

# Oracle9i Application Server Infrastructure: Improved Availability with Red Hat Cluster Manager on Red Hat Linux

*An Oracle White Paper  
November 2002*

## OVERVIEW

This paper is to be used in conjunction with the *Oracle9i Application Server Infrastructure: Improved Availability with Hardware Clusters* paper that describes a generic solution to make the Oracle9iAS Infrastructure more highly available by enabling fail-over capabilities using the hardware vendor's clustering software. This paper provides the necessary steps for this solution on Red Hat Linux using Red Hat Cluster Manager (RHCM) as the clustering software. A separate tar-file, **9iAS\_Linux\_RHCM\_Scripts.tar**, containing example scripts to implement this solution is included with this paper. Note that this solution is for the Oracle9iAS Infrastructure, only, and not for the Oracle9iAS Mid-tier (Portal, Wireless, etc.).

## ASSUMPTIONS

This paper assumes a good working knowledge of Red Hat Linux (Intel based), RHCM, clustering concepts, and Oracle 9i Application Server.

## LIMITATIONS

The following limitations should be considered before implementing this solution.

- 1) Implementation of this solution requires a dedicated two-node RHCM cluster for this solution, only. No other external software pieces will be supported within the configuration detailed in this solution.
- 2) Upon implementation of this solution, the active runtime of one of the two nodes of the cluster will be with the node's hostname set to a virtual hostname and not to the original hostname registered for the cluster. This will result in unpredictable behavior of any installed software that has a dependency on the node's hostname.

## TESTED ENVIRONMENT

- 2 Intel-based servers running as a two-node Red Hat Linux RHCM cluster
- Red Hat Linux Advanced Server 2.1AS (kernel-2.4.9-e.3, glibc-2.2.4-26)
- Red Hat Cluster Manager 1.0.16-7 (clumanager-1.0.16-7)
- 2GB memory on both the cluster nodes
- 18.2 GB shared storage (Dell Power Vault 220S)
- 2940U2W Adaptec SCSI controller (on both the cluster nodes for connection to the shared storage)
- 2 LAN adapters on each of the cluster nodes (1 private and 1 public interface)
- Virtual hostname
- Logical IP address
- Red Hat Linux Advanced Server 2.1AS kernel parameter set to the

following values

- **semnmi** (100), **semmns** (256), **semopm** (100), **semmsl** (100), **shmmax** (2GB), **shmmmin** (1), **shmmni** (100), **shmmseg** (4096), and **shmvmx** (32767)
- Required Red Hat Linux Advanced Server 2.1AS patches for Oracle9iAS 9.0.2.0.1 for Intel-Linux
- Oracle9iAS 9.0.2.0.1 for Intel-Linux

Note: ensure that all SCSI ID's on both cluster nodes and the shared storage are unique and that there is proper SCSI termination on the shared storage

### **EXAMPLES FROM TESTED ENVIRONMENT (USED THROUGHOUT THIS PAPER)**

- Two Intel-based servers, **oracleias1** and **oracleias2**, configured as a two-node RHCM cluster
- Virtual hostname, **oracleias**
- Logical IP address **130.35.149.193** (mapped to the Virtual hostname, **oracleias**)
- Filesystem **/dev/sdd3** (on the shared storage where the Oracle9iAS Infrastructure will be installed)
- Empty **/oracle** directory (where **/dev/sdd3** will be mounted to)

### **SETUP FOR THE ORACLE9IAS INFRASTRUCTURE INSTALL**

From this point onwards in this paper, the two-node Red Hat Linux RHCM cluster environment mentioned above is assumed to be configured and running. All examples will be with respect to this environment.

Log in, as **root**, to one node (**oracleias1**, in this example) and perform the following steps.

- mount the filesystem on the shared storage
  - **# mount /dev/sdd3 /oracle**
- add the logical IP address to the primary network adaptor
  - **# /sbin/ifconfig eth0:1 netmask 255.255.252.0 130.35.149.193**

Note: **255.255.252.0** is an example **netmask** used in the **TESTED**

**ENVIRONMENT**; it should be changed accordingly for your network environment

- change the hostname to the virtual hostname
  - **# hostname oracleias**

## ORACLE9IAS INFRASTRUCTURE INSTALL

- as **root**, on the node where the **SETUP** steps were performed (**oracleias1**, in this example), remove any **iasdb** entries from **/etc/oratab** and completely remove **/etc/emtab**
- as a **non-root** user, on the same node, install Oracle9iAS Infrastructure into **/oracle** (Note: refer to the *Oracle9iAS for Intel-Linux Install Guide* and *Release Note* for this step)
- perform all required Oracle9iAS Infrastructure post-install steps (Note: refer to the *Oracle9iAS for Intel-Linux Install Guide* and *Release Note* for this step)
- shutdown all Oracle9iAS Infrastructure processes (Note: refer to the *Oracle9i Application Server Infrastructure: Improved Availability with Hardware Clusters* paper for this step)
- as **root**, on the same node, copy Oracle configuration files and directories (if they exist) to the other node
  - o **# rcp -p /etc/oratab root@oracleias2:/etc**
  - o **# rcp -p /etc/emtab root@oracleias2:/etc**
  - o **# rcp -p /etc/orainst.loc root@oracleias2:/etc**

## RESTORING TO THE ORIGINAL STATE

Log in, as **root**, to the node where the Oracle9iAS Infrastructure was installed (**oracleias1**, in this example) and perform the following steps.

- change the hostname back to the original hostname
  - o **# hostname oracleias1**
- remove the logical IP address from the primary network interface
  - o **# /sbin/ifconfig eth0:1 down**
- un-mount the filesystem
  - o **# umount /oracle**

## CREATING THE RHCM CLUSTER SERVICE

The provided **9iAS\_Linux\_RHCM\_Scripts.tar** tar-file contains files used in the **TESTED ENVIRONMENT** that will be required for the following steps.

Log in, as **root**, to the node where the Oracle9iAS Infrastructure was installed (**oracleias1**, in this example) and perform the following steps.

- create a directory to host the cluster service files (any directory will work)
  - o **# mkdir /usr/local/bin/9iAS**
- move **9iAS\_Linux\_RHCM\_Scripts.tar** to the **/tmp** directory and

extract from it

- # mv 9iAS\_Linux\_RHCM\_Scripts.tar /tmp
- # cd /tmp
- # tar xvpf 9iAS\_Linux\_RHCM\_Scripts.tar
- copy **start9ias902infra.sh**, **stop9ias902infra.sh**, and **ias\_service** to **/usr/local/bin/9iAS**
  - # cp start9ias902infra.sh /usr/local/bin/9iAS
  - # cp stop9ias902infra.sh /usr/local/bin/9iAS
  - # cp ias\_service /usr/local/bin/9iAS
  - # chmod -R 755 /usr/local/bin/9iAS
- create a **start** and **stop** script within **/usr/local/bin/9iAS** using the example **start** and **stop** files from **9iAS\_Linux\_RHCM\_Scripts.tar** (that are currently in the **/tmp** directory) as references and taking note of the following required changes:
  - **255.255.252.0** should be changed to the netmask value for your network (same value used in the same **SETUP** step)
  - **130.35.149.193** should be changed to the logical IP address
  - **oracleias** should be changed to the virtual hostname
  - **su - oracle** should be changed to **su -** followed by the Red Hat Linux username that performed the Oracle9iAS Infrastructure install
  - **/usr/local/bin/9iAS** should be changed to the directory where the **start9ias902infra.sh** and **stop9ias902infra.sh** were copied to
  - **/oracle/infra90201** should be changed to the **\$ORACLE\_HOME** of the Oracle9iAS Infrastructure install
  - **iasdb** should be changed to the **\$ORACLE\_SID** of the Oracle9iAS Infrastructure install (Note: this is usually always **iasdb**)
  - **admin123** should be changed to the Oracle9iAS Instance Password set during the Oracle9iAS Infrastructure install
  - **oracleias1:0** should be changed to a working display value (usually set via the **\$DISPLAY** environment variable)
  - both the **start** and **stop** scripts should have their execute permissions set for all
    - # chmod a+x /usr/local/bin/9iAS/start

- # **chmod a+x /usr/local/bin/9iAS/stop**
- modify **/usr/local/bin/9iAS/iAS\_service** with the full path location of the **start** and **stop** scripts created in the previous step (default location is **/usr/local/bin/9iAS**)
- create the cluster service using **cluadmin**
  - # **/sbin/cluadmin**
- input as follows (Note: inputs are marked in **bold text**)
  - cluadmin> **service add ias\_infra**
  - Preferred member [None]: **oracleias1**
  - Relocate when the preferred member joins the cluster (yes/no/?) [no]: **no**
  - User script (e.g., /usr/foo/script or None) [None]: **/usr/local/bin/9iAS/iAS\_service**
  - Status check interval [0]: **0**
  - Do you want to add an IP address to the service (yes/no/?) [no]: **no**
  - Do you want to add a disk device to the service (yes/no/?) [no]: **yes**
  - Device special file (e.g., /dev/sdb4): **/dev/sdd3**
  - Filesystem type (e.g., ext2, or ext3): **ext2**
  - Mount point (e.g., /usr/mnt/service1) [None]: **/oracle**
  - Mount options (e.g., rw,nosuid,sync): **[Return]**
  - Forced unmount support (yes/no/?) [yes]: **yes**
  - Would you like to allow NFS access to this filesystem (yes/no/?) [no]: **no**
  - Would you like to share to Windows clients (yes/no/?) [no]: **no**
  - Do you want to (a)dd, (m)odify, (d)elete or (s)how DEVICES, or are you (f)inished adding DEVICES [f]: **f**
  - Add ias\_infra service as shown? (yes/no/?) **yes**
  - choose the node where the Infrastructure is to be started on (**oracleias1**, in this example)
  - cluadmin> **quit**
- remote copy the **/usr/local/bin/9iAS** directory to the other node
  - # **rcp -rp /usr/local/bin/9iAS**

root@oracleias2:/usr/local/bin

## STARTING AND STOPPING THE ORACLE9IAS INFRASTRUCTURE USING THE RHCM CLUSTER SERVICE

At this point, the Oracle9iAS Infrastructure should be up and running on one of the nodes (**oracleias1**, in this example). To stop or start the Oracle9iAS Infrastructure, execute the following command sequences.

(Note: to avoid any confusion by **cluadmin** due to the virtual hostname, any mention of **log in**, below, should be a login via the original hostname (e.g., **oracleias1** or **oracleias2**) and not via the virtual hostname (e.g., **oracleias**))

To stop the Oracle9iAS Infrastructure:

- **log in**, as **root**, to any node of the cluster
- startup **cluadmin**
  - **# /sbin/cluadmin**
- input as follows (Note: inputs are marked in **bold text**)
  - cluadmin> **service disable ias\_infra**
  - Are you sure? (yes/no/?) **yes**
  - cluadmin> **quit**

To start the Oracle9iAS Infrastructure:

- **log in**, as **root**, to any node of the cluster
- startup **cluadmin**
  - **# /sbin/cluadmin**
- input as follows (Note: inputs are marked in **bold text**)
  - cluadmin> **service enable ias\_infra**
  - choose the node where the Infrastructure is to be started on
  - Are you sure? (yes/no/?) **yes**
  - cluadmin> **quit**

For both startup and stop, use **/sbin/clustat** to check the status of the cluster service. Also, logging messages will be appended to the following files during the startup and stop period.

- **/var/log/cluster** (on the startup node)
- **/tmp/iAS\_service.start.log** (on the startup node)
- **/tmp/iAS\_service.stop.log** (on the stop node)

## TESTING THE FAIL-OVER

After the cluster service has been enabled to startup the Oracle9iAS Infrastructure, either one of the following can be performed to simulate the fail-over to the other node. Note that these are two separate command sequences to test the fail-over and that they are not meant to be run one after the other.

(Note: to avoid any confusion by **cluadmin** due to the virtual hostname, any mention of **log in**, below, should be a login via the original hostname (e.g., **oracleias1** or **oracleias2**) and not via the virtual hostname (e.g., **oracleias**))

- manually relocating the service to the other node
  - o **log in**, as **root**, to any node of the cluster
  - o **# /sbin/cluadmin**
  - o input as follows (Note: inputs are marked in **bold text**)
    - cluadmin> **service relocate ias\_infra**
    - Are you sure? (yes/no/?) **yes**
    - cluadmin> **quit**
- shutting down the cluster
  - o **log in**, as **root**, to the node where the Oracle9iAS Infrastructure **is running**
  - o **# /sbin/service cluster stop**

For both cases, notice the Oracle9iAS Infrastructure processes starting up on the other node. To monitor the progress of the fail-over, check for logging messages being appended to the files mentioned above (**/var/log/cluster**, **/tmp/iAS\_service.start.log**, and **/tmp/iAS\_service.stop.log**).

After all Oracle9iAS Infrastructure processes have been started, Oracle9iAS Mid-tier connections to the Infrastructure can be tested to confirm the integrity of the fail-over.



Oracle9i Application Server Infrastructure: Improved Availability with Red Hat Cluster Manager on Red Hat Linux  
November 2002

Author: Ranjan Dutta, Michael Moon, Satish Niradi

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

[www.oracle.com](http://www.oracle.com)

Oracle Corporation provides the software  
that powers the internet.

Oracle is a registered trademark of Oracle Corporation. Various  
product and service names referenced herein may be trademarks  
of Oracle Corporation. All other product and service names  
mentioned may be trademarks of their respective owners.

Copyright © 2002 Oracle Corporation

All rights reserved.