

**The Mentora Practical Solutions Series**

Mentora “Practical Solutions” papers provide insight and concrete suggestions to address technical issues of installing, maintaining and supporting high-availability web-enabled production applications, developed from our *hands-on real-life experience* providing 24/7 uptime for our managed application hosting customers.

Our goal is to help other IT operations teams by sharing lessons learned. Our focus is on practical solutions from an operations perspective – things you can do today, not things that “should be” or that require product enhancements by third-party vendors. We share “what works to make it work”.

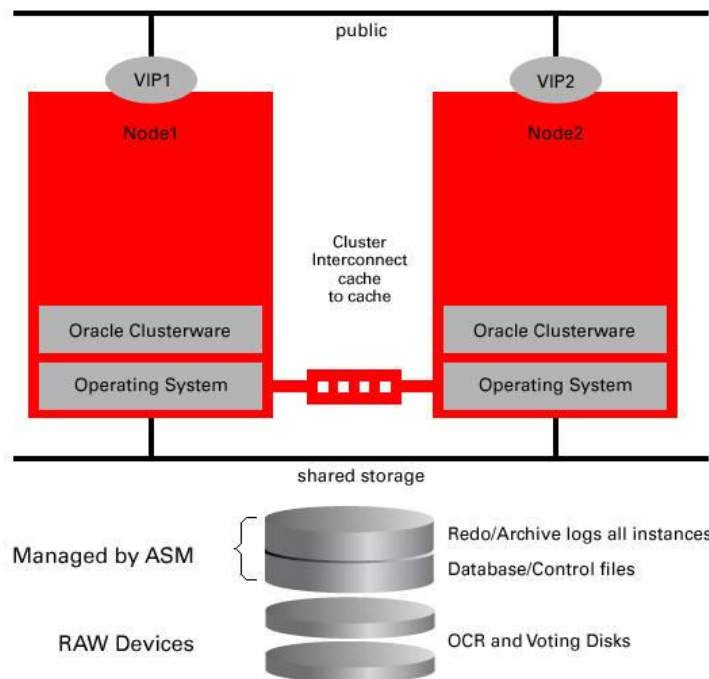
**Practical Solutions: High Availability with Oracle RAC and RedHat 5**

The Challenge

Our customer, an international provider of 24/7 on-line business-to-business data and commercial interchange, needed to expand their production environment in November 2007 by implementing a new Software as a Service application.

To meet redundancy, technology and cost constraints the requirements for the environment included:

- RedHat 5 Linux
- Oracle RAC
- SAN (a NetApp 3020)



Ref: <http://www.oracle.com/technology/pub/images/niemec-racsig-f12.jpg>



## Lessons Learned

1. Oracle RAC installation is currently unproven in the marketplace on RedHat Version 5 and requires workarounds
2. Installing RAC in a SAN environment with RH5 Standard Edition takes some "tweaks"

## **Technical Issues and Solutions**

**Issue:** After installing RedHat 5 (seamlessly, smoother than RH4!), Oracle Clusterware installation failed.

### **Practical Solutions:**

Almost all of the installers had to be edited to get Oracle to recognize RH5.0, and a number of important scripts had to be modified to comment out some RH4-specific information. These also had to be reviewed/re-done after all the patches were applied.

### **Modifications to /install/oraparam.ini for database software installation**

#### **Find this:**

[Certified Versions]

Linux=redhat-3,SuSE-9,redhat-4,UnitedLinux-1.0,asianux-1,asianux-2

#### **Replace with:**

[Certified Versions]

Linux=redhat-3,SuSE-9,redhat-4,UnitedLinux-1.0,asianux-1,asianux-2,redhat-5

### **Modifications to /install/oraparam.ini for Clusterware software installation**

#### **Find this:**

[Certified Versions]

Linux=redhat-3,SuSE-9,redhat-4,UnitedLinux-1.0,asianux-1,asianux-2

#### **Replace with:**

[Certified Versions]

Linux=redhat-3,SuSE-9,redhat-4,UnitedLinux-1.0,asianux-1,asianux-2,redhat-5

**Note: If you want to bypass ALL system prereq checks, simply invoke the installer as ./runInstaller -ignoreSysPrereqs**

### **Modifications to \$ORA\_CRS\_HOME/bin/vipca:**

#### **Find this:**

```
if [ "$arch" = "i686" -o "$arch" = "ia64" -o "$arch" = "x86_64" ]
then
  LD_ASSUME_KERNEL=2.4.19
  export LD_ASSUME_KERNEL
fi
```

#### **Add this line:**

```
unset LD_ASSUME_KERNEL <<== Line to be added
```



## Modifications to \$ORA\_CRS\_HOME/bin/srvctl:

### Find this:

```
LD_ASSUME_KERNEL=2.4.19
export LD_ASSUME_KERNEL
```

### Add this line:

```
unset LD_ASSUME_KERNEL <<== Line to be added
```

## Related Known Issues (Oracle Metalink):

Metalink note 414163.1: 10gR2 RAC Install issues on Oracle EL5 or RHEL5 or SLES10 (VIPCA Failures)

**Issue:** OCFS is not included in Oracle Standard Edition which limits the install to using Raw Partitions. Raw partitions worked "as installed" with local disk but not with remote disk on the NetApp SAN.

### Practical Solution:

Create local raw devices first, then copy them to the SAN.

- create images of equal size to the LUN via dd command
- bind raw device to the images (voting and OCR)
- initialize OCR and Voting disk on those via root.sh
- copy them to NetApp partitions using dd (once again, need to be sure they're exact size)
- bind raw devices on second node to SAN LUNs
- run root.sh on the second node

**Issue:** Redundant fibre paths had to be established from the SAN to the hosts, absent Powerpath or VXDMP. By default RH Dynamic Multipathing uses '/dev/sd\*' nomenclature, dynamically assigning drive names upon each reboot. Oracle RAC expects each logical drive to have a consistent device name, and fails if names are changed upon restart.

### Practical Solutions:

- Do not accept the default "user-friendly" device names ("...sd1,...sd2, ...sd 3,...), instead specify use of the "long device name" in the Linux Multipathing configuration.

```
/etc/multipath.conf
defaults {
    user_friendly_names
    no
}
```

- Exclude non-multipath device names in Oracle asm.config

```
/etc/sysconfig/oracleasm
ORACLEASM_SCANORDER=dm
ORACLEASM_SCANEXCLUDE=sd
```



**Issue:** Our installations were all started from physical node 2, which the cluster automatically treated as "node 1". After ASM instance and disk groups were configured, bringing up a RAC db resulted in the instance #'s not matching the node #'s (i.e. instance 1 was on node2, instance 2 was on node 1).

**Practical Solutions:**

- To correct naming mismatch (visual/logical), create a single instance RAC DB on Node 1, add the next instance to Node 2
- This should be avoided, as this will continue haunting the install down the line with additional databases not being recognized by the cluster (see Metalink note 358124.1: Getting CRS-0215 When Starting Up Database Using svrctl)

**Issue:** RMAN would delete the archive log directory preventing future backups from happening

**Practical Solution:**

A fake directory had to be created for RMAN not to delete the archive destination directory, since RMAN wouldn't delete it if it thought the directory was not empty.

See Metalink note 383784.1: Archive\_log Directory On Asm Being Deleted By Rman Archivelog Backup

---

**ABOUT Mentora [www.mentora.com](http://www.mentora.com)**

Mentora is a pioneer in the new generation of managed application hosting providers, combining support of Linux and the newest technologies, system performance testing and a complete menu of managed services for networks, servers, operating systems, databases, and storage, with our exclusive, personalized **named-engineer-on-call** and **Application Availability SLA**. Headquartered in Atlanta, Mentora has offices in Boston and Washington DC, and delivers services nationwide.