



EXPERIENCES INSTALLING
ORACLE APPLICATION SERVER 10g
(9.0.4)
COLD FAILOVER CLUSTER ON AIX 5L

Revised

May 2005

Contents

Introduction	3
Solution Overview	4
Operating system Preparation	5
Operating System Packages.....	5
OS LEVEL 5200-01	5
HACMP V5.1	5
JCE 1.2.2.....	6
Swap Space	6
Temporary Space.....	7
Storage Preparation.....	8
Check for Available Disk(s)	8
Identify Candidate Disk(s)	8
Prepare Volume Group	9
Prepare Filesystem.....	10
Cluster Preparation	11
Network Planning	11
Cluster Configuration.....	12
Verify and Synchronize Cluster.....	14
Activate Cluster Services	16
Failover Test for Resource Group and Service IP	17
Infrastructure Tier Installation	18
Install Application Server Scripts	22
Middle Tier Installation	24
Summary.....	27
References.....	28
Disclaimer	29
Author(s)	29

INTRODUCTION

The objective of this document is to provide assistance to those involved with the installation of Oracle® Application Server 10g (9.0.4) (AS10g) “Cold Failover Cluster” (CFC) solution on IBM AIX 5L™ (5200) operating system. Two high availability cluster solutions are currently offered by Oracle for AS10g – Cold Failover Cluster and “Active Failover Cluster” (AFC).

CFC provides high availability for the Identity Management (IM) and Metadata Repository (MR) components of an AS10g deployment by relying on an active/passive two-node cluster. AFC delivers high availability of IM and MR functions by leveraging Oracle Real Application Cluster (RAC) technology. As the CFC solution is the focus of this document, please visit Oracle’s High Availability FAQ for Application Server 10g at: http://otn.oracle.com/products/ias/hi_av/904ha_faq.html for additional information on high availability offerings for AS10g, including AFC.

This paper will address the preparation of the hardware infrastructure for a CFC deployment of AS10g, the installation of the Oracle AS10g software stack, and some validation tests for your installation. The configuration presented is for functional validation of this solution only. This paper assumes that the reader is familiar with Oracle Application Server 10g (9.0.4) from a conceptual and functional standpoint, and IBM AIX 5L, and HACMP™.

This document’s primary focus is to supplement “[Oracle Application Server 10g Installation Guide Part No. B13658-02’s](#)” chapter 9 entitled: “[Installing in High availability Environments](#)”. We recommend that you use this document in conjunction with the Oracle Installation Guide during your CFC installation. Also, be sure to examine the “References” portion of this document to identify sources for greater depth on the topics covered in this text.

Your feedback is important for us. We want our technical papers to be as helpful as possible. Please send your comments about this document to the IBM/Oracle International Competency Center (IOICC). The IOICC’s e-mail address is:

ibmoracl@us.ibm.com

SOLUTION OVERVIEW

The diagram below provides an overview of the primary components required to deploy a CFC solution. The primary components involved in this solution are:


1. Middle Tier
2. Infrastructure Tier (IM and MR)
3. Shared Storage

As mentioned in the introduction, an Oracle CFC solution employs an active/passive cluster configuration to support AS10g's Infrastructure Tier components. Clusterware and a shared storage device are the primary features of the cluster. In normal operation mode, the Middle Tier server establishes connection to the Infrastructure Tier via a Virtual IP Address/Hostname (VIP).

This VIP is assigned to the active Infrastructure Tier node. In the event of a failure on the active node (n1), the clusterware detects the service interruption, mounts the shared storage device on the standby node (n2), which starts the AS10g processes, assumes the VIP, and resumes processing requests from the Middle Tier (n3).

It is important to recognize that the only differentiator between a CFC installation and a normal standalone Infrastructure Tier install is that the Oracle Universal Installer will request the Virtual Hostname and the location of a shared disk filesystem. Aside from specifying these parameters, the installation is the same, and most of the effort addressed in the following pages is focused on preparing the underlying hardware and software infrastructure for the cluster.

The nodes involved in your deployment may be standalone servers or, logical equivalents. For our environment, we used logical partitions (LPARs) residing on the same physical IBM @server® pSeries® 650 server. The storage subsystem, IBM 7133-T40 SSA, is directly attached to the two participating cluster nodes and hosts a filesystem for the Oracle software.



OPERATING SYSTEM PREPARATION

This section addresses the preparation of the AIX 5L operating system to support the CFC solution. As we have three nodes in our configuration (Middle Tier and clustered Infrastructure Tier), the following recommendations are typical for all. Please note that at the time of this document's publication, **AS10g (9.0.4) on AIX 5L is certified only on version 5.2 maintenance level (ML) 1 or higher** – oslevel 5200-01.

We recommend that you review the requirements in this section, make any necessary installations, and then apply the most current AIX 5L maintenance level (ML).

To verify fileset levels and APAR numbers in the following section, you may use the following commands:

- For Filesets:
 - #lslpp -L | grep *filesetstring*
 - #lslpp -L *filesetname*
- For APAR:
 - #instfix -iK | grep IYnnnnn (where 'n' is numeric)
 - ex. #instfix -iK "IY45462 IY45707"

OPERATING SYSTEM PACKAGES

The following bos filesets are required on your system and available on your "AIX 5L for POWER™ V5.2 5765-E62" cdroms:

- bos.adt.base
- bos.adt.lib
- bos.adt.libm
- bos.perf.libperfstat
- bos.perf.perfstat
- bos.perf.proctools
- X11.motif.lib
- rsct.basic*
- rsct.compat*
- rsct.core*

OS LEVEL 5200-01

Please verify your operating system (`#oslevel -r`) meets this minimum requirement. If you do not exceed operating system level 5200-01, please apply the following APARs:

- IY43980:libperfstat.h not ANSI-compliant
- IY44810:DSI IN BMRECYCLE
- IY45462:Definition of isnan() in math.h incorrect
- IY45707:J2 READAHEAD/CIO INTERACTION
- IY46214:dropping partial connections leaves them on so_q0
- IY46605:exec of 32 bit application can fail on 64 bit kernel
- IY48525:SDK 1.4.1 32-BIT SR1: CA141-20030930
- IY51801:race condition in aio_nwait_timeout

To obtain an APAR or other AIX 5L fix, please navigate to the IBM Support Site and download:

<http://www-912.ibm.com/eserver/support/fixes/>

HACMP V5.1

Please install "HACMP V5.1 for AIX® Subsystems 5765-F62" on your system (`#smitty installp ->'Install Software'`). Once complete, please apply the following APAR's:

- IY42783 CT:LX: RMC daemon may hang if managed nodes recycle.

- IY43602 Disk Failures causing quorum to be lost is not
- IY45695 (PTF1) Base Fixes for HACMP 5.1.0
- U496124 (PTF2) cluster.es.server.rte.5.1.0.2

Verify that your installed cluster filesets meet or exceed the following fileset levels:

Fileset	Level	State	Type	Description (Uninstaller)
cluster.es.client.lib	5.1.0.4	C	F	ES Client Libraries
cluster.es.client.rte	5.1.0.2	C	F	ES Client Runtime
cluster.es.client.utils	5.1.0.3	C	F	ES Client Utilities
cluster.es.clvm.rte	5.1.0.0	C	F	ES for AIX Concurrent Access
cluster.es.cspoc.cmds	5.1.0.4	C	F	ES CSPOC Commands
cluster.es.cspoc.dsh	5.1.0.0	C	F	ES CSPOC dsh
cluster.es.cspoc.rte	5.1.0.4	C	F	ES CSPOC Runtime Commands
cluster.es.server.diag	5.1.0.3	C	F	ES Server Diags
cluster.es.server.events	5.1.0.4	C	F	ES Server Events
cluster.es.server.rte	5.1.0.4	C	F	ES Base Server Runtime
cluster.es.server.utils	5.1.0.4	C	F	ES Server Utilities
cluster.license	5.1.0.0	C	F	HACMP Electronic License
cluster.man.en_US.es.data	5.1.0.4	C	F	ES Man Pages - U.S. English
cluster.msg.en_US.cspoc	5.1.0.0	C	F	HACMP CSPOC Messages - U.S.

If it is necessary for you to install any of the above filesets, please *reapply any maintenance level(s)* that your system has received since its base operating system install.

JCE 1.2.2

Java™ Cryptography Extension (JCE) 1.2.2 is required for your installation. You may download the package at the following location: <http://java.sun.com/products/jce/downloads/index.html>

Place the jce-1_2_2.zip file in a location writable by the user performing the installation. Note: It is not required that you have Java 2 SDK or RTE v1.2x, v1.3x, or v1.4x installed on your system for the installation of AS10g, as the Oracle Universal Installer (OUI) will meet the java requirements. As our environment featured SDK v1.4x, we placed jce-1_2_2.zip within our \$JAVA_HOME (/usr/java14). You will need to specify the jce-1_2_2.zip file location during the install.

SWAP SPACE

The OUI will validate that your swap space meets or exceeds 1536MB. The following commands will help you to evaluate your current primary paging space and increment if needed (512MB -> 1536MB):

- a) Check your current swap space. Here you can see there is one device that is 512MB:

```
#/usr/sbin/swap -l
device          maj,min      total      free
/dev/hd6        10,  2       512MB      506MB
```

- b) Obtain Physical Partition (PP size) for paging space. The PP size will be used in the next step to increase the swap space:

```
#/usr/sbin/lslv -L hd6 | grep SIZE
MAX LPs:          512
PP SIZE:          32
megabyte(s)
```

- c) Increase the size of your swap device as desired. You specify the number of additional logical partitions (based on PP size) you wish to add to a paging space. For example, we started with 512MB and would need 1024MB of additional space to achieve the 1536MB minimum requirement. Based on our 32MB PP size for our /dev/hd6 swap device, we will add eight additional partitions to it (8 * 32MB= 1024MB additional swap space) with the chps command:

```
#/usr/sbin/chps -s 8 hd6
```

- d) Verify changes

```
#/usr/sbin/swap -l
device          maj,min      total      free
/dev/hd6        10,  2      1536MB     1524MB
```

TEMPORARY SPACE

The OUI will validate that your temporary space meets or exceeds 400MB. Size your temporary space accordingly.

Now that we have satisfied the operating system requirements for the Oracle software across all tiers in our configuration, we will move on to configuring the shared storage device for our Infrastructure Tier nodes.

STORAGE PREPARATION

As mentioned in the Solution Overview, the CFC incorporates a shared storage device, and in this section we will examine the preparation of that device, and a shared filesystem to house the "Identity Management and Oracle Metadata Repository" for Oracle AS10g. This is not a concurrently accessed filesystem, but one that is mounted to the primary node until that node experiences a failure. In the event of a failure on the primary node, the surviving or takeover node can mount the shared storage device and start the Infrastructure Tier services.

Prior to proceeding, please verify that your physical infrastructure has established connectivity to your shared storage device.

We will be performing the following to prepare the storage:

1. Verify configured physical devices visible to CFC nodes
2. Identify an available physical device visible to CFC nodes
3. Use AIX 5L Logical Volume Manager to create a volume group and logical volume
4. Create a filesystem and mount points, one per node, for the filesystem

Note: All of the following operations may be conducted via the command line, or through the AIX 5L SMIT, Software Management Interface Tool, via appropriate menus. Our examples in this section will show command line procedures.

CHECK FOR AVAILABLE DISK(S)

```
#/usr/sbin/lssdev -Cc disk
```

Output from this command will vary according to the number of disk attached to your system. Our two identically configured CFC nodes featured 12 SSA disks, with four disks specific to each node, and eight disks visible to both. Listing the disk devices from both nodes via the lssdev command will enable you to verify that your nodes are seeing the disk you expect them to. Our output:

```
hdisk0 Available 1A-08-L SSA Logical Disk Drive
hdisk1 Available 1A-08-L SSA Logical Disk Drive
hdisk10 Available 1A-08-L SSA Logical Disk Drive
hdisk11 Available 1A-08-L SSA Logical Disk Drive
hdisk2 Available 1A-08-L SSA Logical Disk Drive
hdisk3 Available 1A-08-L SSA Logical Disk Drive
hdisk4 Available 1A-08-L SSA Logical Disk Drive
hdisk5 Available 1A-08-L SSA Logical Disk Drive
hdisk6 Available 1A-08-L SSA Logical Disk Drive
hdisk7 Available 1A-08-L SSA Logical Disk Drive
hdisk8 Available 1A-08-L SSA Logical Disk Drive
hdisk9 Available 1A-08-L SSA Logical Disk Drive
```

If your device listing does not meet expected results, run AIX configuration manager:

```
#/usr/sbin/cfgmgr
```

Re-run the lssdev command from the previous step to check for any newly added disk devices.

IDENTIFY CANDIDATE DISK(S)

Once we have established that physical devices are configured, we need to identify the physical volume identifiers (PVID's) for disks that are not currently associated with a volume group.

```
#/usr/sbin/lspv
```

```
hdisk0      0002c5da584fbc29      rootvg      active
hdisk1      0002c5dadae6ab16      None
hdisk2      0002c5da12794219      None
hdisk3      0002c5daa21ae7f5      None
hdisk4      0002c5da1549b70f      None
hdisk5      0002c5da4f31695a      oravg      active
hdisk6      0002c5dac68e5d2a      None
hdisk7      0002c5dac68e8f40      None
```

hdisk8	0002c5dac68ea192	None
hdisk9	0002c5daa21aee64	none
hdisk10	0002c5da4f31e54d	None
hdisk11	0002c5daec8bd2e9	None
hdisk12	000747ff9178193a	None
hdisk13	000747ff9594f230	None

For each disk, this command shows:

- The disk device name
- Either the 16 character physical volume identifier (PVID) if the disk has one, or 'None'
- Either the volume group identifier to which the disk belongs, or 'None'
- Physical volume state – active, missing or removed

The disks that you want to use might have a PVID, but they must not belong to a volume group. If the disks that you want to use for the disk group do not have PVIDs, enter a command similar to the following for each disk that you want to use:

```
#/usr/sbin/chdev -l hdiskn -a pv=yes (where 'n' is the numeric hdisk identifier)
```

Once you have assigned a PVID on the first node (n1), you may execute this command on the second node (n2), to read the new PVID or you may remove the device via:

```
#/usr/sbin/rmdev -dl hdiskn
Then run AIX 5L configuration manager to add the disk devices:
#/usr/sbin/cfgmgr
Verify the new PVID information:
#/usr/sbin/lspv
```

The PVID's will provide a consistent identifier across the CFC nodes, while "hdisknn" identifiers may or may not map consistently across nodes. After establishing a view of PVID's from the participating CFC nodes, select one or more disks that you wish to use for your Infrastructure Tier filesystem. The disk(s) you choose should not have any association to a volume group, and have the same PVID(s) visible from both nodes.

PREPARE VOLUME GROUP

The commands provided below echo what was configured for our test environment. They are provided as examples only.

On n1:

```
#/usr/sbin/mkvg -y 'sharedvg' -s'16' hdisk3,hdisk9
```

Verify your volume group. Your output should show the volume group you created in an active state:

```
#/usr/sbin/lspv
hdisk0      0002c5da584fbc29      rootvg      active
hdisk1      0002c5dadae6ab16      None
hdisk2      0002c5da12794219      None
hdisk3      0002c5daa21ae7f5      sharedvg    active
hdisk4      0002c5da1549b70f      None
hdisk5      0002c5da4f31695a      None
hdisk6      0002c5dac68e5d2a      None
hdisk7      0002c5dac68e8f40      None
hdisk8      0002c5dac68ea192      None
hdisk9      0002c5daa21aee64      sharedvg    active
hdisk10     0002c5da4f31e54d      None
hdisk11     0002c5daec8bd2e9      None
```

Next, modify the volume group to disable autovaryon, and then export.

```
#/usr/sbin/chvg -a n 'sharedvg' (autovaryon = no)
#/usr/sbin/varyoffvg sharedvg
#/usr/sbin/exportvg sharedvg
```

On n2:

First list the physical volumes on the system:

```
#/usr/sbin/lspv
hdisk0      0002c5da7e6310d6      rootvg      active
hdisk1      0002c5daea7c5474      None
```

hdisk2	0002c5dadae6ab16	None
hdisk3	0002c5da12794219	None
hdisk4	0002c5daa21ae7f5	None
hdisk5	0002c5da1549b70f	None
hdisk6	0002c5dac68e5d2a	None
hdisk7	0002c5dac68e8f40	None
hdisk8	0002c5dac68ea192	None
hdisk9	0002c5daa21aee64	None
hdisk10	0002c5da126c0f27	None
hdisk11	0002c5da4f958273	None

Identify the `hdiskn` corresponding to one of the two PVID's associated with the sharedvg created on n1. Note: if the volume group was created spanning multiple hdisks on n1, you need only specify one of the "hdisknn" devices on n2 to import the volume group.

Import the volume group created from n1 to n2:

```
#/usr/sbin/importvg -y 'sharedvg' hdisk4
```

Note: **hdisk4** is "hdisknn" on n2 corresponding to PVID=0002c5daa21ae7f5.

```
#/usr/sbin/varyoffvg sharedvg
```

On n1:

Re-import the volume group:

```
#/usr/sbin/importvg -y 'sharedvg' hdisk3
```

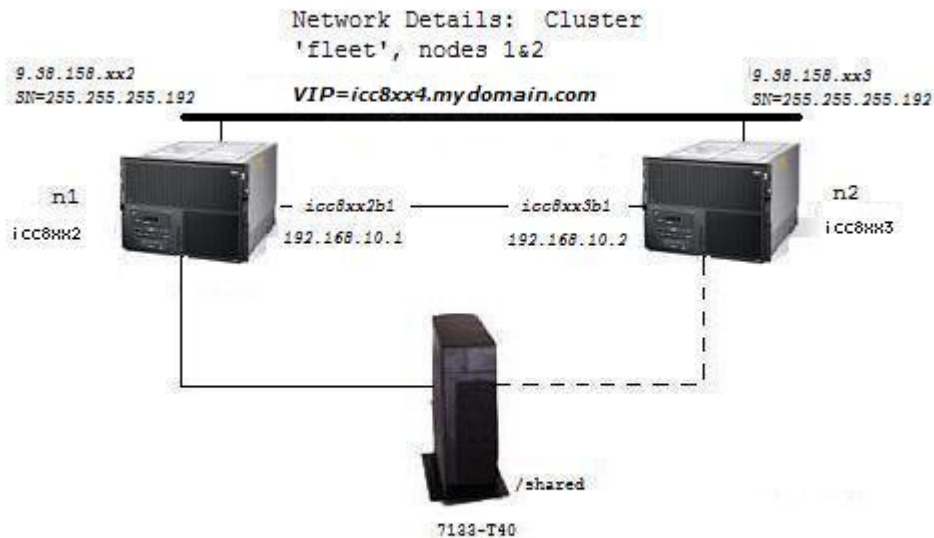
PREPARE FILESYSTEM

Please prepare a filesystem for storing your Infrastructure Tier Oracle software - "\$ORACLE_HOME". We prepared a logical volume, "sharedlv" (`#mklv -y 'sharedlv' oravg nn`), and then created a jfs filesystem (`#smitty jfs`) on this logical volume. Please size your filesystem to accommodate the software requirements, approximately 4.3GB, and any additional capacity you require.

Also, verify that you can mount your filesystem from each node. As this is not a concurrent capable storage configuration, you can only have the shared volume group active on one node at a time.

CLUSTER PREPARATION

This section examines the creation of an HACMP V5.1 cluster to support the Infrastructure Tier. Below is a diagram of our testing environment featuring the hardware and network information for cluster "fleet":



In addition to node down detection, and IP Address takeover, HACMP will provide for an 'Application Server' definition, which we will use to specify the location of two scripts – start and stop for the Infrastructure Tier services. These scripts enable you to invoke a graceful shutdown for maintenance, and provide HACMP the ability to start all of the Infrastructure Tier services in the event of a takeover by the standby node.

For additional information on HACMP services, resource groups, IP aliasing, and more, please refer to the "[HACMP for AIX: Concepts and Facilities \(SC23-4864-03\)](#)". You can access this document from the HACMP Library link provided in the "References" section.

NETWORK PLANNING

You need to provide for HACMP private and public network access within your cluster. HACMP uses the private network to monitor the health of the cluster, and clients access the nodes via the public network. In the diagram above, our private network settings are visible below the heavy black line that identifying the VIP information. Public network information is depicted above the line.

When configuring your network interfaces (IF), the primary (public) and secondary (private) interfaces need to be on different subnets, but using the same subnet mask.

Our `#/etc/hosts` files were consistent across both n1 & n2 and include the following entries:

```

9.38.158.xx2    icc8xx2.sanmateo.ibm.com    icc8xx2
9.38.158.xx1    icc8xx1.sanmateo.ibm.com    icc8xx1
9.38.158.xx3    icc8xx3.sanmateo.ibm.com    icc8xx3
9.38.158.xx4    icc8xx4.sanmateo.ibm.com    icc8xx4
192.168.10.2    icc8xx3b1
192.168.10.1    icc8xx2b1

```

We have defined only one private network interface (192.168.10.x) per cluster node. This is acceptable in our testing environment, but not recommended for production deployments, as it represents a single failure point for this network. You will see later on in this section that HACMP recognizes these configuration issues and will make you aware of them.

CLUSTER CONFIGURATION

The following steps that were used to create the "fleet" cluster to support the Infrastructure Tier will be expressed in terms of SMIT menus. We will use the '>' to express progressive nesting of SMIT screens, to assist you in navigating this interface.

Please follow our procedure for creating cluster "fleet":

1. Configure primary and secondary network adapters

a. Primary

```
#smitty tcpip
>'Minimum configuration & Startup'
>>'Available Network Interfaces'
>>>'Select your primary network adapter (en0 in our configuration)'
>>>>'Minimum Configuration & Startup'
    Provide network configuration information
    Note: When configuring your network interfaces on the participating
    nodes, the primary and secondary interfaces need to be on different subnets, but using the same
    subnet mask.
```

b. Secondary

```
#smitty chinnet
>>'Available Network Interfaces'
>>>'Select your secondary network adapter (en3 in our configuration)'
>>>>'Change / Show a Standard Ethernet Interface'
>>>>>'Provide network configuration information'
    Note: When configuring your network interfaces on the participating
    nodes, the primary and secondary interfaces need to be on different subnets, but using the same
    subnet mask.
```

2. Define the Cluster Topology

When defining the cluster, participating nodes, network information, resource group, and application server to HACMP, you will be performing this activity on one node – typically the primary node n1. It is unnecessary for you to perform this activity on each node, as the complete definition you build on n1 will be captured by the participating cluster nodes during HACMP's 'Verify and Synchronize Cluster' step we perform at the conclusion of the definition phase.

a. Add Cluster

```
#smitty hacmp
>'Extended Configuration'
>>'Extended Topology Configuration'
>>>'Configure an HACMP Cluster'
>>>>'Add/Change/Show an HACMP Cluster'
>>>>>'Cluster Name' ("fleet")
```

b. Add Cluster Nodes

```
#smitty hacmp
>'Extended Configuration'
>>'Extended Topology Configuration'
>>>'Configure HACMP Nodes'
>>>>'Add a Node to the HACMP Cluster'
```

Specify a node name (short name), and its corresponding public internet address as the 'Communication Path to Node'. Perform this for both nodes.

c. Add Cluster Network

```
#smitty hacmp
>'Extended Configuration'
```

```
>>'Extended Topology Configuration'
>>>'Configure HACMP Networks'
>>>>'Add a Network to the HACMP Cluster'
      Select '# Pre-defined IP-based Network Types' ("ether") from the resulting pop-up.
```

Specify a 'Network Name' or accept default, provide your Netmask, and accept the default "Yes" value for 'Address Takeover'. Our example:

```
* Network Name                [net_ether_010]
* Network Type                ether
* Netmask                     [255.255.255.192]
* Enable IP Address Takeover via IP Aliases [Yes]
  IP Address Offset for Heartbeating over IP Aliases []
```

d. Add Network Interfaces

```
#smitty hacmp
>'Extended Configuration'
>>'Extended Topology Configuration'
>>>'Configure HACMP Communication Interfaces/Devices'
>>>>'Add Communication Interfaces/Devices'
      Select 'Add Pre-defined Communication Interfaces and Devices' from the pop-up
      Select 'Communication Interfaces' from the pop-up
      Choose your defined network ("net_ether_010" in our example) from 'Select a Network'
>>>>>'Add a Communication Interface'
```

Provide your network information for your network adapters. Perform this for each node's adapters. Our information for private network interfaces (one per node):

Note: "n1" data provided for 'Entry Fields' and "n2" data within "<>" for reference:

```
* IP Label/Address           [Entry Fields]
                             [icc8xx2b1]          <icc8xx3b1>
* Network Type              ether
* Network Name              net_ether_010
* Node Name                  [icc8xx2]          <icc8xx3>
  Network Interface          [en3]            <en3>
```

e. Add Resource Group

```
#smitty hacmp
>'Initialization and Standard Configuration'
>>'Configure HACMP Resource Groups'
>>>'Add a Resource Group'
      Select a Resource Group Management Policy" (Rotating)
>>>>Specify a 'Resource Group Name' ("rtrg" in our example), and 'Participating Node
      Names / Default Node Priority' ("icc8xx2", "icc8xx3")
```

Note: Default Node Priority decreases from left to right, and is used to determine which standby node acquires the resource in a multiple standby configuration. For our example, we list our primary node first, n1, followed by our secondary node, n2.

f. Add Service IP Labels/Addresses

```
#smitty hacmp
>'Initialization and Standard Configuration'
>>'Configure Resources to Make Highly Available'
>>>'Configure Service IP Labels/Addresses'
>>>>'Add a Service IP Label/Address'
>>>>>Select an 'IP Label/Address' (F4 in the Entry Field)
```

Note: The Service IP Label/Address you specify here should be the VIP for your CFC – 'icc8xx4' in our example. This address should reside on the same subnet as the public IP Address for the nodes and should have been defined in your /etc/hosts during your network configuration steps for the cluster. Select the 'Network Name' field and depress F4. You should see the network ("net_ether_010") we defined earlier.

Note: The Service IP Label will reside on the same network you have defined for your cluster "net_ether_010" in our example.

g. Add Service IP Labels/Addresses to Resource Group

```
#smitty hacmp
>'Initialization and Standard Configuration'
>>'Configure HACMP Resource Groups'
>>>'Change/Show Resources for a Resource Group (standard)'
>>>'Select a Resource Group' ("rtrg")
>>>>specify a 'Service IP Labels/Addresses' ("icc8xx4"), & 'Application Servers' ("ias")
```

Note: This final menu should display the 'Resource Group Name' ("rtrg"), and 'Participating Node Names' ("icc8xx2","icc8xx3"), defined with the Resource Group definition in the previous step.

Our information:

	[Entry Fields]	
Resource Group Name	rtrg	
Participating Node Names (Default Node Priority)	icc8xx2 icc8xx3	
* Service IP Labels/Addresses	[icc8xx4]	+
Volume Groups	[sharedvg]	+
Filesystems (empty is ALL for VGs specified)	[]	+
Application Servers	[ias]	+

h. Add Application Server

As mentioned at the start of the Cluster Preparation section, an HACMP Application Server will provide you the ability to specify the Infrastructure Tier service start/stop scripts. You will need to locate these scripts locally on each participating node.

The scripts you will eventually need are contained within the *.zip file you downloaded containing this document. As we will be verifying our cluster configuration, and testing with a failover once we complete the configuration, we recommend that you touch a file and provide it with execute permissions within a directory local to each node. Following the installation of the Oracle software, you will overlay these empty, "place holder", files with the provided service start and stop scripts.

```
#touch /home/oracle/scripts/infra_start
#touch /home/oracle/scripts/infra_stop
#chmod +x infra_start infra_stop
```

Now add the Application Server:

```
#smitty hacmp
>'Extended Configuration'
>>'Extended Resource Configuration'
>> 'HACMP Extended Resources Configuration'
>>>'Configure HACMP Application Servers'
>>>>'Add an Application Server'
>>>>>Provide the 'Server Name' <"ias">,
'Start Script' <"/home/oracle/scripts/infra_start">, and
'Stop Script' <"/home/oracle/scripts/infra_stop">
```

VERIFY AND SYNCHRONIZE CLUSTER

```
#smitty hacmp
>'Initialization and Standard Configuration'
>>'Verify and Synchronize HACMP Configuration'
```

The output from the synchronization of our environment is provided below for your reference:

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

[TOP]

Verification to be performed on the following:

Cluster Topology
Cluster Resources

Retrieving data from available cluster nodes. This could take a few minutes...

Verifying Cluster Topology...

WARNING: There may be an insufficient number of communication interfaces defined on node: icc8xx2, network: net_ether_010. Multiple communication interfaces are recommended for networks that will use IP aliasing.

WARNING: There may be an insufficient number of communication interfaces defined on node: icc8xx3, network: net_ether_010. Multiple communication interfaces are recommended for networks that will use IP aliasing.

Verifying Cluster Resources...

WARNING: Error notification stanzas will be added during synchronization for the following:

Node: icc8xx2
en_label: LVDD
en_resource: NONE
en_class: LVM_SA_QUORCLOSE

WARNING: Error notification stanzas will be added during synchronization for the following:

Node: icc8xx3
en_label: LVDD
en_resource: NONE
en_class: LVM_SA_QUORCLOSE

Remember to redo automatic error notification if configuration has changed.

Verifying additional pre-requisites for Dynamic Reconfiguration...
...completed.

Updating ODM errnotify on node icc8xx2.
Updating ODM errnotify on node icc8xx3.

Verification has completed normally.

Cldare: No changes detected in Cluster Topology or Resources.

```

...completed.

Updating ODM errnotify on node icc8xx2.
Updating ODM errnotify on node icc8xx3.

Verification has completed normally.

Clsnapshot: Creating file /usr/es/sbin/cluster/snapshots/active.0.odm...

clsnapshot: Succeeded creating Cluster Snapshot: active.0.

cldare: Requesting a refresh of the Cluster Manager...
0513-095 The request for subsystem refresh was completed successfully.
...completed.

```

As we mentioned at the beginning of this section, HACMP will assess your configuration. The above output highlights a limitation of our cluster configuration. Specifically, the two warnings we received during the "Verifying Cluster Topology" portion indicate that we should include additional communication interfaces for reliability. Please be sure to configure your cluster appropriately to meet your service level targets.

ACTIVATE CLUSTER SERVICES

Perform the following procedure on n1:

```

#smitty hacmp
>' System Management (C-SPOC)'
>>' Manage HACMP Services'
>>>' Start Cluster Services'

```

Tail the `"/tmp/hacmp.out"` log and review messages. It's a good practice to monitor this log when starting/stopping HACMP. SMIT will indicate promptly to you that services have been started, however this process takes time, and the log will provide greater detail on where HACMP is in the process of starting your cluster.

You can also view the status of HA Group Services using the following command:

```
#/usr/bin/lssrc -ls grpsvcs
```

With one n1 running, you should see output similar to:

```

#/usr/bin/lssrc -ls grpsvcs
Subsystem      Group          PID           Status
grpsvcs        grpsvcs        573536        active
2 locally-connected clients. Their PIDs:
561398(haemd) 585926(clstrmgr)
HA Group Services domain information:
Domain established by node 1
Number of groups known locally: 3

```

Group name	Number of providers	Number of local providers/subscribers
ha_em_peers	1	1
CLRESMGRD_1067971441	1	1
CLSTRMGR_1067971441	1	1

The "Number of providers" value above indicates the number of active nodes within the cluster.

Once you have received positive confirmation that HACMP is up on n1, please activate HACMP cluster services on n2 in the same manner. Note: n2 will log an entry to `"/tmp/hacmp.out"` on n1 when it has joined the cluster successfully. You may tail the local `"/tmp/hacmp.out"` on n2 for greater detail on the status of HACMP services.

`#/usr/bin/lssrc -ls grpsvcs` should deliver an output similar to the following indicating n2 is up:

```

#/usr/bin/lssrc -ls grpsvcs
Subsystem      Group          PID           Status
grpsvcs        grpsvcs        573536        active
2 locally-connected clients. Their PIDs:
561398(haemd) 585926(clstrmgr)

```

```

HA Group Services domain information:
Domain established by node 1
Number of groups known locally: 3
      Number of   Number of local
Group name   providers   providers/subscribers
ha_em_peers           2             1             0
CLRESMGRD_1067971441  2             1             0
CLSTRMGR_1067971441  2             1             0

```

Please verify that the shared filesystem you created is mounted on n1:

```

#/usr/bin/df | grep shared
/dev/sharedlv 42598400 34220360 20% 31104 1% /shared

```

Verify network:

Check for availability of port 1521:

```

#/usr/bin/netstat -an | grep 1521

```

You can view your network interface status and verify your public and private interfaces:

```

- #/usr/bin/netstat -in
Name Mtu Network Address Ipkts Ierrs Opkts Oerrs Coll
en0 1500 link#2 0.9.6b.4e.58.68 98480 0 72349 0 0
en0 1500 9.38.158.12 9.38.158.xx2 98480 0 72349 0 0
en3 1500 link#3 0.9.6b.4e.58.6b 217972 0 196217 0 0
en3 1500 192.168.10 192.168.10.1 217972 0 196217 0 0
en3 1500 9.38.158.12 9.38.158.xx4 217972 0 196217 0 0

```

Above we see the two public addresses, VIP included, and the private network address active on n1

FAILOVER TEST FOR RESOURCE GROUP AND SERVICE IP

Prior to installing Oracle AS10g Infrastructure Tier software, we encourage you to perform a failover / fallback test of your cluster.

On Primary node n1:

```

#smitty hacmp
>' System Management (C-SPOC)'
>>' Manage HACMP Services'
>>>' Stop Cluster Services'
      In the 'Shutdown mode' field, press F4, and select 'takeover'

```

On Secondary node n2:

Tail `"/tmp/hacmp.out"` and look for HACMP to acquire resources from n1. Once takeover is complete, perform the validation steps mentioned in the preceding step to verify that the shared storage filesystem has been mounted and that the Service IP Label (VIP) is active. You should be able to ping the VIP by fully qualified host name (`#ping icc8xx4.sanmateo.ibm.com`), and IP address.

On Primary node n1:

```

Re-start your cluster services
#smitty hacmp
>' System Management (C-SPOC)'
>>' Manage HACMP Services'
>>>' Start Cluster Services'

```

Once n1 has re-joined the cluster, stop cluster services on n2 in 'takeover' mode. The primary node should acquire the resources. Perform verification of storage and VIP. Start cluster services on n2. We are now ready to install the Infrastructure Tier software.

INFRASTRUCTURE TIER INSTALLATION

With the supporting HACMP cluster configuration completed, and the HA Group Services up and running, we have satisfied nearly all of the installation pre-requisites for the AS10g Infrastructure Tier software installation. The remaining requirements are detailed in section two of [Oracle Application Server 10g Quick Installation and Upgrade Guide10g \(9.0.4\) for AIX-Based Systems # B13662-01](#). This document, available at http://download-west.oracle.com/docs/cd/B12428_09/quickinstall.904/quickinstallax.pdf, will assist you in defining a group and user to own the Oracle software, and set your environment variables. Additionally, you can locate further Infrastructure Tier installation details starting on page 254 of the "[Oracle Application Server 10g Installation Guide Part No. B13658-02](#)".

Oracle Universal Installer (OUI)

Prior to invoking the OUI, we set, performed or verified the following within our environment:

- created a "dba" group and an "oracle" user to own the Oracle software
- exported our DISPLAY to an X terminal
- checked for port availability (`#/usr/bin/netstat -an | grep 1521`) Note: Portal and Wireless developer desire this port. If unavailable on your system, please refer to [Oracle Application Server 10g Installation Guide](#) for additional discussion on this topic.
- executed the `rootpre.sh` on both cluster nodes. Note: you need to execute this script on all participating cluster nodes.

1. Invoke the OUI

```
#!/runInstaller
```

- Confirm that you have executed the `rootpre.sh` (both nodes)
- Provide the path to the location of your `jce-1_2_2.zip` file
- Verify that you pass the installation checks and address any deficiencies

Once you have passed all pre-installation checks, you should see a message similar to the following:

```
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2004-04-12_10-46-38PM.  
Please wait ...
```

Note: Though not encouraged, you may be in a situation where you wish to bypass the pre-installation checks performed by the OUI:

```
#!/runInstaller -ignoreSysPrereqs
```

Once you have passed all pre-installation checks, you should see a message similar to the following:

2. Welcome

If you have performed an installation earlier that you wish to clean-up select "Deinstall Products" otherwise select Next to continue.

3. Specify Inventory Directory

If this is the first Oracle product installation on your machine, you will see this menu, and need to provide a full path for the inventory directory. Enter a path different from your intended \$ORACLE_HOME (location of the AS10g software), but residing on the shared storage device.

Example: `/shared/product/oraInventory`

4. UNIX® Group Name

If this is the first Oracle product installation on your machine, you will need to specify the name of the operating system group that will have permission to upgrade/install Oracle software on your system.

We created a "dba" group to own and manage Oracle software

5. Run `orainstRoot.sh`

If this is the first Oracle product installation on your machine, you will need to execute this shell script as the `root` user. The prompting message will specify the location of the script in your Oracle inventory directory.

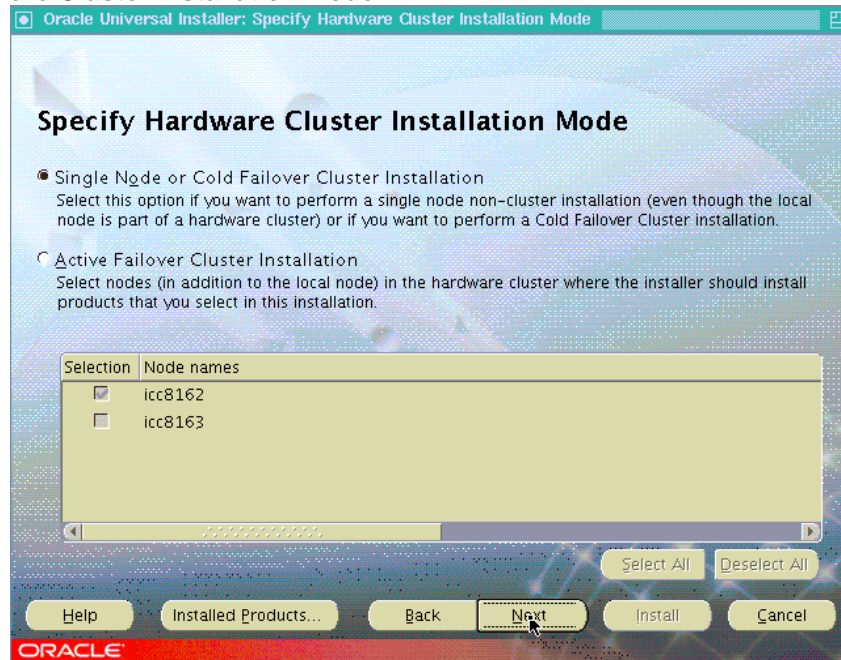
6. Specify File Locations

Specify a new name for your \$ORACLE_HOME, and be sure to specify the path to your shared storage filesystem as the location.

Example:

Name: iasHome
Path: /shared/product/904

7. Specify Hardware Cluster Installation Mode



Choose the Cold Failover Cluster (CFC) option. As you can see on this screen snapshot, our two clustered nodes were identified by the OUI. If you do not see this screen, the OUI was unable to detect HACMP running.

8. Select a Product to Install

Choose the "OracleAS Infrastructure 10g 9.0.4.0.0" installation option. Click Next to continue.

9. Select Installation Type

Select; "Identity Management and OracleAS Metadata Repository (4.27 GB)" as your installation type. Click Next to continue.

10. Preview of Steps for Infrastructure Installation

This screen provides an overview of the pending configuration screens the OUI will display. Click Next to continue.

11. Confirm Pre-Installation Requirements

If you have not verified the availability of port 1521 yet, please check it now using this command: `#/usr/bin/netstat -an | grep 1521`

You will also need to indicate that you have root privileges to proceed with the installation. Click in the option boxes to indicate you have satisfied these requirements and click Next to continue.

Refer to section 2.8 of [Oracle Application Server 10g Quick Installation and Upgrade Guide10g \(9.0.4\) for AIX-Based Systems](#) for port availability issues.

12. Select Configuration Options

Select all components **except** "OracleAS Certificate Authority". Be sure to notice the status of the "High Availability Addressing" option – it should be pre-selected and by default.

13. Specify Namespace in Internet Directory

Verify the suggested namespace or customize as necessary. Please refer to the aforementioned Oracle Installation Guides for further information on Namespace specification.

14. Specify High Availability Addressing

This is the point in the installation where you specify your fully qualified **VIP**. This is the value we specified as a "Service IP Labels/Addresses" during our Cluster Configuration step.

Example: `icc8xx4.sanmateo.ibm.com`

15. Specify Privileged Operating System Groups

If you are executing this installation as a user who is not in the `dba` privileged group, this screen will appear and prompt you for the name of that privileged group.

16. Database Identification

Specify your global database name, and append your "*mydomain.com*" to it. Your SID may not exceed eight characters.

Example: Global Database Name: `iasdb.sanmateo.ibm.com`
 SID: `iasdb`

17. Set SYS and SYSTEM Passwords

Provide passwords and click "Next" to continue.

18. Database File Location

Specify a location on your shared filesystem for the datafiles.

Example: Database File Location: `/shared/product/oradata`

19. Database Character Set

Choose your Character Set. Click "Next" to continue.

20. Specify Instance Name and `ias_admin` Password

Provide the requested information and click "Next" to continue.

21. Summary

Review the installation summary and click "Install" to continue.

22. Install -> Prompt to run `root.sh`

When prompted, run the `root.sh` script from a separate window. The script will prompt you for your local bin directory.

```
#root@icc8xx2@/shared/product/904>./root.sh
Running OracleAS root.sh script...

The following environment variables are set as:
ORACLE_OWNER= oracle
ORACLE_HOME= /shared/product/904

Enter the full pathname of the local bin directory: [Default:/usr/local/bin]:

Adding entry to /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
Verifying SunJCE installation ....
..Logging messages in /shared/product/904/sso/log/jcetest.log

SunJCE installation is verified to be OK.
You can proceed with the rest of the Oracle AS10g installation

Entering Oracle Internet Directory Root Installation Section
```

```
Oid Server Installation
Checking LDAP binary file protections
Setting oidmon file protections
Setting oidldapd file protections
Setting oidrepld file protections
Setting oidpasswd file protections
Setting oidempasswd file protections
Setting oidstats.sh file protections
Setting remtool file protections
Leaving Oracle Internet Directory Root Installation Section
```

23. Configuration Assistants

This screen will show you the progress of the configuration assistants. You can tail the "\$ORACLE_INVENTORY/logs/installActions*.log" for additional detail on how the install is progressing.

24. End of Installation

Review the provided URL summary information, and note the location of the setupinfo.txt file for future reference. Click "Exit" to end the installation process.

INSTALL APPLICATION SERVER SCRIPTS

Within the `as10g904_cfc_aix.zip` file containing this document, there are four scripts:

- `infra_start`
- `startias904infra.sh`
- `infra_stop`
- `stopias904infra.sh`

The "infra_*" scripts are invoked by HACMP during a cluster services start/stop, or as a result of a takeover by a standby cluster node. They in turn execute the AS10g service, "startias*" / "stopias*", script as the object owner – oracle in our environment.

1. Place the "infra_start" script into the location you specified under 'Start Script' in the Application Server definition. Recall that during step 2h of our "Cluster Configuration" process, we specified a placeholder file within our "Application Server" definition for our HACMP resource group. Now, we will provide the actual service start/stop scripts for the Oracle Infrastructure Software we have just installed.

You can Show the configuration of your Application in the following manner:

```
#smitty hacmp
>Extended Configuration
>>Extended Resource Configuration
>>>HACMP Extended Resources Configuration
>>>>Configure HACMP Application Servers
>>>>>Change/Show an Application Server
```

Our configuration data:

Change Application Server

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]
Server Name	ias
New Server Name	[ias]
Start Script	[/home/oracle/scripts/infra_start>
Stop Script	[/home/oracle/scripts/infra_stop>

2. Upload the `infra_start` to the location specified in your 'Start Script' parameter
For example: `"/home/oracle/scripts/infra_start"`

```
infra_start:
#!/bin/ksh
su - oracle << !
/home/oracle/scripts/startias904infra.sh
echo
echo Completed iAS 904 Infrastructure Startup @ `date`
!
```

3. Perform same procedure for `infra_stop` using the path specified in your 'Stop Script' Parameter.
4. Notice in the `infra_start` script above that "startias904infra.sh" is called from within `infra_start` from `/home/oracle/scripts` by the oracle user. Upload `startias904infra.sh` to the location specified by `infra_start` or one of your choosing. Verify that `infra_start` points to the specified location.
5. Perform the same procedure for `stopias904infra.sh`.
6. Change the mode of the scripts to make them executable (`#!/usr/bin/chmod +x filename`) and ownership (`#!/usr/bin/chown oracle:dba filename`) on the scripts. Your listing should appear similar to the following:

```
root@icc8xx2@/home/oracle/scripts>ls -la
```

```
total 48
-rwxr-xr-x 1 oracle dba          134 May 19 12:42 infra_start
-rwxr-xr-x 1 oracle dba          130 Apr 21 12:52 infra_stop
-rwxr-xr-x 1 oracle dba          473 Apr 19 19:52 start9ias904infra.sh
-rwxr-xr-x 1 oracle dba          527 Apr 21 12:32 stop9ias904infra.sh
```

MIDDLE TIER INSTALLATION

With your Infrastructure Tier active, you may proceed with the installation of your Middle Tier.

Oracle Universal Installer (OUI)

Prior to invoking the OUI, we set, performed or verified the following within our environment:

- created a "dba" group and an "oracle" user to own the Oracle software
- checked for port availability (`#/usr/bin/netstat -an | grep 1521`) Note: Portal and Wireless developer desire this port. If unavailable on your system, please refer to [Oracle Application Server 10g Installation Guide](#) for additional discussion on this topic.
- exported our `DISPLAY` to an X terminal
- executed the `rootpre.sh` on local node

1. Invoke the OUI

```
#!/runInstaller
```

- Provide the path to the location of your `jce-1_2_2.zip` file.
- Verify that you pass the installation checks and address any deficiencies.

Once you have passed all pre-installation checks, you should see a message similar to the following:

```
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2004-04-12_10-46-38PM.  
Please wait ...
```

2. Welcome

If you have performed an installation earlier that you wish to clean-up select "Deinstall Products" otherwise select "Next" to continue.

3. Specify Inventory Directory

If this is the first Oracle product installation on your machine, you will need to provide a full path for the inventory directory.

4. UNIX Group Name

If this is the first Oracle product installation on your machine, you will need to specify the name of the operating system group that will have permission to upgrade/install Oracle software on your system. We created a 'dba' group to own and manage Oracle software

5. Run `orainstRoot.sh`

If this is the first Oracle product installation on your machine, you will need to execute this shell script as the root user.

6. Specify File Locations

Specify a new name and path for your `$ORACLE_HOME`. The Middle Tier installation will require approximately 3GB of space.

Example:

```
Name: ias904Home  
Path: /oracle/product/904
```

7. Select a Product to Install

Select Oracle Application Server 10g and click "Next" to continue.

8. Select Installation Type

Select your Installation Type. We chose to install all Middle Tier components: "Business Intelligence and Forms (2.59GB)"

9. Preview of Steps for Middle Tier Installation

Click "Next" to continue.

10. Confirm Pre-Installation Requirements

Confirm that you have root user privileges, click the option box, and click "Next" to continue.

11. Select Configuration Options

Select the Oracle configuration assistants you wish to invoke during the install. We de-selected all products, as we intend to invoke the configuration assistants at a later time.

12. Register with Oracle Internet Directory

You will need to supply your Infrastructure Tier fully qualified hostname and port number to enable this new Oracle instance to register with an existing Internet Directory. You specify your VIP address here.

Example:

```
Host:   icc8xx4.sanmateo.ibm.com
Port:   389
```

13. Specify Login for Oracle Internet Directory

Supply your Internet Directory username and password.

14. Select OracleAS Metadata Repository

Identify your Infrastructure Tier as your Metadata Repository for this installation.

Example: `icc8xx4.sanmateo.ibm.com:1521:iasdb.sanmateo.ibm.com:iasdb...`

15. Specify Instance Name and ias_admin Password

Provide the new instance name and ias_admin password.

16. Summary

Click "Install" to continue.

17. Install -> Prompt to run root.sh

When prompted, run the `root.sh`. It will be located in your `$ORACLE_HOME`.

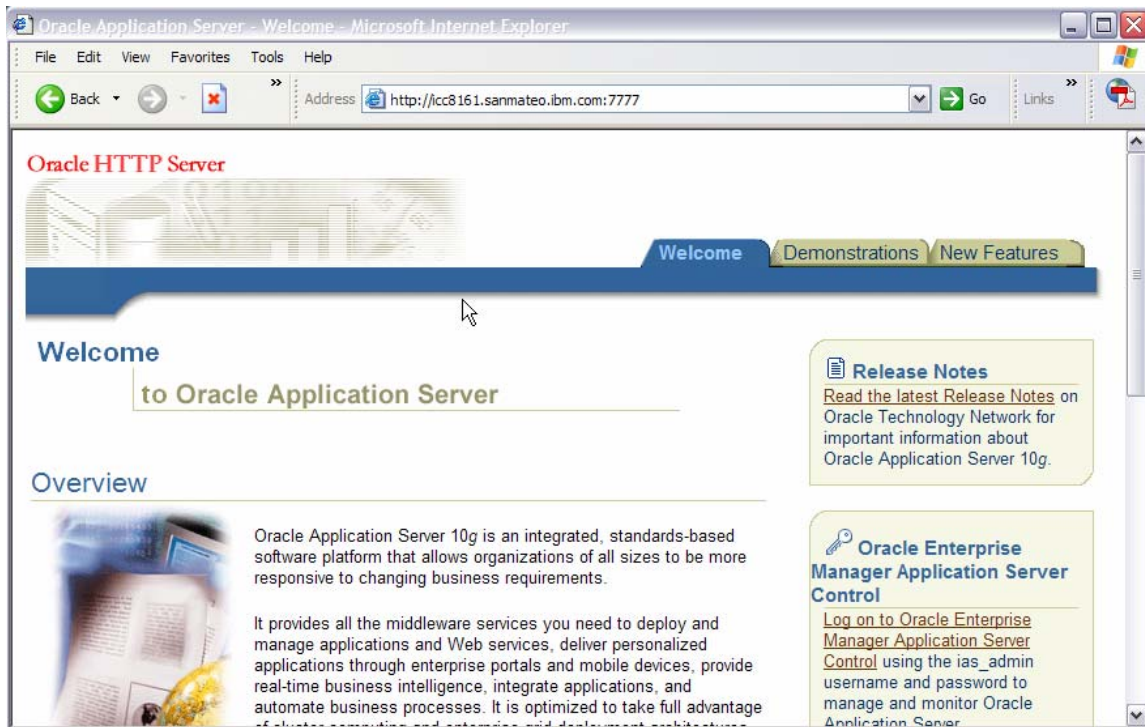
18. End of Installation

Click `Exit` to end the Installation process. Then we recommend you validate that you can access the HTTP Server and Welcome Page via the URL specified on the "End of Installation" screen: <http://icc8xx1.sanmateo.ibm.com:7777>

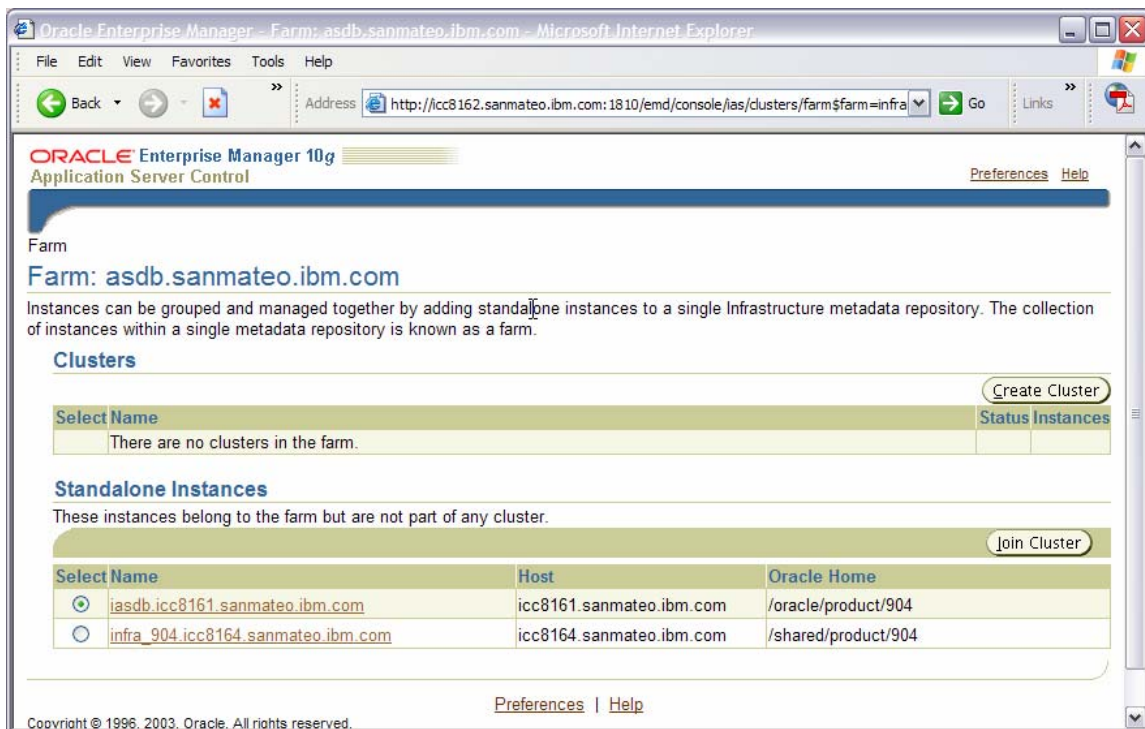
Next, validate you can connect to the Enterprise Manager console: <http://icc8xx2.sanmateo.ibm.com:1810>

You should see results similar to the following two screen snapshots:

HTTP Server:



EM Website:



SUMMARY

The installation of the Infrastructure Tiers and Middle Tier components for the Cold Failover Cluster solution on AIX 5L is complete. We recommend that you validate connectivity from your Middle Tier to both Infrastructure Tier nodes. Please follow the procedure you exercised in "[FAILOVER TEST FOR RESOURCE GROUP AND SERVICE IP](#)" on page 17 of this document to conduct a failover of the Infrastructure Tier. Verify that your Middle Tier can access the VIP from each Infrastructure Tier node. If you experience issues with Middle Tier connectivity following a failover, verify that all the Infrastructure Tier services have started on the takeover node, and review the "hacmp.out" log for errors.

Once you are comfortable with base functionality of your installation, please follow the recommendations in "[What to Do Next](#)" of "[Oracle Application Server 10g Installation Guide](#)" and prepare a backup of your installation. Specific information on backups and other Application Server administrative operations may be located in the "[Oracle Application Server 10g Administrators Guide](#)"

This concludes the topics for this document. Please refer to the listings in the "References" section for sources of additional information on questions that you may have.

REFERENCES

Oracle Application Server 10g Installation Guide 10g (9.0.4) for AIX-Based Systems, hp HP-UX PA-RISC (64-bit), Tru64 UNIX, and Linux x86 Part No. B13658-02

http://download-west.oracle.com/docs/cd/B12428_10/install.904/install.pdf

Oracle Application Server Documentation Library 10g for AIX-Based Systems:

http://download-west.oracle.com/docs/cd/B12428_09/docs.htm

Oracle Application Server 10g Release Notes 10g(9.0.4) for AIX-Based Systems Part No. B13666-03:

http://download-west.oracle.com/docs/cd/B12428_09/relnotes.904/relnotesax.pdf

Oracle Application Server 10g Quick Installation and Upgrade Guide 10g (9.0.4) for AIX-Based Systems Part No. B13662-01:

http://download-west.oracle.com/docs/cd/B12428_09/quickinstall.904/quickinstallax.pdf

Oracle Application Server 10g (9.0.4) Documentation:

<http://otn.oracle.com/documentation/appserver10g.html>

High Availability Cluster Multi-Processing (HACMP) for AIX Library:

http://www-1.ibm.com/servers/eserver/pseries/library/hacmp_docs.html

DISCLAIMER

This document reflects the IBM / Oracle International Competency Center's understanding about Oracle products running on IBM hardware servers and storage. It was produced and reviewed by the members of the IBM organization called the IBM / Oracle International Competency Center and others.

This document is presented "As-Is" and IBM does not assume responsibility for the statements expressed herein. It reflects the experiences of the IBM / Oracle International Competency Center. If you have questions about the contents of this document, please direct them to the IBM / Oracle International Competency Center (ibmoracl@us.ibm.com).

AUTHOR(s):

Tim Spencer, IBM Corporation,
2929 Campus Drive, San Mateo, CA
94403.
E-mail: spencert@us.ibm.com.

Nitin Sharma, IBM Corporation,
401 Island Parkway, Belmont, CA
94002.
E-mail: sharmani@us.ibm.com.



© IBM Corporation 2005
IBM Corporation
Marketing Communications
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
April 2005
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the e-business logo, @server, AIX, AIX 5L, HACMP, POWER and pSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark of The Open Group in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

The IBM home page on the Internet can be found at <http://www.ibm.com>.

The pSeries home page on the Internet can be found at <http://www.ibm.com/servers/eserver/pseries>.