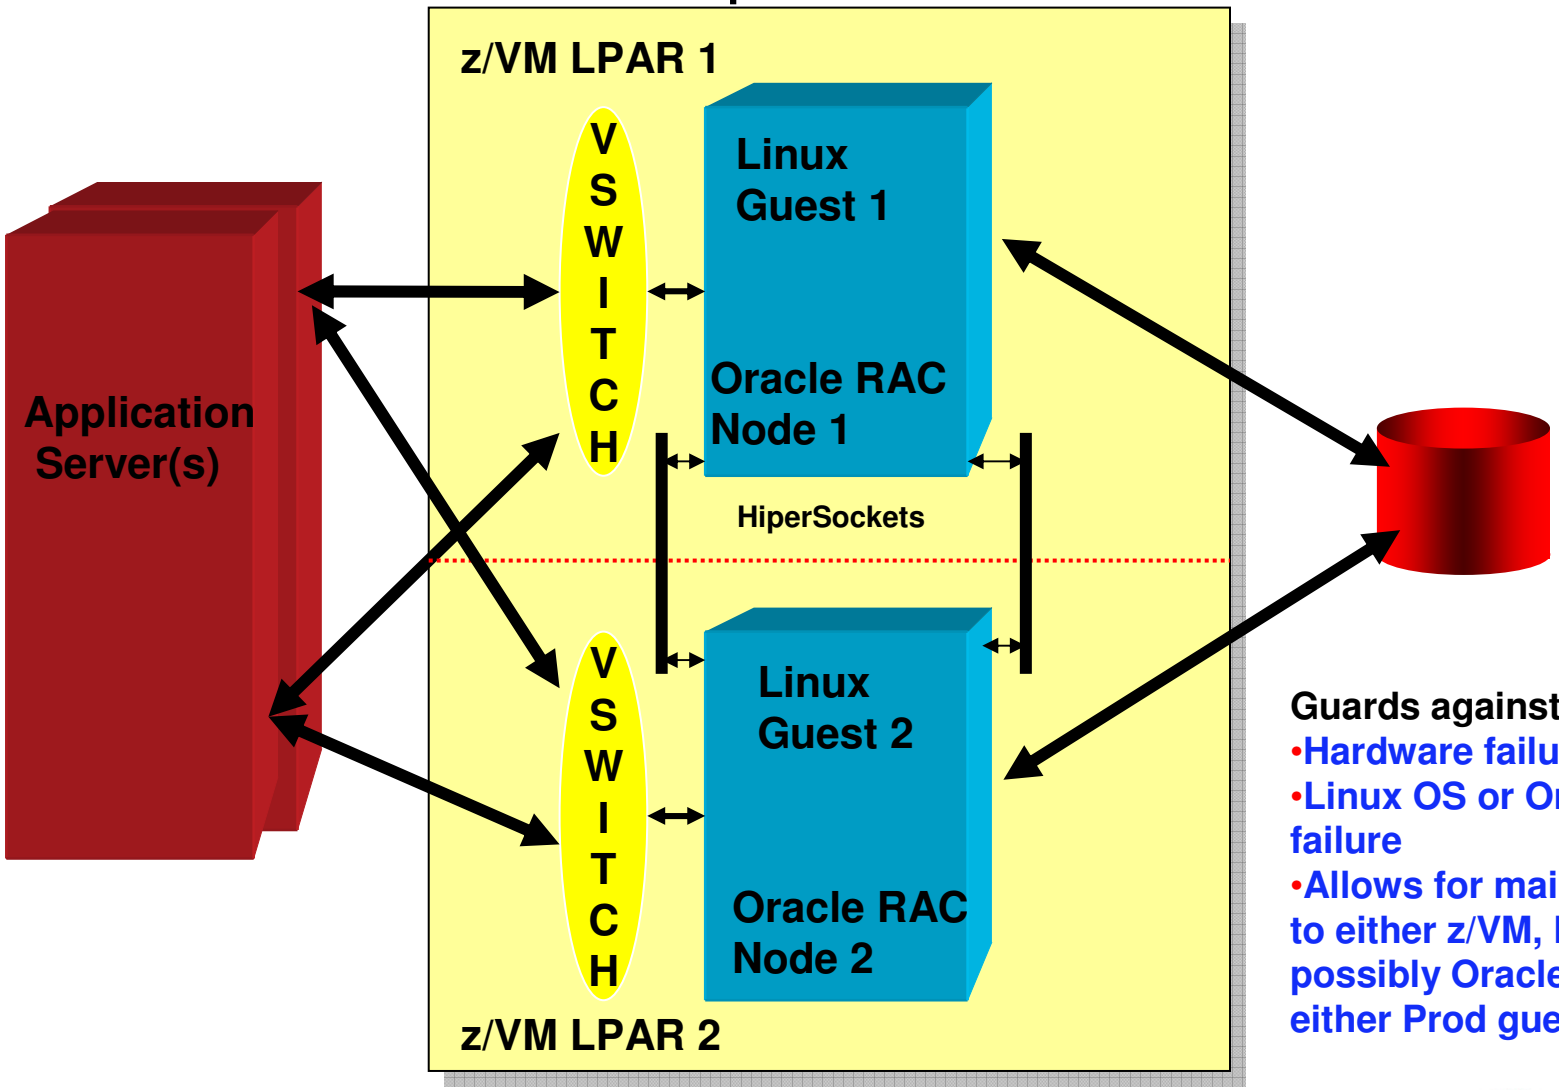

Guards against:

- **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

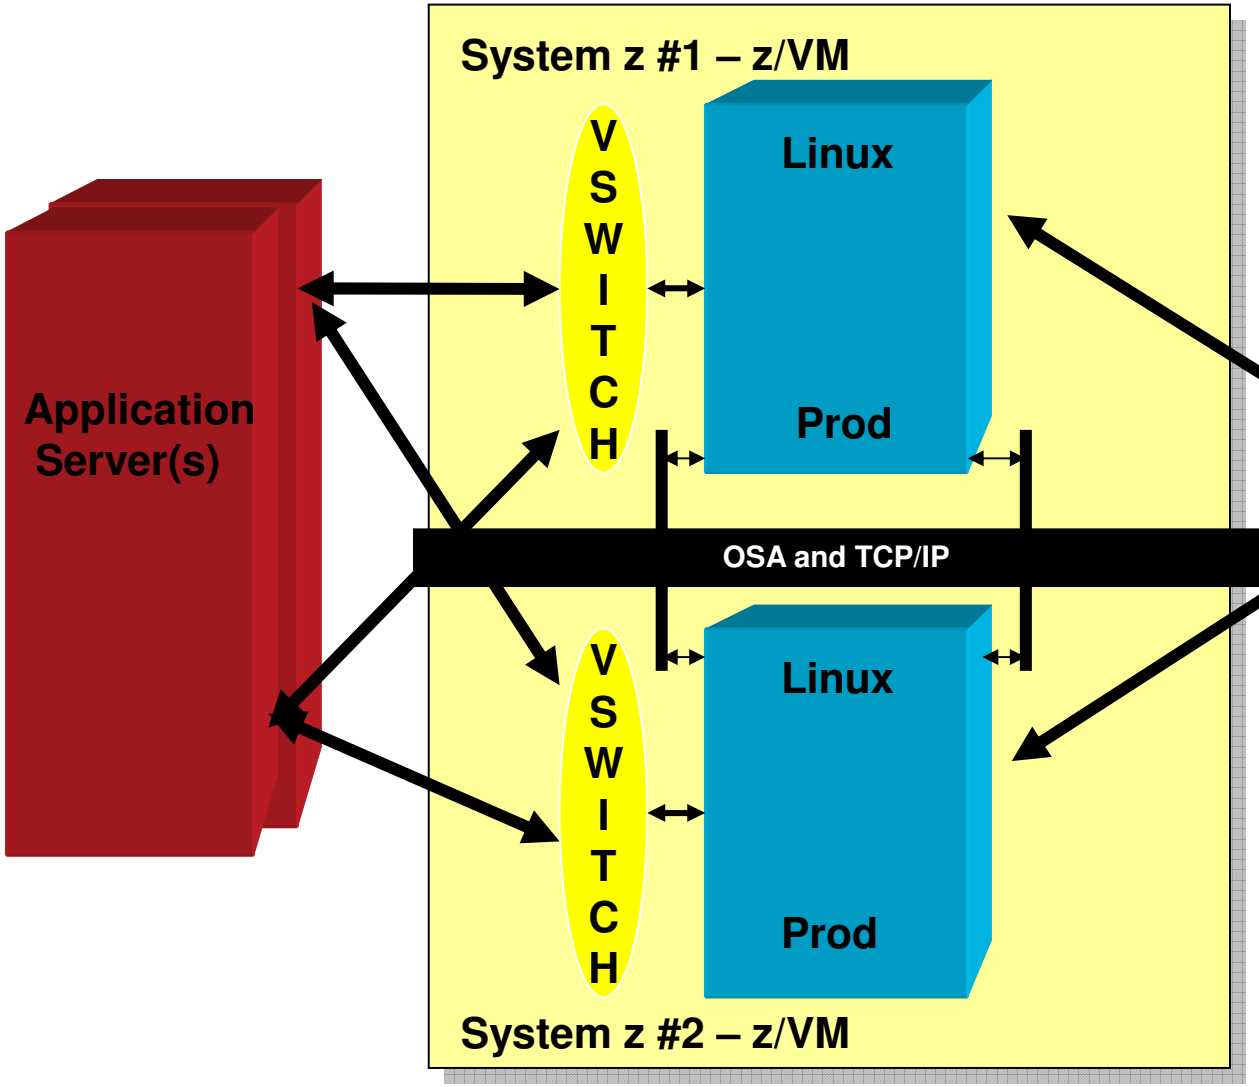

Guards against

- 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_ipcmax 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.**

Oracle Interconnect with VSWITCH (Active/Passive):

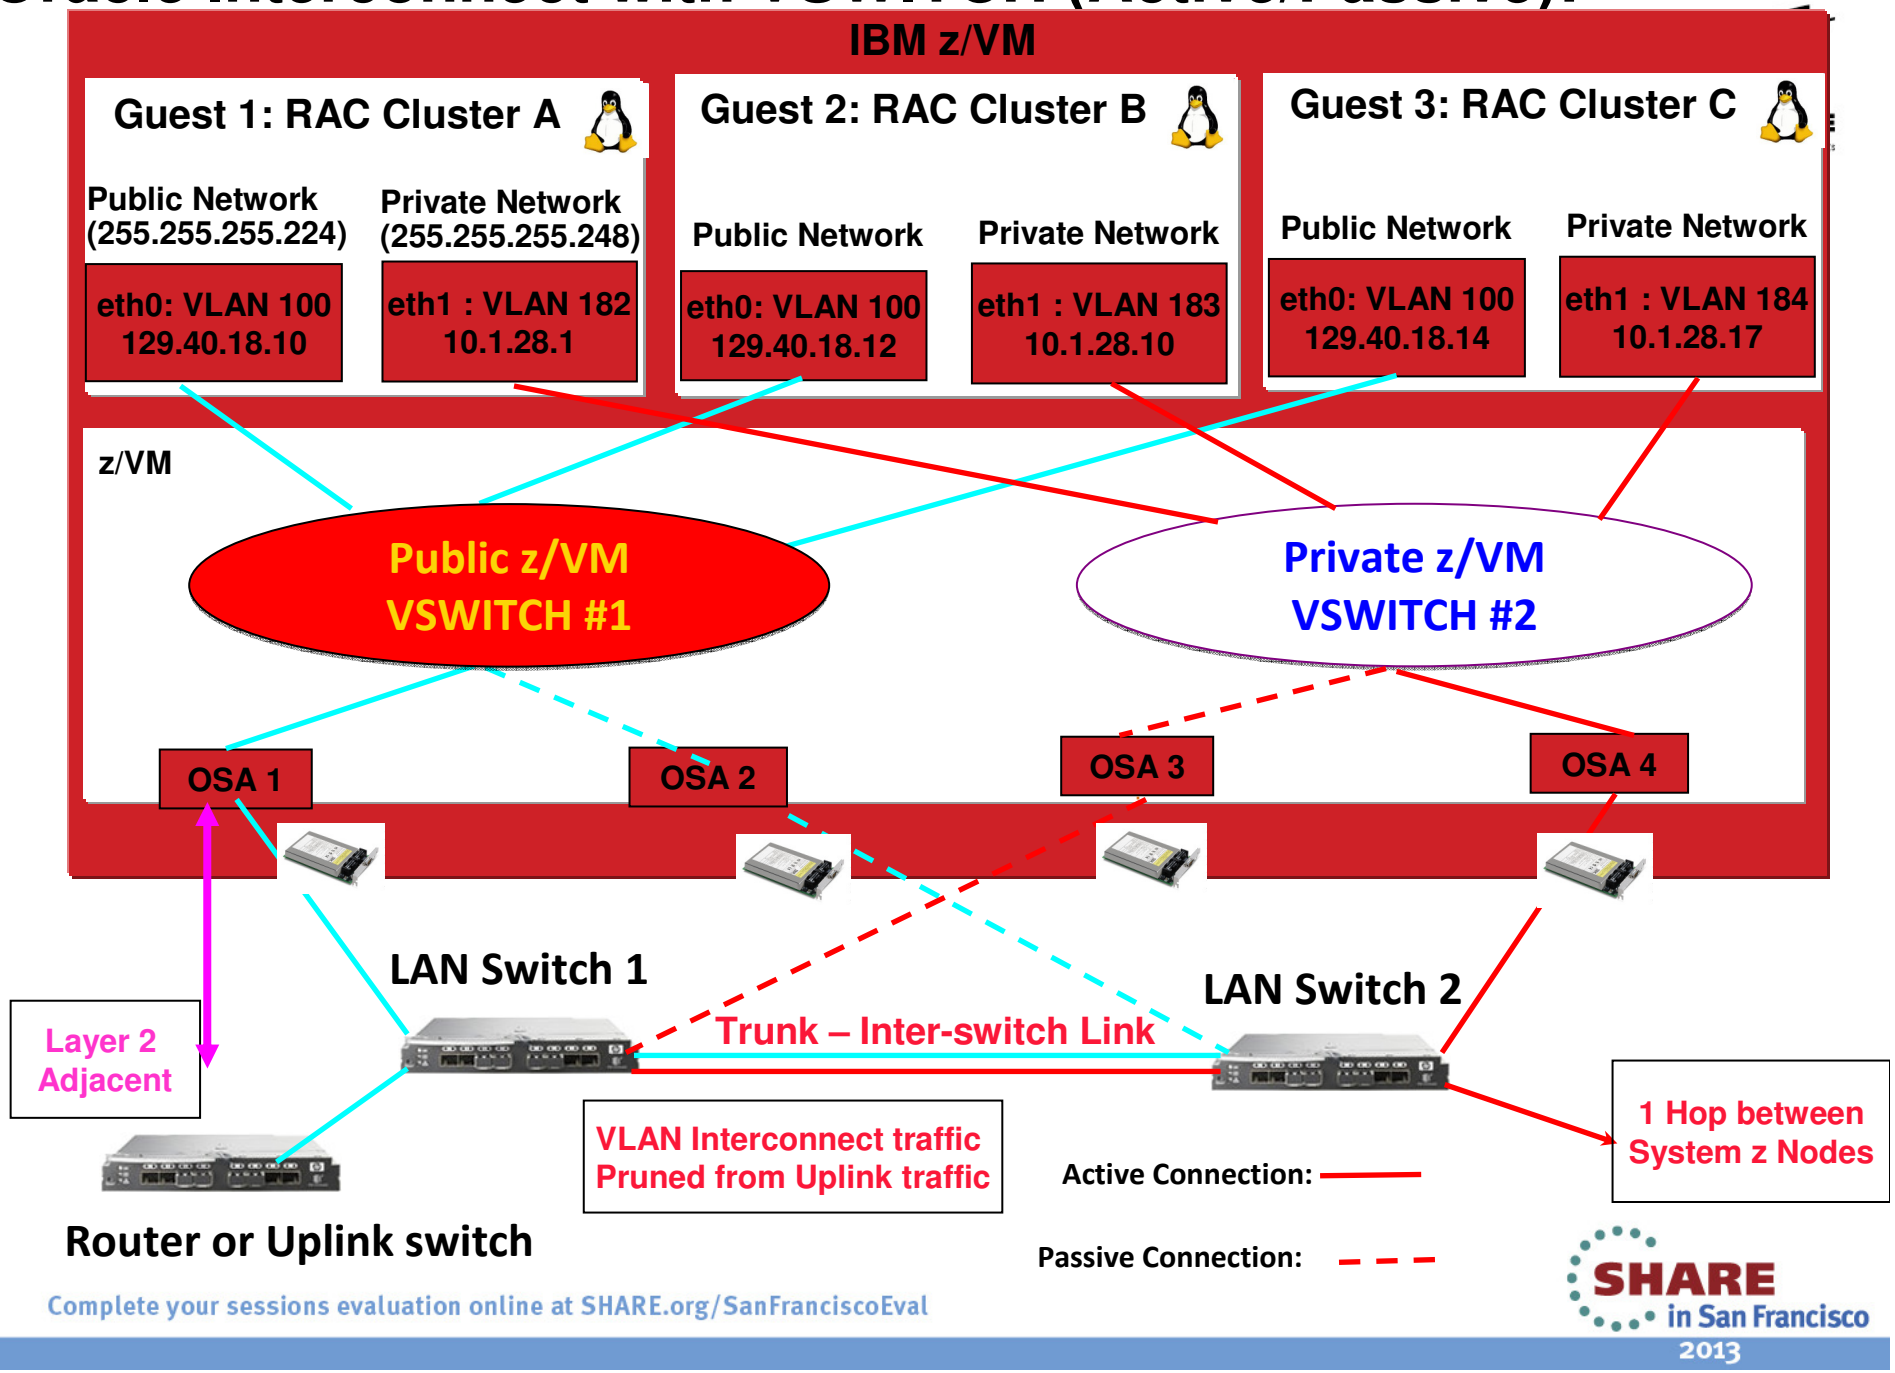

Oracle Interconnect with z/VM Link Aggregation

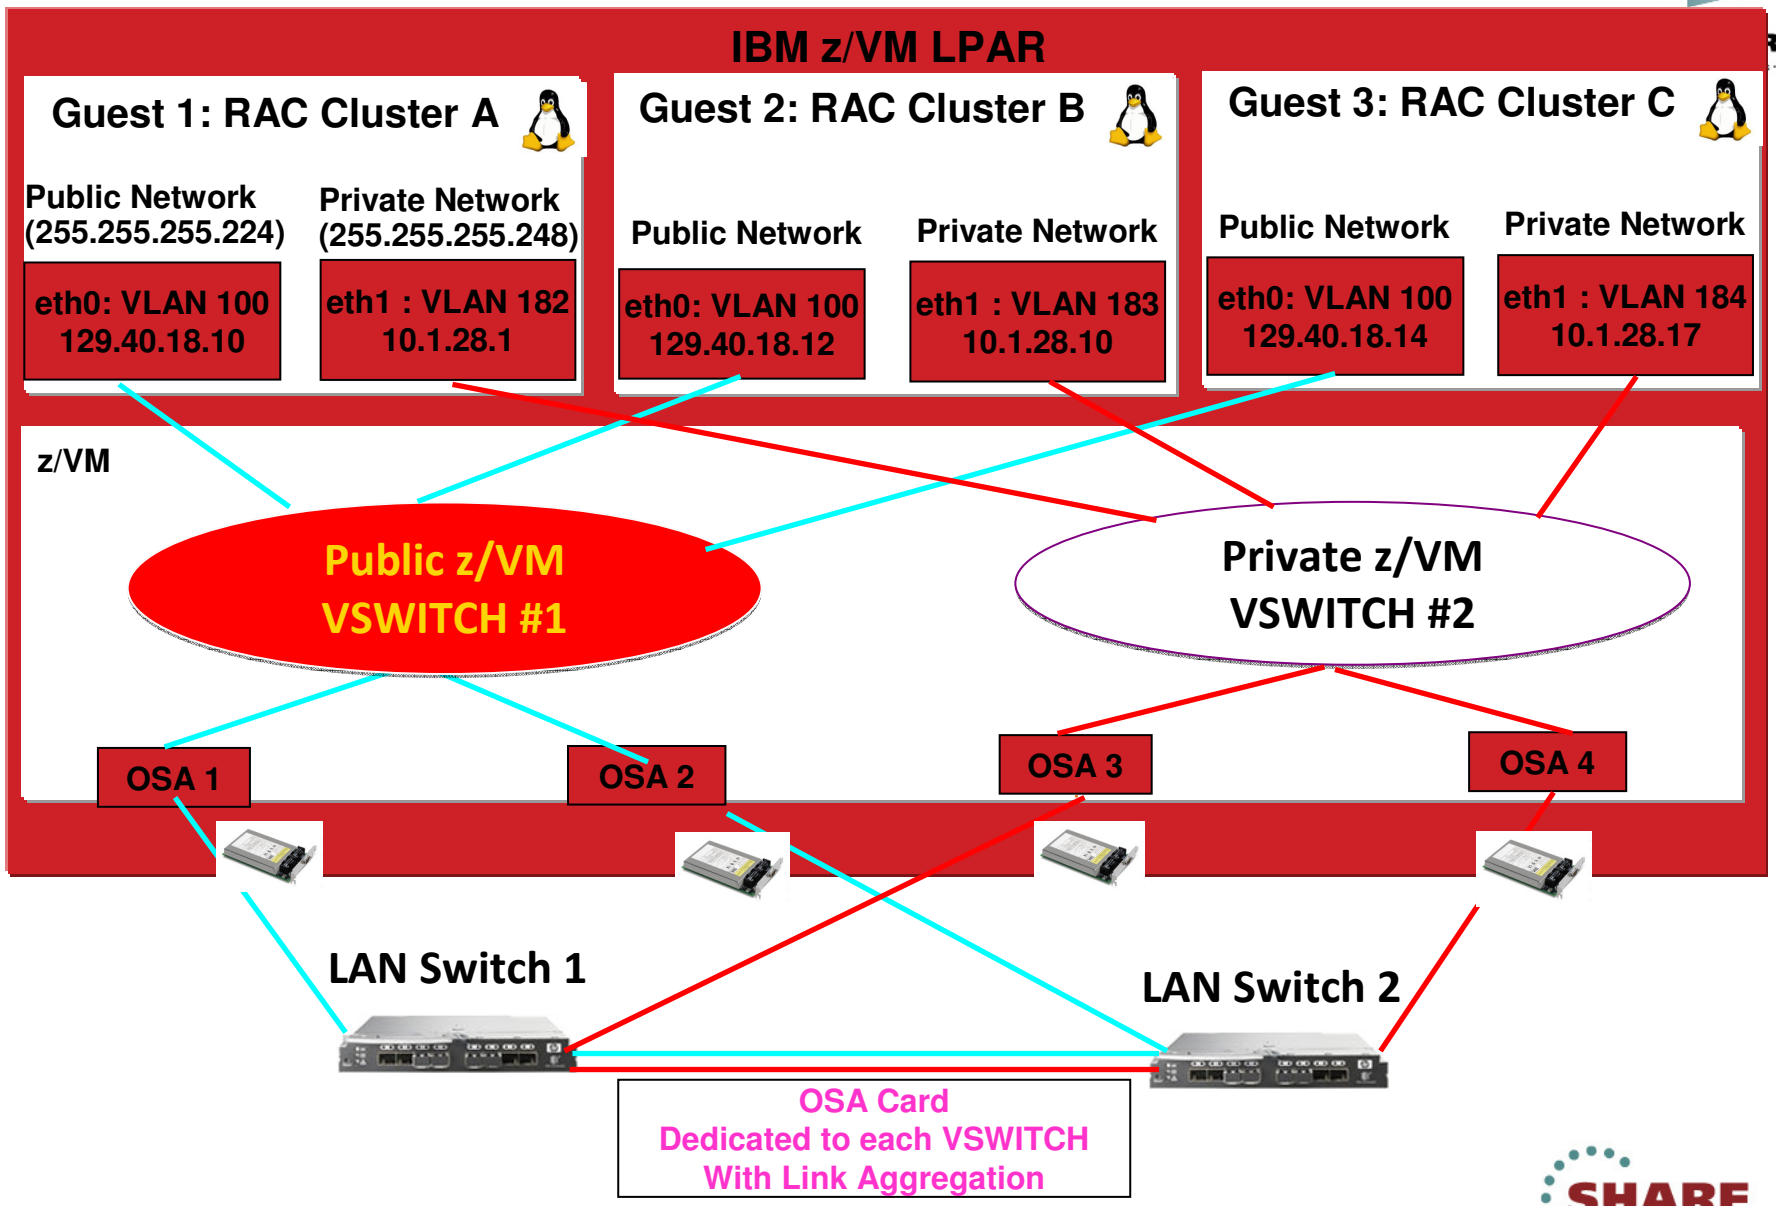

Oracle Interconnect with Linux Bonding

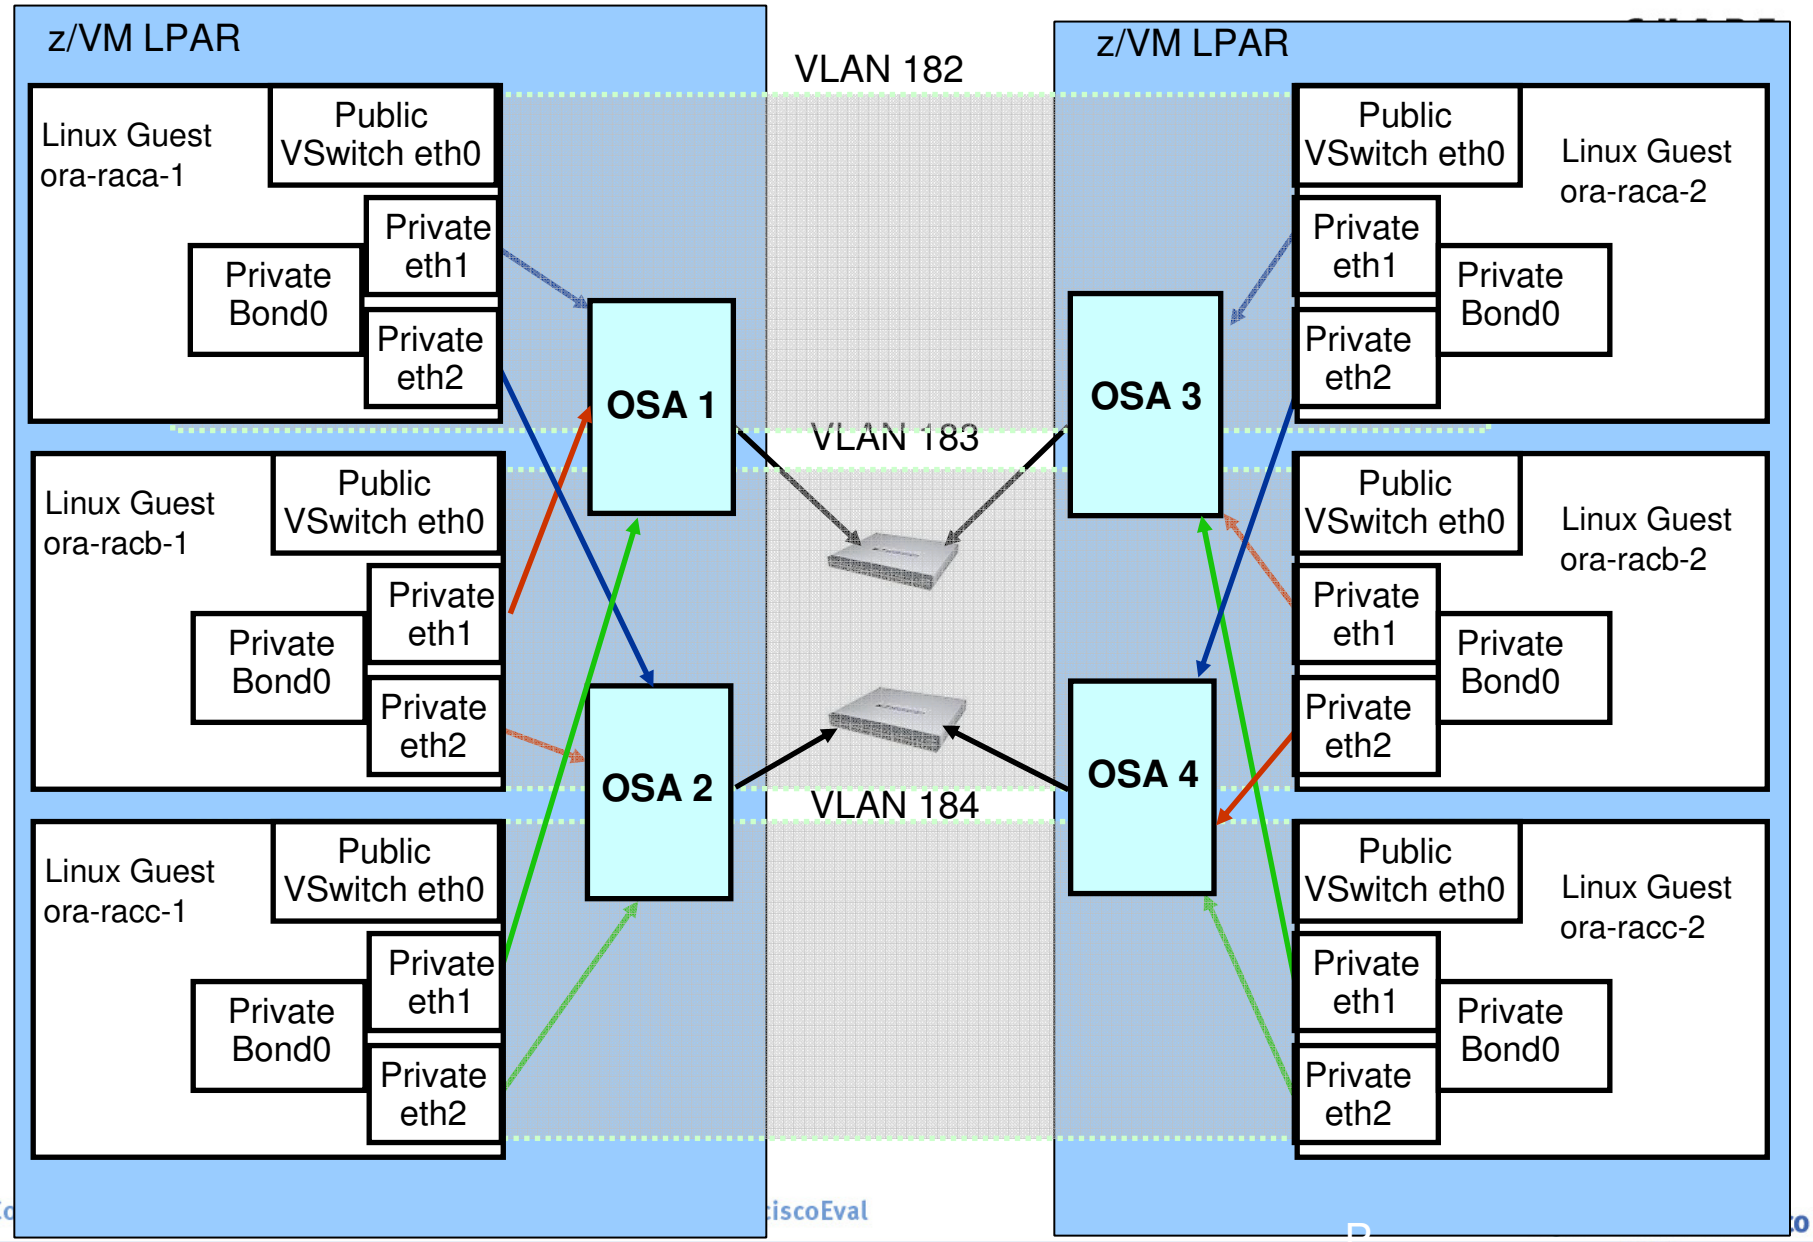

Oracle Interconnect with Oracle's HAIP Redundancy

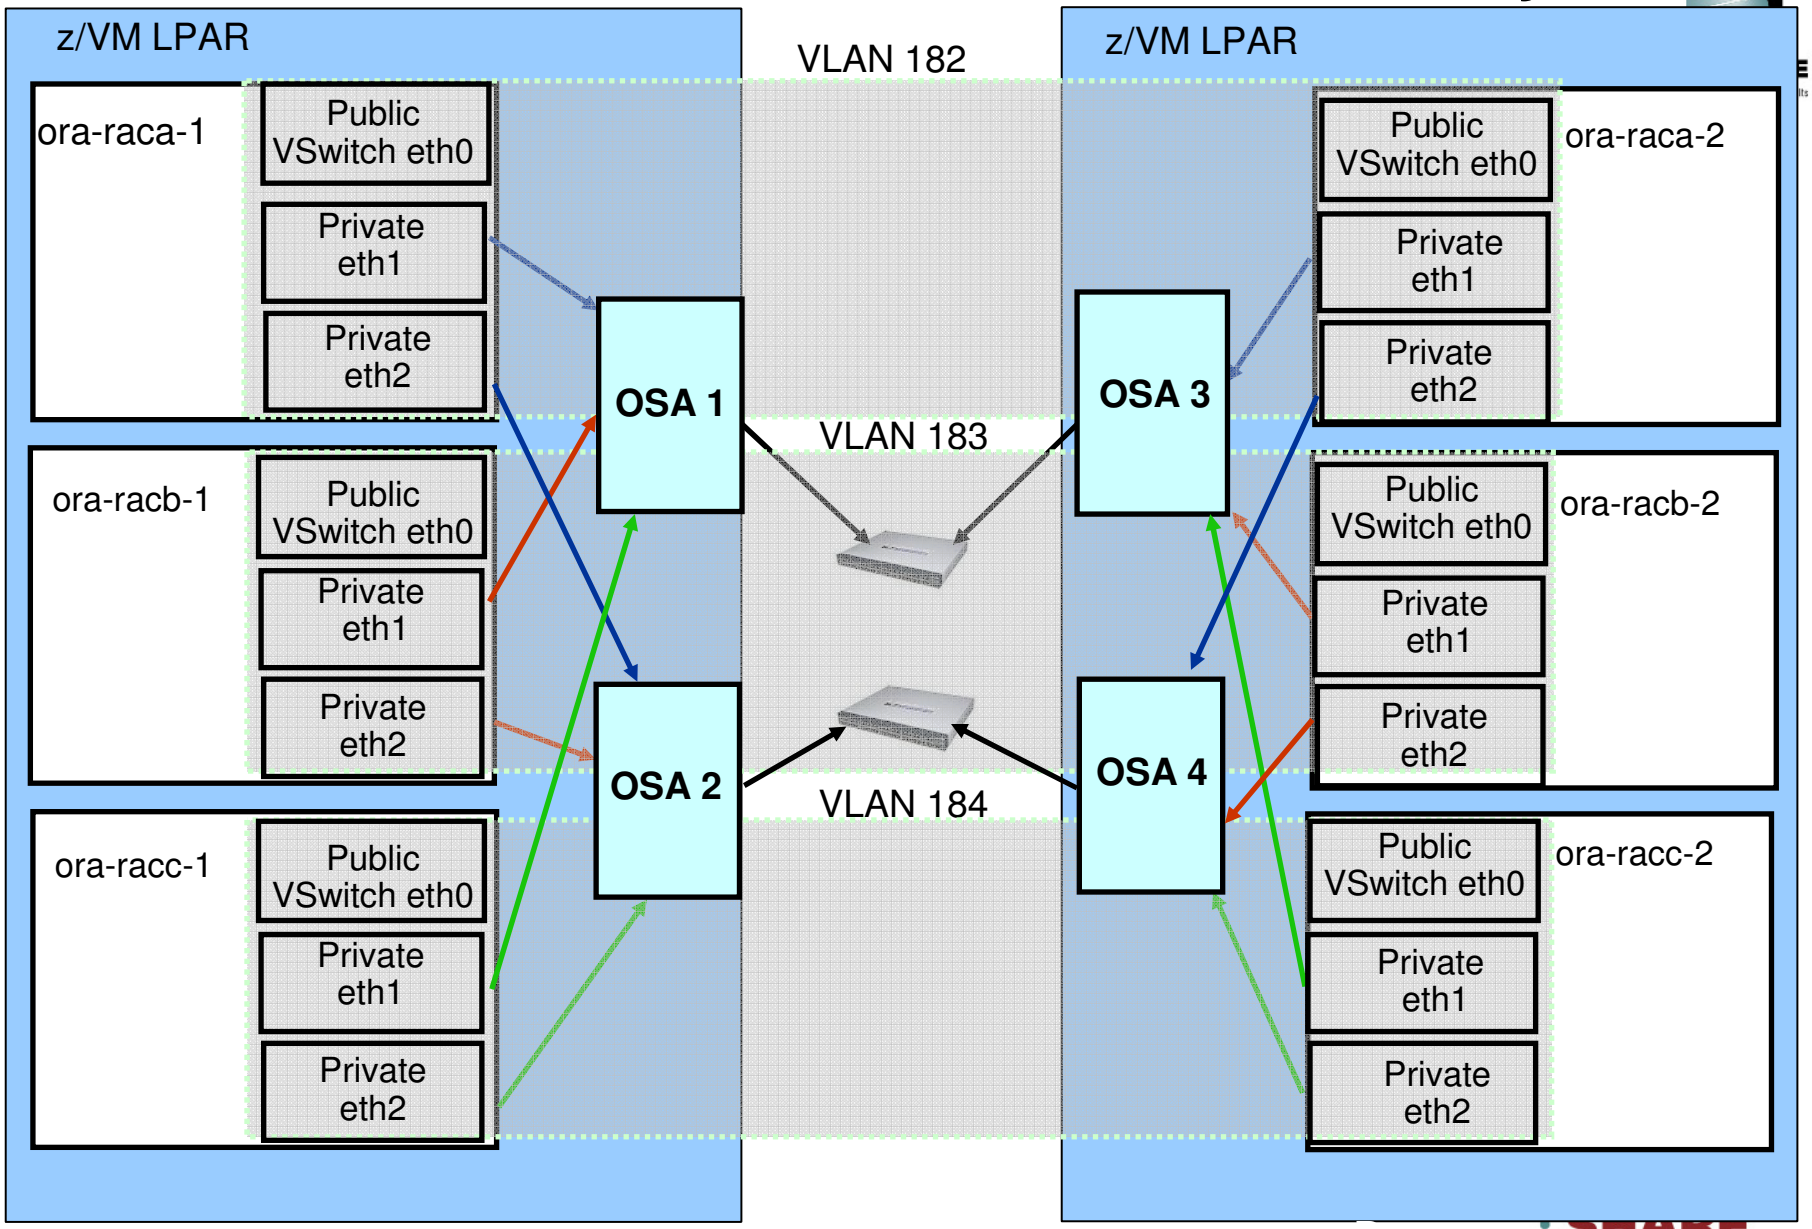

Collecting Data from Oracle AWR Reports

- **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/racc1  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

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 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	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

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	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

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

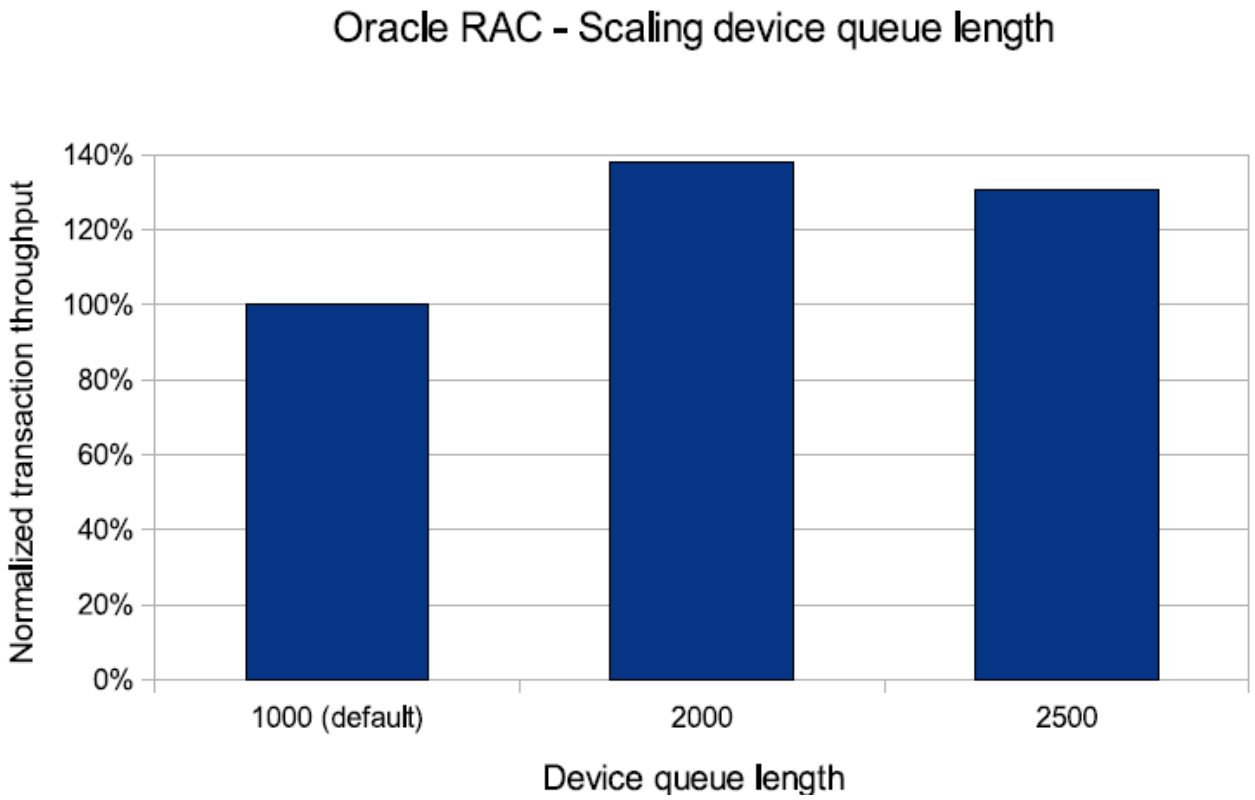

Choose the Correct Network MTU size

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 ARP Required Oracle RAC Installs

- Oracle Grid Install when the network interfaces are set with NOARP you can encounter **BUG – 10173295** when running the root.sh script on the first node.

  Error: Did not succssfully configure and start ASM at **/opt/oracle/11gR2/crs/install/crsconfig_lib.pm line 6470**.

  /opt/oracle/11gR2/perl/bin/perl -I/opt/oracle/11gR2/perl/lib -I/opt/oracle/11gR2/crs/install/opt/oracle/11gR2/crs/install/rootcrs.pl execution failed

  CRS-1013:The OCR location in an ASM disk group is inaccessible. Details in /opt/oracle/11gR2/crs/log/dhsora1/client/clscfg.log

  Oracle Database 11g Clusterware Release 11.2.0.2.0 - Production Copyright 1996, 2010 Oracle. All rights reserved.

  2011-03-16 20:01:53.085: [CLSCFG][53553008]clscfg_main: Configuration type [4]

  ibctx: Failed to read the whole bootblock.

- Update the network interfaces to have ARP enabled (the following is incorrect)
  - **ifconfig –a**
  - eth0 Link encap:Ethernet HWaddr 02:00:02:00:00:A2
    inet addr:130.35.55.234 Bcast:130.35.55.255 Mask:255.255.252.0
    inet6 addr: fe80::200:200:100:a2/64 Scope:Link
    UP BROADCAST RUNNING **NOARP** MULTICAST MTU:1492 Metric:1
    RX packets:5749678 errors:0 dropped:0 overruns:0 frame:0
    TX packets:2799431 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:1414260847 (1.3 GiB) TX bytes:2735238017 (2.5 GiB)

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 Recommended Configurations for System z

Architecture	Oracle Private Network (interconnect)	Oracle Public Network
All z/VM Linux guests in one LPAR	•Private Layer2 VSwitch Guest LAN OSA recommended •Real layer 2 Hipersocket possible •Guest LAN Hipersocket not supported	•Shared Public VSwitch recommended •Shared or dedicated OSA card is possible
z/VM Linux guests on different LPARs	•Real Layer 2 Hipersocket recommended •Private Layer 2 Gigabit OSA card possible	•Shared Public VSwitch recommended •Shared or dedicated OSA card
z/VM Linux guests on different physical machines	•Private Layer 2 Gigabit OSA card recommended with physical switch in between (one hop)	•Dedicated OSA card Possible

Oracle RAC One Node - deployment Omotion

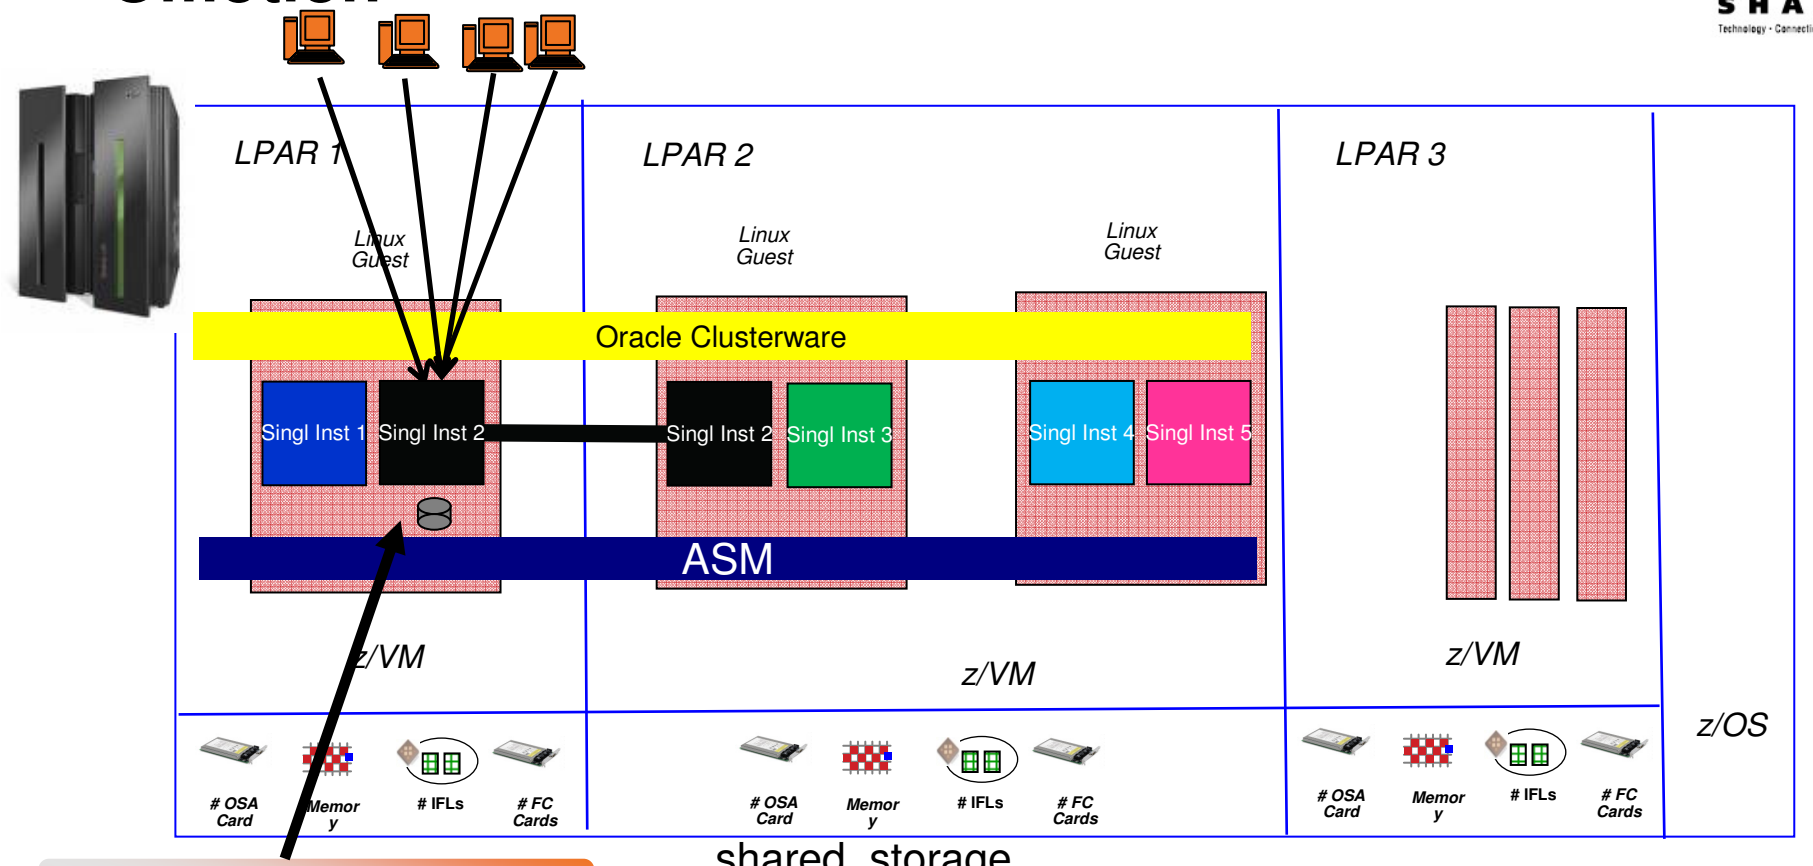

- Patch Oracle binaries, modify Linux parameters, etc..

Groups

shared storage

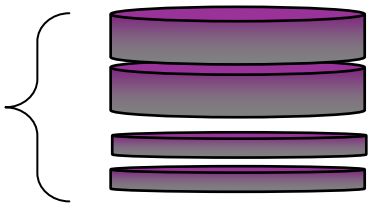

OCR and Voting

Oracle RAC One Node –Omotion

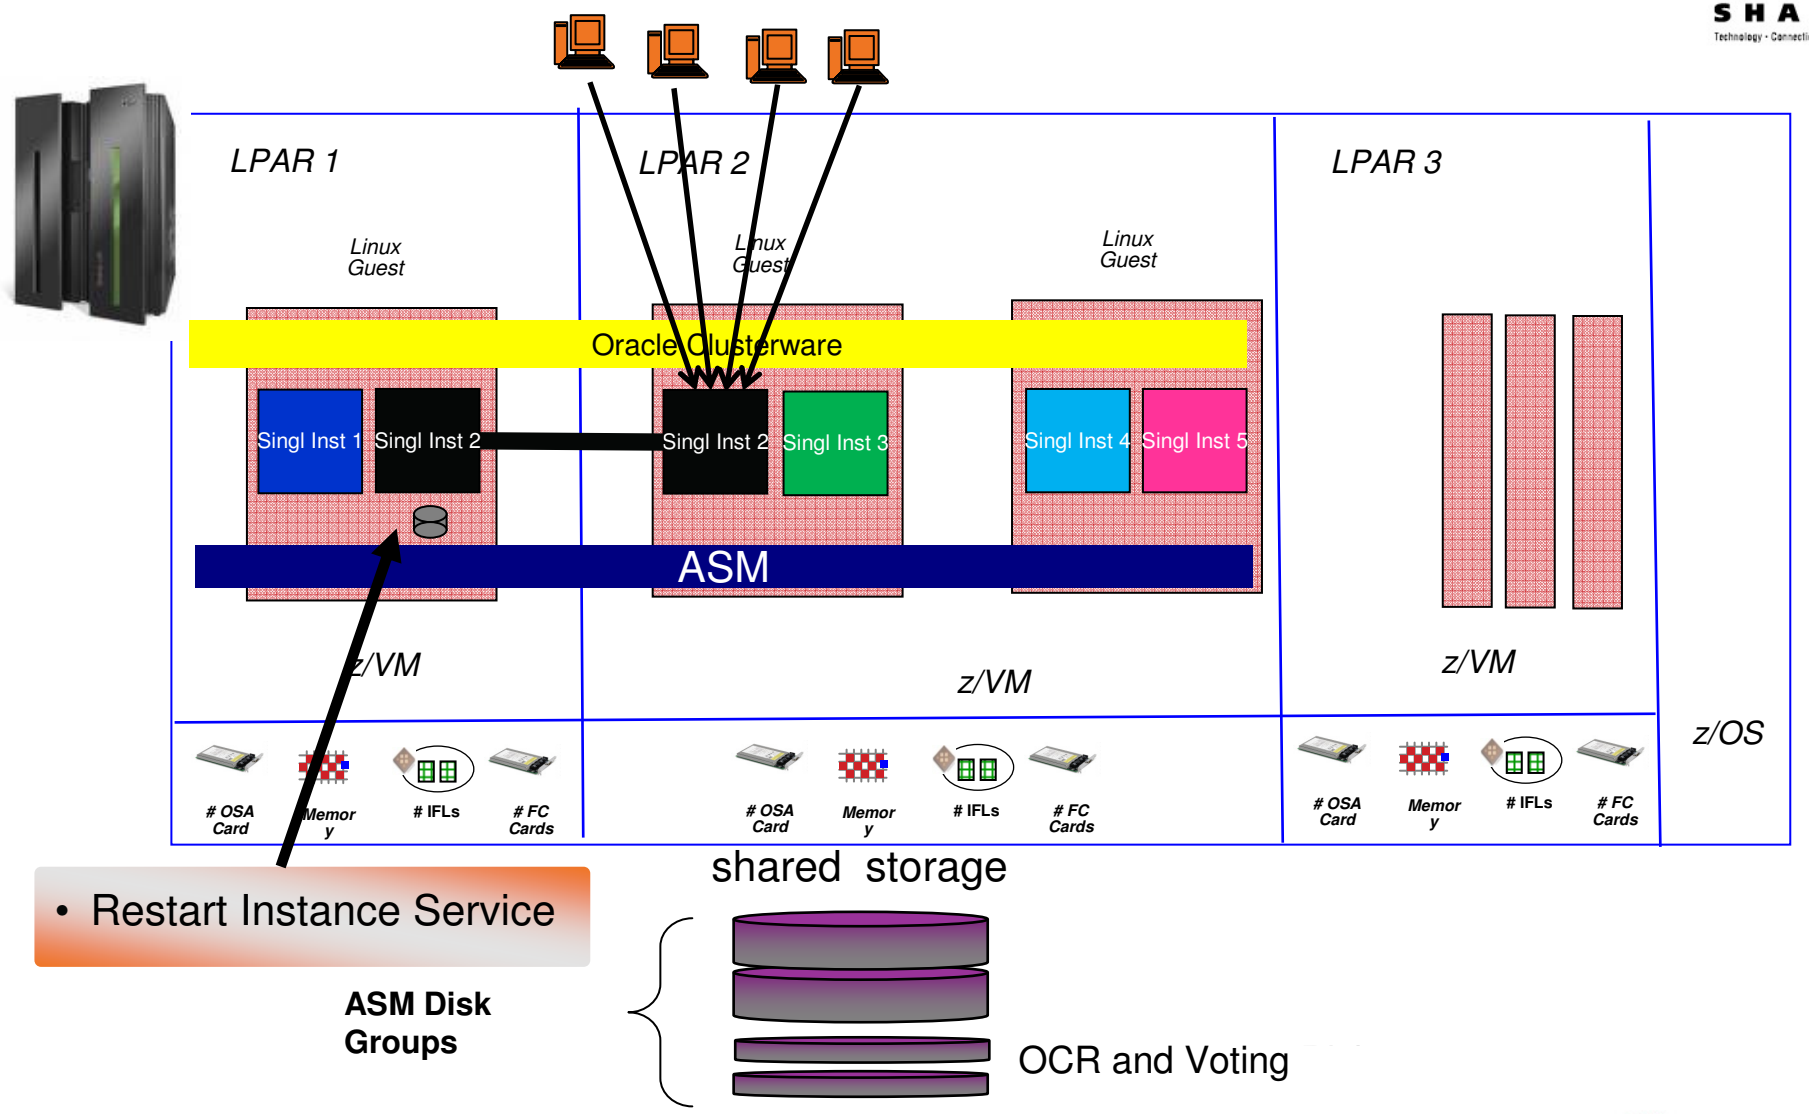

- Restart Instance Service

ASM Disk Groups

RAC ONE – Transparent Application Failover (TAF)

- Oracle Support Note – **453293.1**

- TAF is required to be configured post Install or else failover will not work properly.

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

Suggestion: check that your connections show up as Failover eligible.

IBM z/VM 6.2 Live Guest Relocation

- z/VM Live Guest Relocation Works for Oracle RAC……but is not **yet** Officially Certified to relocate while running….

- You can shutdown one RAC Node – Relocate the Linux Guest then restart – this is Certified.

Oracle Red Hat 6 Certifications – IBM System z

- Grid Agent 12c (now certified)
- WebLogic 12c (now certified)
- Oracle Database 11.2.0.3.3+
  (Supported, Certified 02/2013)

Support of RedHat 6.2+ with Oracle 11gR2 on Linux on System z

- **Now Certified by Oracle!**

- My Oracle Support documentation
  - Linux package rpmchecker is available (used to require an SR)
  - Note 1470834.1 - Requirements for Installing Oracle 11gR2 on RHEL 6 on IBM: Linux on System z (s390x) is available
  - Note 1514012.1 - runcluvfy stage -pre crsinst generates Reference Data is not available for verifying prerequisites on this operating system distribution on Redhat 6 - IBM: Linux on System z

- Start now in a test or development environment

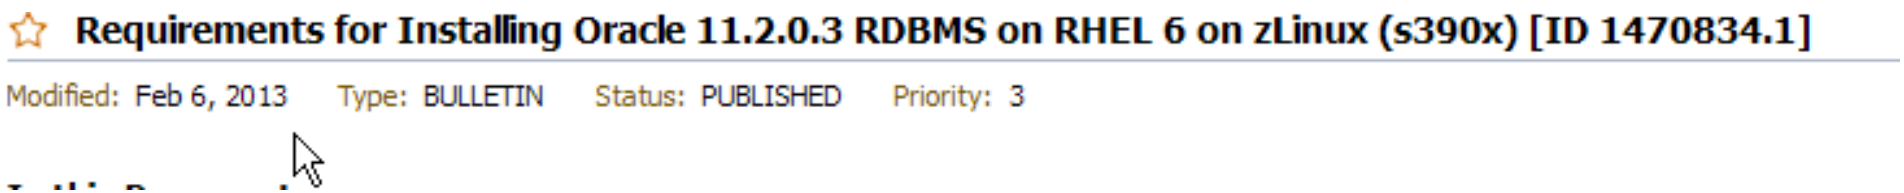

Oracle 12c Grid Monitoring Agent

- New Oracle **12.1.0.2 grid** agent is SuSe 11 & RHEL 6 certified for Linux on System z

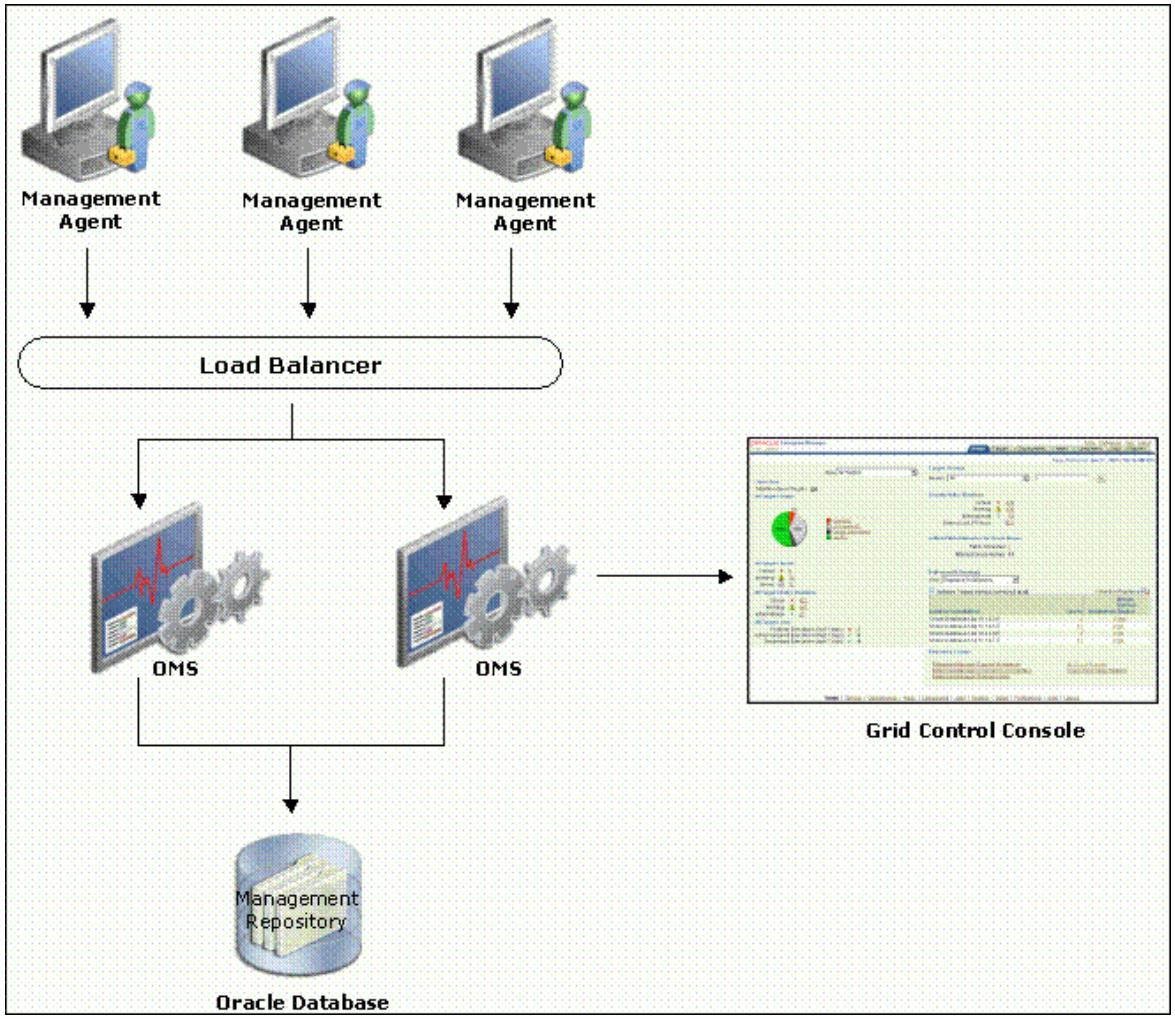

Oracle Cloud Control 12c for Monitoring System z databases – Centralized Deployment

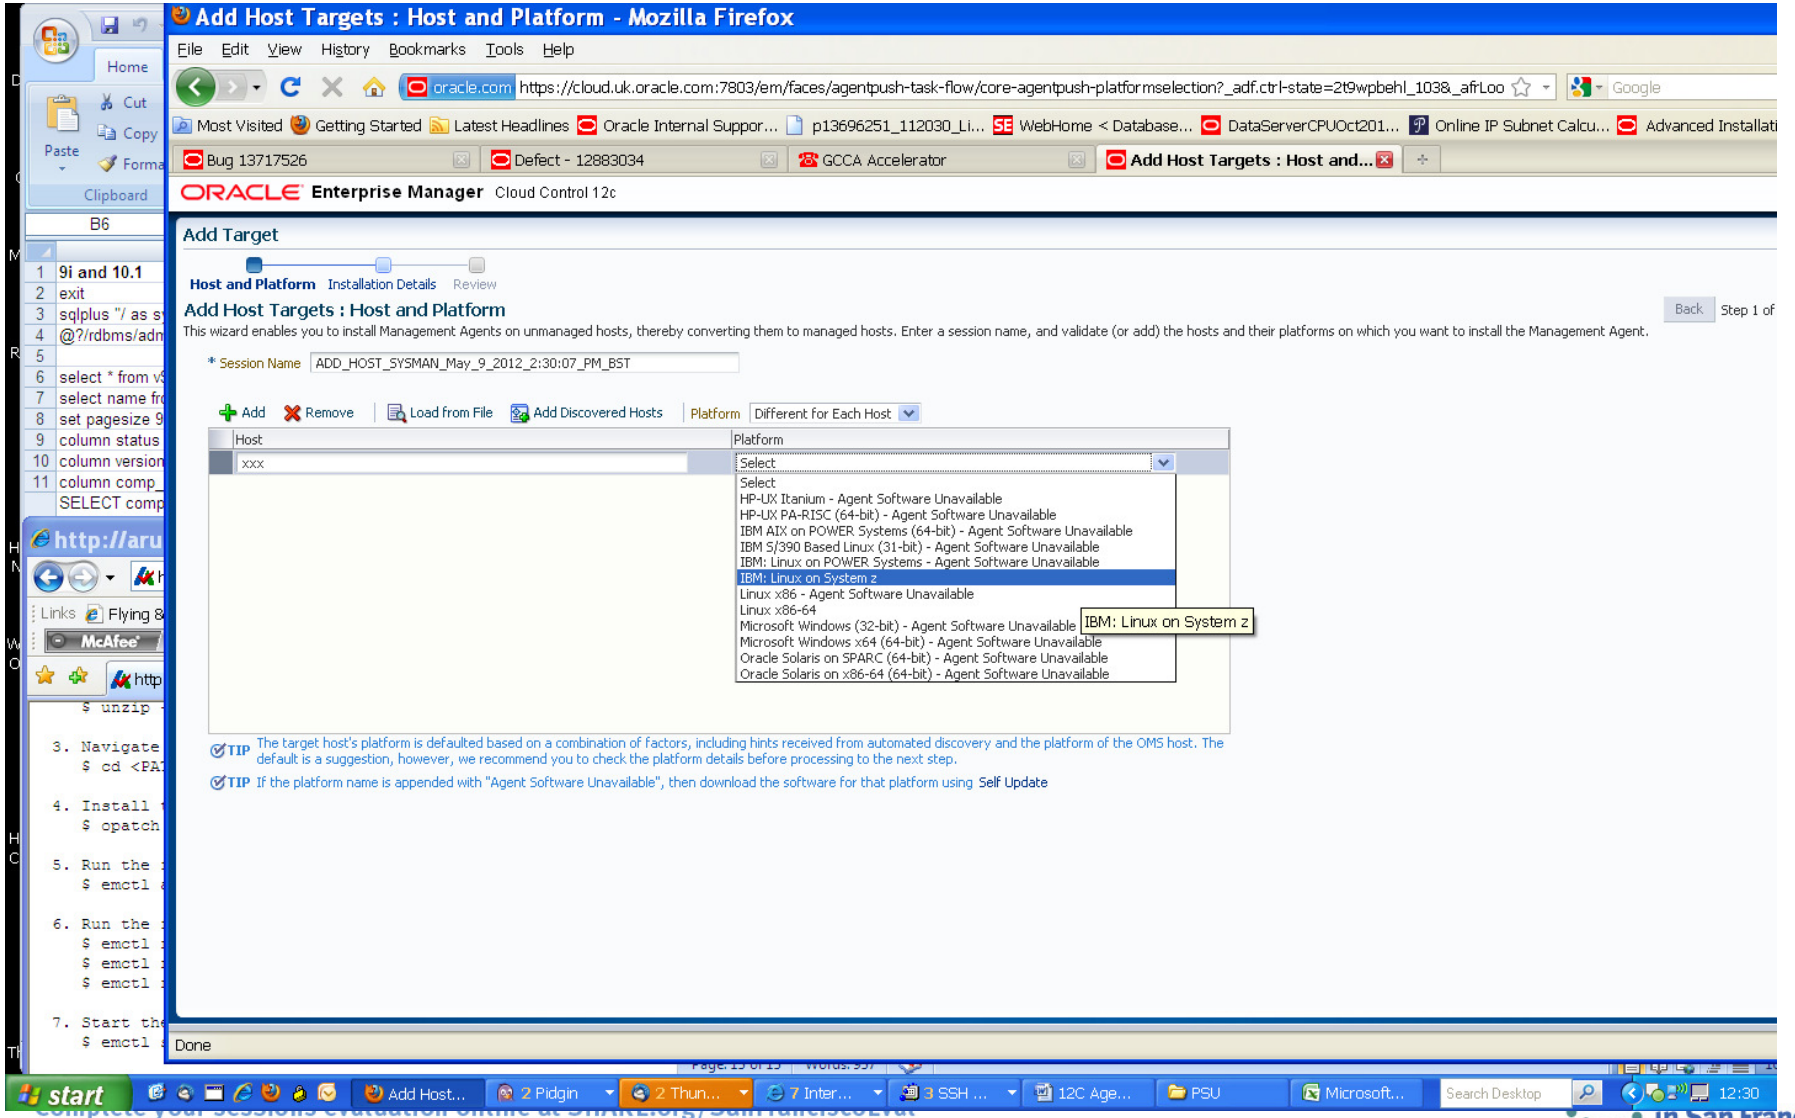

WebLogic Server 12c Release 1 Version 12.1.1x

- Oracle WebLogic Server 12c Certified for SLES 11 SP1+ & Red Hat 6.1+
- Java 7 SDK (64 bit) certified.
- IPv6 Certified
- Download latest Java 7 SDK from IBM Developer works.

				Server Certification		
Installation Type	Version Supported	Processor Type	OS Version	OS 32/64 bit	Oracle FM 32/64 bit	JDK Vendor Version* 12c (12.1.1.0.0)
Oracle WebLogic Server	12c (12.1.1.x)	IBM z/Linux	SLES 11 (SP 1+)	64	64	IBM JDK 7 (SR1+)
Oracle WebLogic Server	12c (12.1.1.x)	IBM z/Linux	Red Hat EL 6 (UL 1+)	64	64	IBM JDK 7 (SR1+)

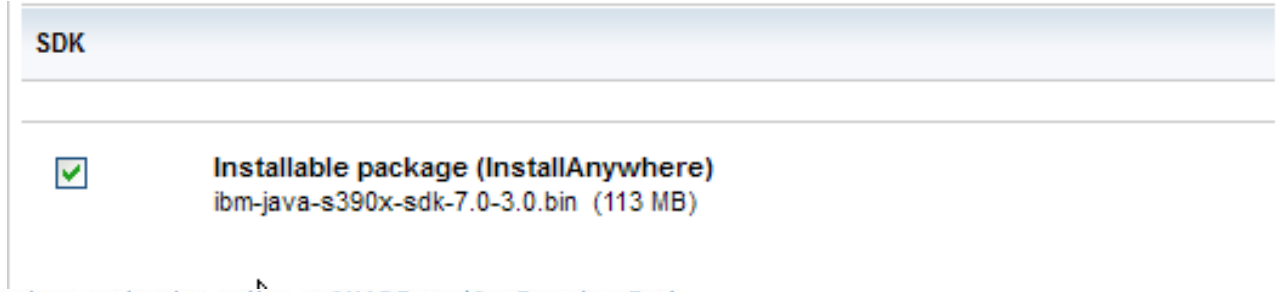

Installing WebLogic Server 12c Release 1

Start Installer with:

export JAVA_HOME=/opt/ibm/java-s390x-70

export PATH=$JAVA_HOME/bin:$PATH

java -Djava.io.tmpdir=/tmp -d64 -jar ./wls1211_generic.jar

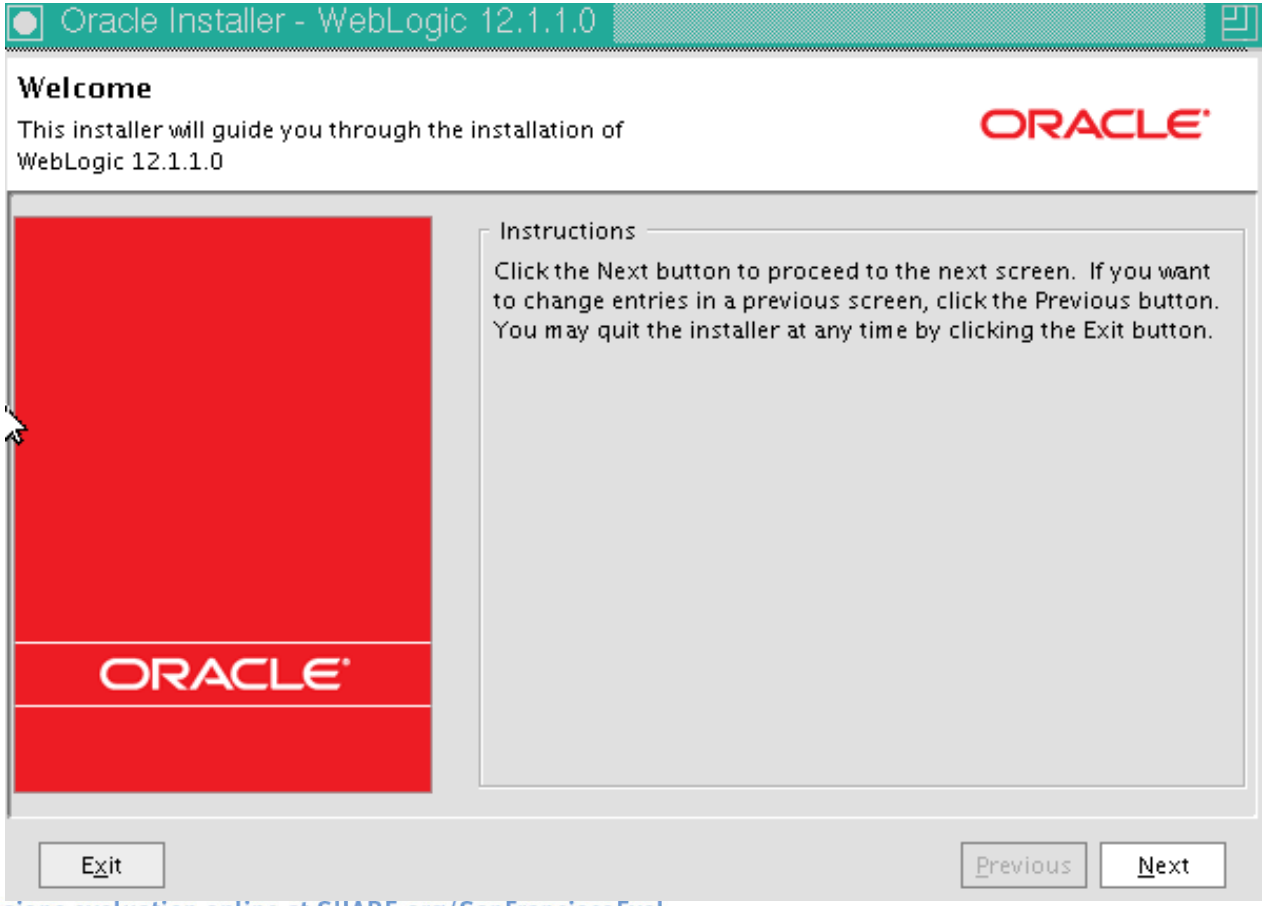

Oracle Database – Upgrades Red Hat 5 -> 6

- Upgrade sequence will be to upgrade 10gR2 to 11.2.0.3 on RedHat 5.x
- Move the database to a clean install of RH 6.2+ (Oracle services have changed from 10gR2->11gR2 and Red Hat 5 vs Red Hat 6)
- RAC Customers can do one node at a time.
- Single Instance - fresh Red Hat 6 install, pre- Install Oracle ASM & DB Code – with the latest Patchsets, then run deconfigure (old) and configure on the new Red Hat 6 system.

  Oracle Note: 1434351.1 - Alternative Way To Upgrade An ASM Standalone Configuration From Release to release 11.2.0.3.

Red Hat Oracle Grid Installs – iptables

[Problem] Oracle Grid Install on Red Hat– cluvfy for the cluster, OR during the install / adding nodes copying the Oracle Grid Home to another system hangs/fails.

ERROR: PRVF-7617 : Node connectivity between "xxxxxxx : 10.0.0.1" and "yyyyyyy : 10.0.0.2" failed

Result: TCP connectivity check failed for subnet "10.0.0.0"

- **[Solution] The problem may be attributed to the firewall/iptables**

chkconfig --list iptables
iptables 0:off 1:off 2:on 3:on 4:on 5:on 6:off

or you can check the service with

service iptables status
Firewall is stopped.

service iptables stop
chkconfig iptables off

Red Hat & SuSe Disk & Memory Configuration

- HugePages
- FCP / SCSI Storage
- DASD Storage

How to Setup Linux Large Pages

- Obtain Huge Pagesize from '**cat /proc/meminfo**' for System z Linux (SLES 11, Red Hat 6 – 1mb, SLES 10, RHEL5 – 2mb page size)

- Setup the oracle user to be able to use hugepages in /etc/security/limits.conf

  /etc/security/limits.conf ** example of 70GB Linux Guest

  *** soft memlock 75161927680**

  *** hard memlock 75161927680**

- Set the following kernel parameters in /etc/sysctl.conf
  - vm.hugetlb_shm_group=`id -g oracle`
  - vm.nr_hugepages= <# of Huge pages needed for Oracle SGA>

- In 11.2.0.3 Huge pages are used automatically when set correctly but sometimes requires a reboot to get all the HugePages contiguous pages (may want to use oracle parameter – use_large_pages='only')

multipath.conf SCSI/FCP Device Persistence

```
defaults {
      user_friendly_names yes
rr_min_io=15
dev_loss_tmo 90
fast_io_fail_tmo 5
      }
blacklist {
       devnode "^(dasd)[0-9]*"
}

multipaths {
      multipath {
            wwid 36005076303ffcbbf000000000000ef00
            alias    ASMFCP1
            path_grouping_policy    multibus
}
}
```

- SuSe 10 -> 11 or Red Hat 5 -> 6 may want to revisit settings (failover -> multibus)
- rr_min_io is storage array specific
- fast_io_fail_tmo – Red Hat 6 (length of time to wait before failing)
- Use the /dev/mapper/<alias name> for ASM Diskstring
- If Using ASMLib you can use a disk string '/dev/oracleasm/disks/*
- Follow Oracle Note:1377392.1 for udev rule device permission example.

ECKD / DASD Requires a Partition

- **ECKD/DASD Disks it's mandatory to create at least one partition**
- **example below shows an Oracle session's memory region being**
- **Linux was not expecting this and wrote "e5e5e5e5e5" to sector 0 (as Linux assumes disks are partitioned)**
- **Moral of the story: Always Partition DASD disk for LVM, ASM and swap**

```
147a1d000:  ffffffffffffffff ffffffffffffffff  ................
147a1d010:  ffffffffffffffff ffffffffffffffff  ................
147a1d020:  ffffffffffffffff ffffffffffffffff  ................
147a1d030:  ffffffffffffffff 00b38f0000000081  ................
147a1d040:  0000000205d5fdb0 00000000864d3a3c  ............M:<
147a1d050:  0e575248245f5347 41535441545f5503  .WRH$_SGASTAT_U.
147a1d060:  535953015404504f 4f4c044e414d4504  SYS.T.POOL.NAME.
147a1d070:  4442494407534e41 505f49440f494e53  DBID.SNAP_ID.INS
147a1d080:  54414e43455f4e55 4d424552e5e5e5e5  TANCE_NUMBER....
147a1d090:  e5e5e5e5e5e5e5e5 e5e5e5e5e5e5e5e5  ................
...
147a1dff0:  e5e5e5e5e5e5e5e5 e5e5e5e5e5e5e5e5  ................
```

DASD Storage (use HyperPAV)

- HyperPAV - Available for SLES 11 & Red 5.9 & 6.1
- For Oracle need to create a UDEV rule for ASM disk
- Oracle 11.2.0.3 RAC Installs may hang at 75% while doing a disk check (see Oracle Note: **1459030.1 – exectask patch)**

Create a **/etc/udev/rules.d/99-udev-oracle.rules** file to assign permissions for DASD devices.

```
vi /etc/udev/rules.d/99-udev-oracle.rules
```

Result:

```
KERNEL=="dasd*1",ID=="0.0.0300",OWNER="grid",GROUP="oinstall",MODE="0660",SYMLINK+="ASM0300"
KERNEL=="dasd*1",ID=="0.0.0305",OWNER="grid",GROUP="oinstall",MODE="0660",SYMLINK+="ASM0305"
```

Make an entry for each device you plan to use with Oracle ASM.

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

- **Collaboration:**
  IBM / Oracle / Velocity Software
- **Due out - Q1 - 2013**

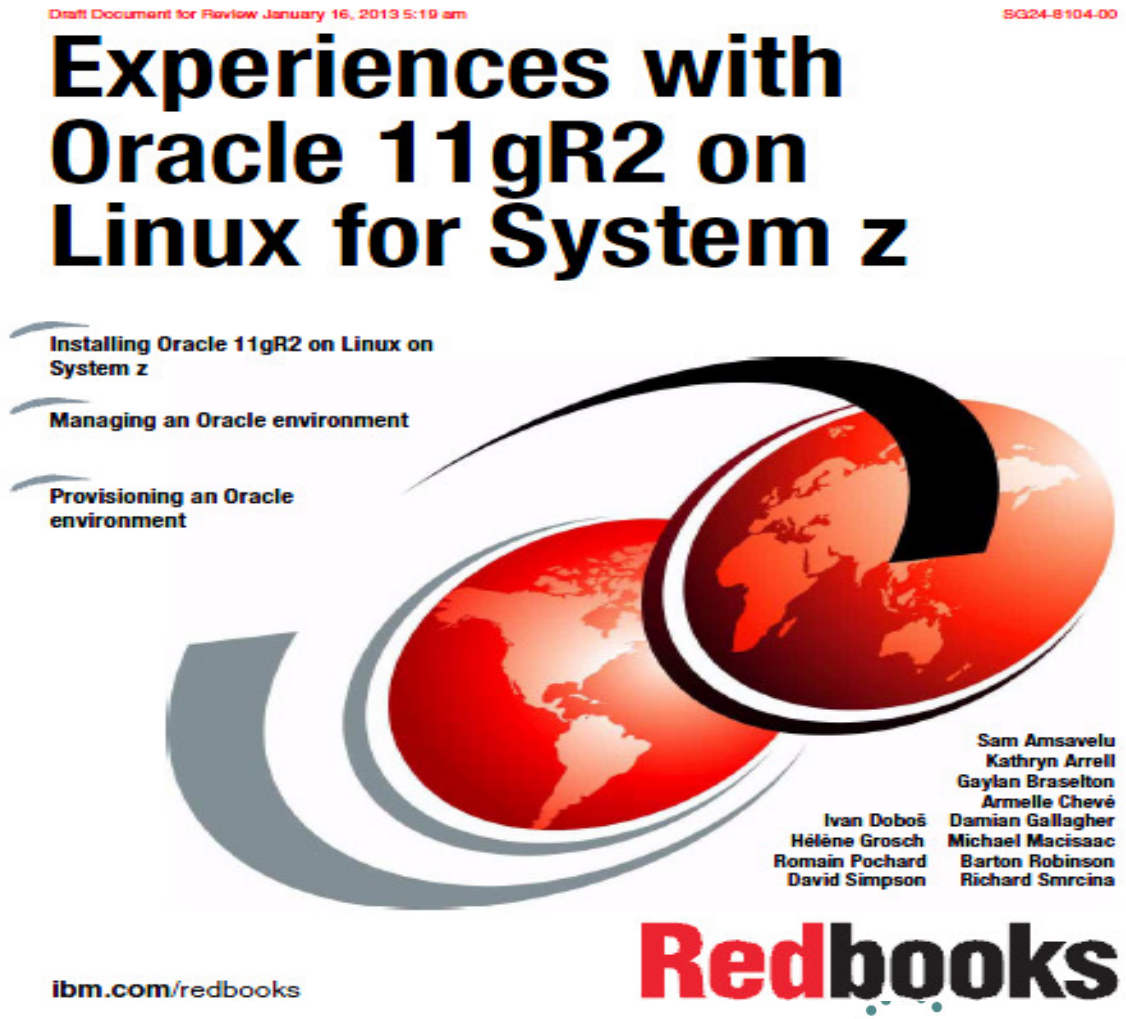

References - Key Oracle & IBM Whitepapers:

- **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

System z Social Media Channels

- Top Facebook pages related to System z:
  - IBM System z
  - IBM Academic Initiative System z
  - IBM Master the Mainframe Contest
  - IBM Destination z
  - Millennial Mainframer
  - IBM Smarter Computing
- Top LinkedIn groups related to System z:
  - System z Advocates
  - SAP on System z
  - IBM Mainframe- Unofficial Group
  - IBM System z Events
  - Mainframe Experts Network
  - System z Linux
  - Enterprise Systems
  - Mainframe Security Gurus
- Twitter profiles related to System z:
  - IBM System z
  - IBM System z Events
  - Millennial Mainframer
  - Destination z
  - IBM Smarter Computing
- YouTube accounts related to System z:
  - IBM System z
  - Destination z
  - IBM Smarter Computing
- Top System z blogs to check out:
  - Mainframe Insights
  - Smarter Computing
  - Millennial Mainframer
  - Mainframe & Hybrid Computing
  - The Mainframe Blog
  - Mainframe Watch Belgium
  - Mainframe Update
  - Enterprise Systems Media Blog
  - Dancing Dinosaur
  - IBM Destination z

#SHAREorg

SHARE
Technology · Connections · Results

Oracle RAC Networking Alternatives on Linux on System z and Red Hat 6 Oracle DB Support

Speaker Names: David Simpson & Kathryn Arrell
Speakers Company: IBM

Date of Presentation: **Wednesday, February 6, 2013 (9:30am)**
Franciscan D, Ballroom Level

Session Number: **12758**
Twitter -> @IBMandOracle

Insert Custom Session QR if Desired.

SHARE in San Francisco 2013