

Configuring Highly Available Oracle  
Collaboration Suite with F5 BIG-IP®  
Application Traffic Manager

*Oracle Maximum Availability Architecture White Paper  
January 2006*

# Maximum Availability Architecture

Oracle Best Practices For High Availability

# Configuring Highly Available Oracle Collaboration Suite with F5 BIG-IP Application Traffic Manager

- Introduction ..... 2
- F5 BIG-IP Application Traffic Manager Terms..... 3
  - Pool..... 3
  - Virtual server ..... 3
  - Rule ..... 4
  - Monitor..... 4
- Oracle Collaboration Suite HA Architecture and Install Overview ..... 4
  - Architecture Overview ..... 4
  - Install Overview ..... 7
- Oracle Collaboration Suite Load Balancer Setup ..... 8
  - Prerequisites..... 8
  - Configure Load Balancer..... 9
    - Well Known Ports..... 9
    - Create Pools ..... 10
    - Create Rules ..... 11
    - Create Virtual Servers and Associate them with Pools or Rules..... 12
    - Create Monitors and Associate them with the Nodes ..... 13
    - Propagate Information to the Redundant BIG-IP..... 15
- Oracle Collaboration Suite High Availability Installation ..... 15
  - Pre-Installation Tasks..... 15
    - Validate that ports are not in use ..... 15
    - Static Ports Files ..... 16
    - Oracle Internet Directory Virtual Server Load Balancer Configuration..... 16
  - Installation Tasks ..... 17
  - Post-Installation Tasks..... 17
    - Verify Port Settings..... 17
    - Enable Oracle Internet Directory Traffic to Both Nodes ..... 17
    - Verify the Applications Service Registry..... 17
  - Validation Step ..... 18
- Appendix ..... 19
  - A. Static Ports..... 19
    - staticports.ini Template File..... 19
    - OID/DIP staticports.ini (static\_oid.ini) ..... 20
    - SSO/DAS staticports.ini (static\_sso.ini)..... 20
    - Applications staticports.ini (static\_apps.ini) ..... 20
  - B. Dumping Oracle Collaboration Suite Service Registry ..... 20
  - C. References..... 21

# Configuring Highly Available Oracle Collaboration Suite with F5 BIG-IP Application Traffic Manager

## INTRODUCTION

The availability of the collaboration system directly affects business processes, user productivity, and cost. Oracle Collaboration Suite 10g Release 1 (10.1.1) and Release 2 (10.1.2) is an integrated, standards-based collaboration solution. Oracle Collaboration Suite consists of different components that are deployed on multiple tiers. The availability of each component has a direct impact on the availability of the system.

A highly available Oracle Collaboration Suite deployment requires a highly available database, Infrastructure services, and Applications. Oracle Collaboration Suite 10g Release 1 (10.1.1) with a high availability foundation built on Oracle Database Real Application Clusters, Oracle Application Server 10g Release 2 (10.1.2.0.1) Cluster Identity Management, and multiple Oracle Collaboration Suite Applications nodes, provides High Availability (HA) architectures that are suited to different customer requirements. The primary Oracle Collaboration Suite High Availability architecture solutions are:

- Single Cluster
- Colocated Identity Management
- Distributed Identity Management

In all the preceding high availability deployments, a hardware load balancer distributes the incoming Oracle Internet Directory (OID) and Directory Integration & Provisioning (DIP), the Single Sign On (SSO) and Delegated Administrative Service (DAS), and Oracle Collaboration Suite Applications service requests across these simultaneously active instances. Failure of any one of the instances causes the load balancer to direct the subsequent requests to the remaining active instances. In addition, it's recommended that the F5 BIG-IP Application Traffic Manager be deployed redundantly for a more robust HA implementation.

This implies that the hardware load balancer is an integral part of the architecture and provides load balancing as well as failover capabilities. Also, many load balancers provide an SSL accelerator feature that can also be used to act as the Secure Socket Layer (SSL) proxy for HTTPS-based connections from SSO and DAS. The SSL acceleration feature is outside the scope of this paper. For high availability, the load balancer is always deployed redundantly.

This paper has been jointly written by Oracle Corporation and F5 Networks and describes the configuration and operational best practices for using F5 BIG-IP as

**The hardware load balancer is an integral component for providing high availability.**

**F5's BIG-IP provides the necessary load balancer features for Oracle Collaboration Suite high availability load balancing and monitoring**

the load balancer with an Oracle Collaboration Suite 10g HA Distributed Identity Management deployment.

### **F5 BIG-IP APPLICATION TRAFFIC MANAGER TERMS**

This document assumes that you are familiar with F5 BIG-IP. This section discusses the basic terminology to help with further discussion. For a detailed discussion of these terms, please refer to the BIG-IP Solutions Guide and the BIG-IP Reference Guide, see <http://tech.f5.com/home/bigip/manuals/-bigip>.

The version of BIG-IP kernel assumed for the rest of the discussion is BIG-IP Kernel 4.5.10 Build84. The version of iControl software assumed is 4.5.

#### **Pool**

A pool is a set of nodes grouped together to receive traffic according to a load balancing method. Members of the pool can be one or more machines (*node*) or they can be one or more node:port (a specific port on a node is also referred to as *node address*).

We generally advocate port-specific members in pools. Also note that if a non-port-specific virtual server is used then port translation will not happen on the pool layer, which can cause issues.

Each pool has its own characteristic for persistence definition and the load-balancing algorithm used. Certain types of applications may require the same client returning to the same pool member (node), this can be configured using a persistence setting on the pool. For ORACLE COLLABORATION SUITE, the DAS pool is the only pool that requires persistence to be configured.

Pools are associated with specific virtual servers directly or by rules (see later). As a result, the traffic coming to a virtual server is directed to one of the associated pools. After a pool receives a request, it directs the request to a member of the pool based on a chosen load balancing method. After a pool receives traffic, either directly from a virtual server or through a rule, the pool can optionally perform a number of different operations, such as inserting a header into an HTTP request, setting the Quality of Service or Type of Service level within a packet, or redirecting a request to a fallback destination.

#### **Virtual server**

A virtual server with its virtual address is the client addressable hostname or IP address through which nodes in a load balancing pool are made available to a client, either directly, or indirectly through a rule. Therefore a virtual server is the host name or IP address used by clients to access the devices that BIG-IP is load-balancing traffic for.

Before creating a virtual server, you must configure a load balancing pool of the actual physical devices you wish to forward the traffic to. You can then create the virtual server, specifying that pool as the destination for any traffic coming from

this virtual server. Also, if you want some of the traffic from that virtual server to go to multiple pools based on a pre-determined criteria, then you can create a rule specifying the criteria and BIG-IP would forward the traffic to a pool meeting the rule's criteria. A virtual server can also be configured to a specific port or to accept "ANY" ports.

A given load balancer device may contain one or more virtual servers.

### Rule

A rule is a user-written script that chooses among one or more load balancing pools. In other words, for an incoming request for a virtual server, a rule selects the pool to send it to. Therefore, rules allow a more granular level of control over traffic routing.

The specific pool selected by a rule for a request consists of one or more members.

### Monitor

Monitors are used to verify the state of a node or a node address. Monitors verify connections and services on nodes that are members of load-balancing pools. A monitor is designed to check the status of a node or service on an ongoing basis, at a set interval. If the node or service being checked does not respond within a specified timeout period, or the status of the node indicates that the performance of the node has degraded, the BIG-IP system automatically takes it out of the pool and redirects the traffic to the other members of the pool. When the node or service becomes available again, the monitor detects this and the node or service is automatically accessible to the pool.

## ORACLE COLLABORATION SUITE HA ARCHITECTURE and INSTALL OVERVIEW

### Architecture Overview

As mentioned earlier, there are three primary Oracle Collaboration Suite HA architectures that are described in detail in the *Oracle Collaboration Suite 10.1.2 HA Guide*<sup>1</sup>,

1. Single Cluster,
2. Colocated Identity Management, and
3. Distributed Identity Management.

For this paper we will describe the Oracle Collaboration Suite configuration with a load balancer with the Distributed Identity Management solution as depicted in Figure 1.

One point not clearly visible in Figure 1 is that the F5 BIG-IP Application Traffic Manager should also be redundant for a more robust HA implementation.

The load balancer configuration should not vary much between these different architectures other than the port numbers and the F5 BIG-IP pool members.

## Maximum Availability Architecture

Regardless of which Oracle Collaboration Suite HA architecture you use, the same general configuration of F5 virtual servers is required. However, the pool and pool member configurations will vary among the different architectures. The primary virtual server names required are:

1. ldap.mydomain.com
2. sso.mydomain.com
3. ocsapp.mydomain.com
4. ocsapp\_s2s.mydomain

The fourth virtual server, ocsapp\_s2s, is not depicted in Figure 1 but is required to workaround an issue that is documented in the Install guide under section 9.2.10.13. To quickly summarize this issue, certain Oracle Collaboration Suite Applications that communicate with each other need a special virtual server that has persistence set on it' BigIP pool. This complete setup is described in detail later.

Each of the “root” virtual server names can be associated with multiple ports. This paper will use this strategy. Optionally, you can use different virtual server names instead of the same virtual server root name with different ports.

As mentioned previously, the Oracle Collaboration Suite HA architectures are in the Oracle Collaboration Suite HA Guide and Install Guides.

For purposes of the Oracle Collaboration Suite HA configuration with the F5

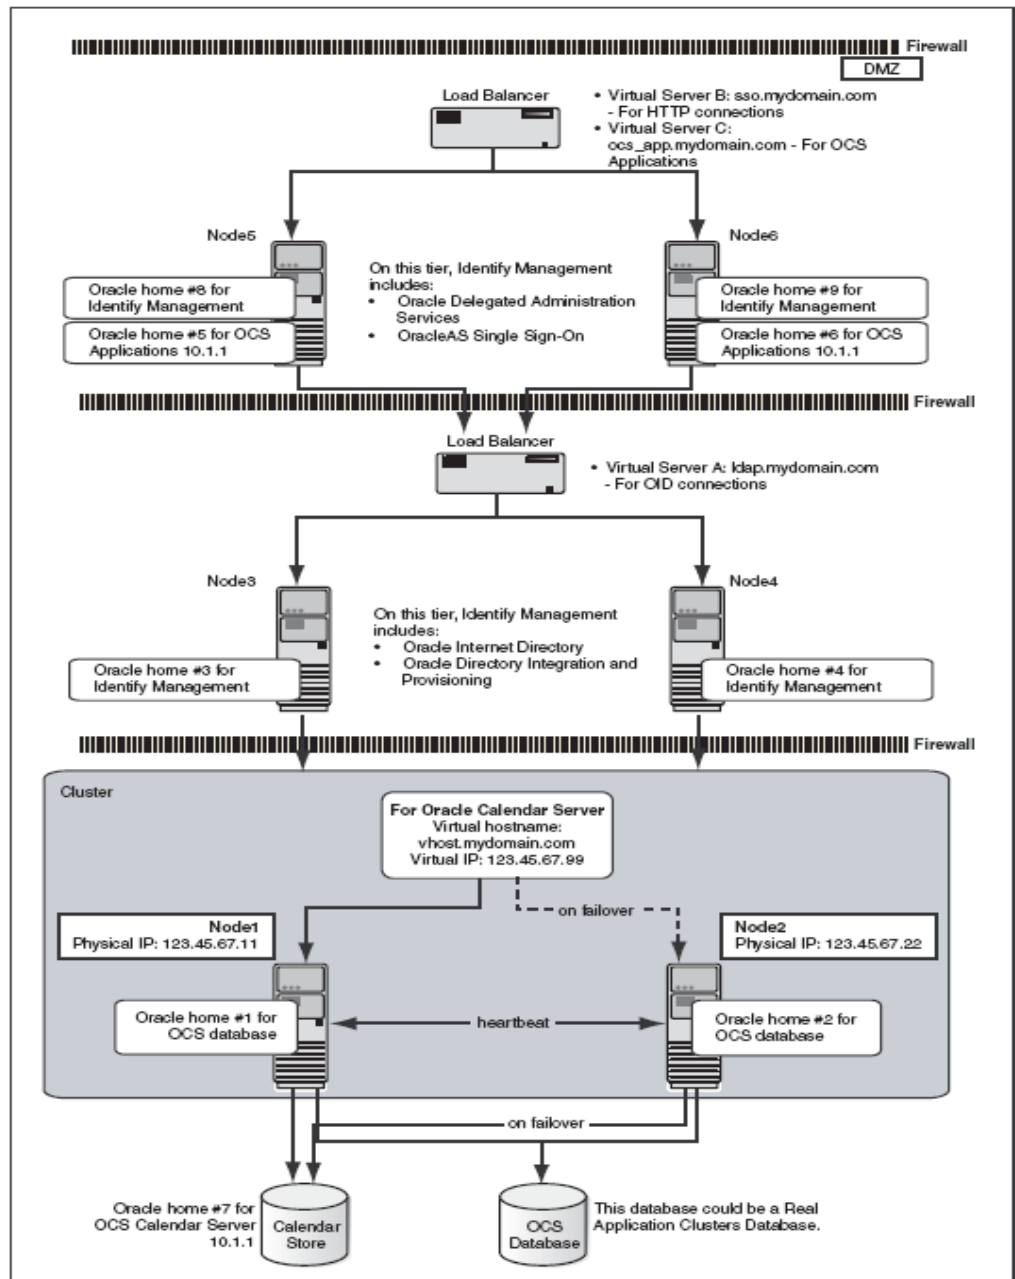


Figure 1 OCS Distributed Identity Management Deployment

BIG-IP, it's not necessary to discuss the Database tier and the Calendar Server-only tier because neither of those requires a load balancer. The focus of this paper will be on the Identity Management tier and the Oracle Collaboration Suite Applications tier, also called the Applications tier.

To quickly differentiate these 3 architectures, the Single Cluster install is just that, a single cluster that houses the complete Oracle Collaboration Suite installation on each node in the cluster. As a result, for this configuration, two nodes are load balanced against for each component.

The Colocated IM architecture separates the Collaboration Suite Database tier from the IM and Oracle Collaboration Suite Applications tiers rather than sharing nodes as in the Single cluster architecture. The IM tier has the OID/DIP and SSO/DAS components co-located on the same nodes in the same Oracle\_home's. This configuration has a total of four nodes.

Finally, the Distributed IM architecture is similar to the Co-located architecture except that it divides the OID/DIP component from the SSO/DAS component to separate nodes. As a result, two additional nodes are added and the OID/DIP IM components are segregated on a separate set of nodes from the SSO/DAS nodes that also house the Oracle Collaboration Suite Applications nodes.

For purposes of discussion and example, here is the high-level Distributed IM configuration of the physical nodes:

Server type	Hostname
Hosts for the database tier	d1.mydomain.com
Node1 and Node 2 in Figure 1	d2.mydomain.com
Hosts for the OID/DIP tier	o1.mydomain.com
Node3 and Node 4 in Figure 1	o2.mydomain.com
Hosts for SSO/DAS tier and Oracle Collaboration Suite Applications tier (separate Oracle home's)	s1.mydomain.com
Node5 and Node 6 in Figure 1	s2.mydomain.com

**Install Overview**

The load balancer setup must be done prior to the Oracle Collaboration Suite HA install. Then during the Oracle Collaboration Suite HA install you use the staticports.ini file that specifies the ports that will be used by the installer. In this manner you can ensure that the ports agree with the load balancer configuration and are consistent across multiple nodes.

Here is a high-level view of the steps involved in setting up an Oracle Collaboration Suite HA installation that will use the F5 BIG-IP Application Traffic Manager.

1. Complete the prerequisites for the installation

2. Configure the F5 BIG-IP configuration.
  - i. Create pools
  - ii. Create rules
  - iii. Create virtual servers and associate them with pools or rules
  - iv. Create monitors and associate them with the nodes
  - v. Propagate information to the redundant BIG-IP
3. Create the staticports.ini file, for each Oracle Collaboration Suite HA install step based on the load balancer configuration.
4. If a firewall separates your load balancer and the servers or cluster nodes, then ensure that appropriate ports are open for the two way traffic across the firewall. (This point is outside the scope of this paper)
5. Mark the non-install node or nodes down in the load balancer as required. This will be detailed in the HA Install section.
6. Perform the Oracle Collaboration Suite HA installs using the staticports.ini file that was created earlier.

The details of these steps follow.

## ORACLE COLLABORATION SUITE LOAD BALANCER SETUP

### Prerequisites

1. Decide on the virtual server names and ports (ensure the ports are free on the Oracle Collaboration Suite hosts). For this paper example we will use the following virtual server names:
  - ldap.mydomain.com
  - sso.mydomain.com
  - ocsapp.mydomain.comThe ports are summarized in Table 1
2. Get the IP addresses assigned to the virtual servers and ensure that they are part of your Domain Name Server (DNS).
3. Plan the load balancer configuration described in Table 1.

**Understanding the load balancer components, planning out the deployment, and walking through it are key to a successful implementation.**

<u>Virtual Server:port</u>	<u>Pool</u>	<u>Pool Nodes</u>	<u>Monitor</u>	<u>Purpose / staticports.ini setting</u>
ldap.mydomain.com:389	oid_pool	o1.mydomain.com:389 o2.mydomain.com:389	ldapm	IMHA OiD <i>Oracle Internet Directory port</i>
ldap.mydomain.com:636	oidssl_pool	o1.mydomain.com:636 o2.mydomain.com:636	ldapssl	IMHA OiD SSL <i>Oracle Internet Directory (SSL) port</i>
sso.mydomain.com:7777 <i>Note there are 2 pools that are routed to by a load balancer rule, we'll create a rule named sso_das_rule. The DAS pool has persistence set. See details in configuration steps.</i>	sso_pool das_pool	s1.mydomain.com:7777 s2.mydomain.com:7777	ssohttp	IMHA SSO / DAS <i>Oracle HTTP Server port</i>
ocsapp.mydomain.com:80	mt_app_pool	s1.mydomain.com:7778 s2.mydomain.com:7778	ocshttp	Applications mid-tier 7778 should match the <i>Oracle HTTP Server port</i>
ocsapp.mydomain.com:25	mt_smtp	s1.mydomain.com:25 s2.mydomain.com:25	ocssmtp	eMail SMTP port <i>Oracle Mail SMTP port</i>
ocsapp.mydomain.com:143	mt_imap	s1.mydomain.com:143 s2.mydomain.com:143	ocsimap	eMail IMAP port <i>Oracle Mail IMAP4 port</i>
ocsapp.mydomain.com:110	mt_pop	s1.mydomain.com:110 s2.mydomain.com:110	ocspop	eMail POP3 port <i>Oracle Mail POP3 port</i>
ocsapp.mydomain.com:9401	mt_wci	s1.mydomain.com:9401 s2.mydomain.com:9401	tcp	Webcache Invalidation <i>Web Cache Invalidation port</i>
ocsapp.mydomain.com:7778	mt_wclsnr	s1.mydomain.com:7778 s2.mydomain.com:7778	N/A	Webcache HTTP Listener <i>Web Cache HTTP Listen port</i> (this is the same as the <i>Oracle HTTP Server port</i> , it can't be different. when using staticports.ini) Same nodes are monitored by the ocshttp monitor.
ocsapp_s2s.mydomain.com:80	mt_s2s_pool	s1.mydomain.com:7779 s2.mydomain.com:7779	ocshttp	Applications mid-tier Service-to-service workaround <i>Oracle HTTP Server Listener port</i>

**Table 1 Load Balancer Configuration Summary**

### Configure Load Balancer

The following subsections describe the steps to configure a load balancer in detail.

#### Well Known Ports

Note that well-known port numbers get set to their textual well-known service name as listed in the `BigIP /etc/services` file. Table 2 summarizes the well-known port service labels that are set. These port labels are automatically set for pools and virtual servers.

<u>Port Number</u>	<u>Service Label</u>
389	ldap
25	smtp
143	imap2 or imap
110	pop

**Table 2 Port Service Labels**

**Create Pools**

To create a new pool using the BIG-IP configuration tool, connect to the active device of the redundant load balancer configuration and click **Pools** and then click the **+**. Each pool has to be created separately. The characteristics of these pools are described in the Table 3.

<u>Pool Name</u>	<u>Pool Members</u>	<u>Persistence</u>
oid_pool	o1.mydomain.com:389 o2.mydomain.com:389	No persistence
oidssl_pool	o1.mydomain.com:636 o2.mydomain.com:636	No persistence
sso_pool	s1.mydomain.com:7777 s2.mydomain.com:7777	No persistence
das_pool	s1.mydomain.com:7777 s2.mydomain.com:7777	Active HTTP cookie Method: Insert Expiration: Null
mt_app_pool	s1.mydomain.com:7778 s2.mydomain.com:7778	No persistence
mt_smtp	s1.mydomain.com:25 s2.mydomain.com:25	No persistence
mt_imap	s1.mydomain.com:143 s2.mydomain.com:143	No persistence
mt_pop	s1.mydomain.com:110 s2.mydomain.com:110	No persistence
mt_wci	s1.mydomain.com:9401 s2.mydomain.com:9401	No persistence
mt_wclsnr	s1.mydomain.com:7778 s2.mydomain.com:7778	No persistence
mt_s2s_pool	s1.mydomain.com:7779 s2.mydomain.com:7779	Active HTTP cookie Method: Insert Expiration: Null

**Table 3 Load Balancer Pool Summary**

In addition, for each pool the following should be enabled (they are enabled by default),

- Enable SNAT
- Enable NAT

After you create the pool, set the persistence by selecting the Persistence tab under the pool screen as shown in Figure 2.

The screenshot shows the configuration interface for the persistence of a pool named 'das\_pool'. The 'Persistence' tab is active. Under 'Persistence Type', the 'Active HTTP Cookie' radio button is selected and circled in blue. The 'Method' dropdown menu is also circled in blue and set to 'Insert'. Other persistence types include None, SSL (with a 300 sec timeout), SIP (with a 32 sec timeout), Simple (with a 0 sec timeout and a mask field), Destination Address Affinity (with a mask field), Expression (with a 0 sec timeout and an expression field), and Passive HTTP Cookie. The 'Cookie Name' field is empty at the bottom.

Figure 2 Setting Persistence Screenshot

**Create Rules**

To create a rule, click Rules and then click the + to add a new rule. Figure 3 shows the SSO/DAS rule.

```

if (http_uri starts_with "/oiddas/") {
    use pool das_pool
}
else {
    use pool sso_pool
}
    
```

Figure 3 SSO\_DAS\_rule example

**Create Virtual Servers and Associate them with Pools or Rules**

Create the virtual servers and associate each with its respective pool or rule.

To create a virtual server, click on **Virtual Servers** and then click on **+** to add a new virtual server.

**Note** that the Oracle Collaboration Suite Applications tier virtual server, ocsapp.mydomain.com:80, uses a different port than it’s associated pool nodes do. This provides standard client HTTP port access, port 80, without requiring the Oracle Collaboration Suite Applications tier to use a privileged port (< 1024) on the Oracle Collaboration Suite Applications nodes. The default entries for other virtual servers are fine for creating the virtual server.

There are three primary steps for creating a virtual server:

1. Configure Virtual IP Address and Service – Here, you enter the virtual host name and the port (service).
2. Configure Basic Properties – Here, accept the default settings and do not change anything.
3. Select Physical Resources – Here, select the pool or rule from the drop down lists.

There are also 2 other optional steps, “Configure Redundant Properties” and “Configure Outbound Properties” which we did not configure for any of the pools.

Address	Pool	Rule
ldap.mydomain.com:389	oid_pool	Not applicable
ldap.mydomain.com:636	oidssl_pool	Not applicable
sso.mydomain.com:7777	Not applicable	sso_das_rule
ocsapp.mydomain.com:80	mt_app_pool	Not applicable
ocsapp.mydomain.com:25	mt_smtp	Not applicable
ocsapp.mydomain.com:143	mt_imap	Not applicable
ocsapp.mydomain.com:110	mt_pop	Not applicable
ocsapp.mydomain.com:9401	mt_wci	Not applicable
ocsapp.mydomain.com:7778	mt_wclsnr	Not applicable
ocsapp_s2s.mydomain.com:80	mt_s2s_pool	Not applicable

**Table 4 Load Balancer Pool Rule Summary**

For the virtual servers related to the OID service - ldap.mydomain.com:389 and ldap.mydomain.com:636 - ensure that **TCP Enabled** is selected and set **Idle Connection Timeout TCP (seconds)** to a very large value (for example. 345600). This is available on the **Virtual Services Properties** tab for a virtual server as shown in Figure 4.

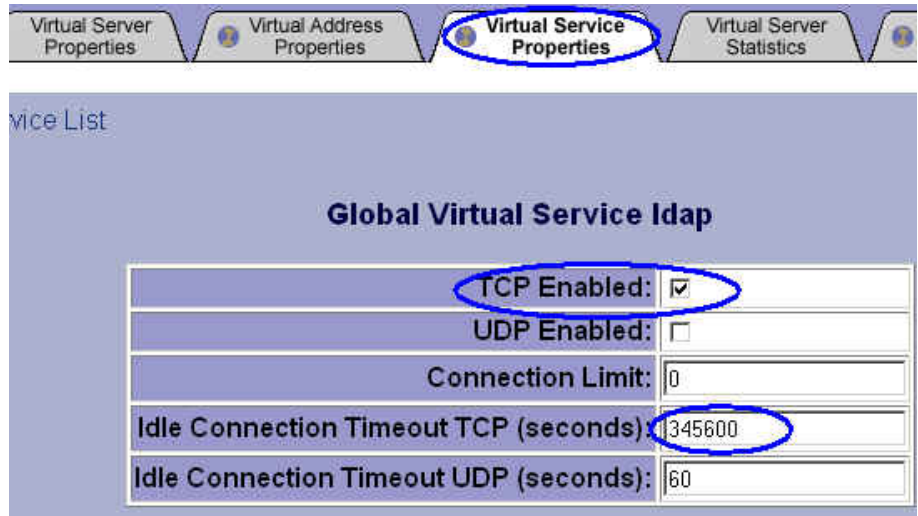


Figure 4 Virtual Server - Virtual Service Properties

**Create Monitors and Associate them with the Nodes**

Create the monitors listed in Table 5.

To create a monitor, click **Monitors** and then click + to add a new monitor.

Monitor name	Configuration
ldapm	Inherit from LDAP <b>Interval:</b> 20 <b>Timeout:</b> 61 <b>Username:</b> <a username full directory name (DN) > <b>Password:</b> <username password> <b>Filter:</b> cn= <i>datasename</i> <i>Note: It is recommended that a dedicated account be used to monitor the LDAP service to prevent operational conflicts with other uses of the account. In particular, administrative accounts such as orcladmin should <b>not</b> be used for the username.</i>
ldapssl	Inherit from tcp <b>Interval:</b> 20 <b>Timeout:</b> 61
ssohttp	Inherit from http <b>Interval:</b> 20 <b>Timeout:</b> 61 <b>Send String:</b> GET /sso/status <b>Receive Rule:</b> OC4J_Security is running
ocshttp	Inherit from http <b>Interval:</b> 20 <b>Timeout:</b> 61
ocssmtp	Inherit from smtp <b>Interval:</b> 30 <b>Timeout:</b> 91 <b>Domain:</b> <Your email domain>

Monitor name	Configuration
ocsimap	Inherit from imap <b>Interval:</b> 30 <b>Timeout:</b> 91 <b>Username:</b> orclguest (or some other health check user) <b>Password:</b> < password for above user> <b>Folder:</b> INBOX
ocspop	Inherit from pop3 <b>Interval:</b> 30 <b>Timeout:</b> 91 <b>Username:</b> orclguest (or some other health check user) <b>Password:</b> < password for above user>

**Table 5 Monitor Summary**

For the *ldapm* monitor, it is recommended that a dedicated account be used to monitor the LDAP service to prevent operational conflicts with other uses of the account. In particular, administrative accounts such as orcladmin should not be used for the monitor username. The username field should be similar to:

```
cn=ldapmUser ,cn=Users ,dc=mydomain ,dc=com
```

where *ldapmUser* is an Oracle Collaboration Suite account that was provisioned with minimum privileges. The validity of the user DN can be verified at the operating system level by executing an *ldapbind* command for the user DN as follows:

```
ldapbind -h ldap.mydomain.com -p 389 -D \  
"cn=ldapmUser ,cn=Users ,dc=mydomain ,dc=com" -w welcome1
```

The interval and timeout for the monitors should be adjusted according to your requirements.

**Interval** is the frequency at which BIG-IP pings the service and **timeout** is the maximum time it waits each time before determining whether the service is down.

A low interval time implies frequent pings but faster automatic failover in case of the service going down.

The timeout value should be a minimum of  $interval * 3 + 1$ . For slow backend servers or servers with higher load, it should be adjusted higher to prevent false alarms.

**Note:** The Oracle Process Manager and Notification server (OPMN) process monitors the application server component processes and restarts them. Therefore, it is important that the interval and timeout values here should work well with the ping and restart timeout values specified in *\$ORACLE\_HOME/opmn/conf/opmn.xml* for a component. The values in Table 5 are the recommended default values.

After the monitors have been created, associate the monitors to the nodes (using the **Node Associations** tab) as Table 6 indicates.

<b>Monitor Name</b>	<b>Nodes</b>	<b>Purpose</b>
ldapm	o1.mydomain.com:389 o2.mydomain.com:389	<i>Oracle Internet Directory port</i>
ldapssl	o1.mydomain.com:636 o2.mydomain.com:636	<i>Oracle Internet Directory (SSL) port</i>
ssohttp	s1.mydomain.com:7777 s2.mydomain.com:7777	IMHA SSO / DAS <i>Oracle HTTP Server port</i>
ocshttp	s1.mydomain.com:7778 s2.mydomain.com:7778	Applications mid-tier <i>Oracle HTTP Server port</i>
ocshttp	s1.mydomain.com:7779 s2.mydomain.com:7779	Applications mid-tier Service-to-service workaround <i>Oracle HTTP Server Listener port</i>
ocssmtp	s1.mydomain.com:25 s2.mydomain.com:25	<i>Oracle Mail SMTP port</i>
ocsimap	s1.mydomain.com:143 s2.mydomain.com:143	<i>Oracle Mail IMAP4 port</i>
ocspop	s1.mydomain.com:110 s2.mydomain.com:110	<i>Oracle Mail POP3 port</i>
tcp (system supplied)	s1.mydomain.com:9401 s2.mydomain.com:9401	<i>Web Cache Invalidation port</i>

**Table 6 Monitor Node Association Summary**

**Propagate Information to the Redundant BIG-IP**

Because a redundant load balancer is highly recommended for the deployment, the preceding configuration done performed on the active load balancer should be propagated to the standby load balancer in the redundant configuration. To do so using the BIG-IP Configuration Utility, click **Redundant Properties** on the home page and then click **Synchronize Configuration**.

This will propagate the newly created configuration to the redundant load balancer, which will then be ready to service the new configuration in the event of a failure of the active load balancer.

**ORACLE COLLABORATION SUITE HIGH AVAILABILITY INSTALLATION**

Here, we describe the pre-install and install steps common to the various HA deployments. These steps are relevant to the BIG-IP Application Traffic Manager usage in this configuration. For other detailed install steps refer to the Oracle Collaboration Suite10g Installation Guides.

**Pre-Installation Tasks**

The following sections describe the pre-installation steps for installing Oracle Collaboration Suite in high availability environment.

**Validate that ports are not in use**

Before proceeding with the installation and using the load balancer ports as previously described, you should ensure that the ports are free on the appropriate

nodes. This can be done by using the `netstat` command or by verifying with your network or system administrator. A simple `netstat` command to verify port 7777 is not in use would be as follows:

```
netstat -a | grep 7777
```

The preceding command should not return any line in response. If it returns a line with a “tcp ... \*:7777 LISTEN” then the queried port is in use.

**Static Ports Files**

Instead of using default ports, you can assign custom port numbers for Oracle Collaboration Suite components during the installation. For this, you must create a file containing the component names and port numbers. This file is referred to as the static ports file or `staticports.ini`.

The static ports feature of Oracle Universal Installer (OUI) ensures that the only specific ports will be used for the install. However, for this, these ports must be free on all the relevant nodes. Your planning process should take this into account while deciding the various ports.

Please refer to the [Appendix A](#) for the template and sample files that go with the install steps.

**Oracle Internet Directory Virtual Server Load Balancer Configuration**

A requirement for the OID/DIP install is to point the load balancer to only one OID node during the install. This is because the load balancer must direct traffic to only the first node until all OID nodes are installed.

To disable traffic to non-install node(s), click **Nodes** and then click the non-install node(s) and perform the following tasks,

1. Disable “*Enable Session*” for the non-install OID node `o2.mydomain.com:389` and the SSL node `o2.mydomain.com:636`
2. Disable “*Enable Connections*” for the non-install OID node `o2.mydomain.com:389` and the SSL node `o2.mydomain.com:636`

Following this, each node window will have a “**FORCED DOWN**” status as illustrated in Figure 5.

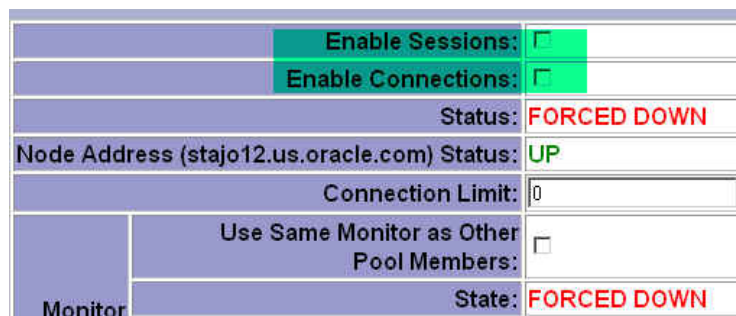


Figure 4 F5 BIG-IP Screen Snapshot of Node

### Installation Tasks

During the installation one or more installation sessions may be required. In each case, make sure that you start the installation process with the arguments required for it to use the correct staticports.ini file.

Table 7 summarizes the installation step and the command to use to initiate the installation process.

<u>Static Ports File In /test directory</u>	<u>Install Step</u>	<u>Install Command</u>
static_oid.ini	OID / DIP	<code>./runInstaller oracle.ocs.infrastructure:s_staticPorts=/test/<b>static_oid.ini</b></code>
static_sso.ini	SSO / DAS	<code>./runInstaller oracle.ocs.infrastructure:s_staticPorts=/test/<b>static_sso.ini</b></code>
static_apps.ini	Applications	<code>./runInstaller oracle.ocs.midtier:s_staticPorts=/test/<b>static_apps.ini</b></code>

**Table 7 Installation Steps and Commands**

### Post-Installation Tasks

The post-installation tasks described here are relevant only to the BIG-IP Application Traffic Manager usage in this configuration and common to all installations. For other detailed post-install steps refer to the Oracle Collaboration Suite 10g Installation Guide and the Oracle Collaboration Suite 10g Release Notes.

#### Verify Port Settings

After each installation, verify the port settings match your load balancer configuration. To verify, check the `$ORACLE_HOME/install/portlist.ini` file.

#### Enable Oracle Internet Directory Traffic to Both Nodes

Following the installation of both OID/DIP and SSO/DAS, enable OID traffic to all nodes. Click **Nodes** and then click the “FORCED DOWN” nodes and do the following for each:

1. Enable “*Enable Session*” for the non-install OID nodes o2.mydomain.com:389 and the SSL node o2.mydomain.com:636
2. Enable “*Enable Connections*” for the non-install OID nodes o2.mydomain.com:389 and the SSL node o2.mydomain.com:636

#### Verify the Applications Service Registry

As a part of the Applications tier post-installation steps the Oracle Collaboration Suite Service Registry in the Oracle Internet Directory has to be modified to contain the correct load balancer URL’s as detailed in the Oracle Collaboration Suite Install Guide.

Verify that the URL's are set correctly as described under the "Update the Oracle Collaboration Suite Service Registry Entries in Oracle Internet Directory to Use the Load Balancer" section and the "Configure Applications Tier Service-to-Service Operations with a Dedicated Load Balancer Virtual Server" section.

There is a sample script to dump this portion of OID contained in [Appendix B](#).

### **Validation Step**

Perform the following tasks to validate if the installation was successful:

1. Access `http://sso.mydomain.com:7777/oiddas` multiple times and validate that everything is working.
2. Access `http://sso.mydomain.com:7777/pls/orasso` multiple times and validate that everything is working.
3. Access `http://ocsapp.mydomain.com` multiple times and validate that everything is working.

## APPENDIX

### A. Static Ports

#### staticports.ini Template File

This template can be found on DVD 1, in the response directory

```
#Oracle HTTP Server Listen port = port_num
#Oracle HTTP Server SSL port = port_num
#Oracle HTTP Server Listen (SSL) port = port_num
#Oracle HTTP Server Diagnostic port = port_num
#ASG port = port_num
#Application Server Control port = port_num
#Application Server Control RMI port = port_num
#Java Object Cache port = port_num
#Log Loader port = port_num
#DCM Discovery port = port_num
#Oracle Notification Server Request port = port_num
#Oracle Notification Server Local port = port_num
#Oracle Notification Server Remote port = port_num
#Oracle Management Agent Port = port_num

# Ports specific to Infrastructure install

#Oracle Internet Directory port = port_num
#Oracle Internet Directory (SSL) port = port_num
#Enterprise Manager Console HTTP Port = port_num
#Enterprise Manager Agent Port = port_num

# Ports specific to Applications install

#Web Cache HTTP Listen port = port_num
#Web Cache HTTP Listen (SSL) port = port_num
#Web Cache Administration port = port_num
#Web Cache Invalidation port = port_num
#Web Cache Statistics port = port_num
#Oracle Net Listener = port_num
#Oracle Mail IMAP4 port = port_num
#Oracle Mail IMAP4 Secure port = port_num
#Oracle Mail POP3 port = port_num
#Oracle Mail POP3 Secure port = port_num
#Oracle Mail SMTP port = port_num
#Oracle Mail NNTP port = port_num
#Oracle Mail NNTP Secure port = port_num
#Oracle Calendar server = port_num
#Oracle Calendar server manager (CSM) = port_num
#Wireless PIM Notification Dispatcher = port_num
#Wireless PIMAP UDP Dispatcher = port_num
#RTC redirector Server port=port_num
#RTC redirector MX port=port_num
#RTC redirector XMPP port=port_num
#RTC redirector Secure XMPP port=port_num
#RTC process monitor port=port_num
#RTC messenger directory server first port=port_num
#RTC messenger directory server second port=port_num
```

```
#RTC messenger multiuser chat port=port_num
#RTC messenger connection manager port=port_num
#RTC messenger statistics collection port=port_num
#RTC messenger server to server connection port=port_num
#RTC messenger group service port=port_num
#RTC messenger voice proxy listener port=port_num
```

**OID/DIP staticports.ini (static\_oid.ini)**

```
Oracle Internet Directory port = 389
Oracle Internet Directory (SSL) port = 636
```

**SSO/DAS staticports.ini (static\_sso.ini)**

```
Oracle HTTP Server port = 7777
```

**Applications staticports.ini (static\_apps.ini)**

```
#
# "Oracle HTTP Server port" will get set to the "Web
# Cache HTTP Listen port"
#Oracle HTTP Server port = 7778
Oracle HTTP Server Listen port = 7779
# eMail
Oracle Mail IMAP4 port = 143
Oracle Mail POP3 port = 110
Oracle Mail SMTP port = 25
# Webcache
Web Cache HTTP Listen port = 7778
Web Cache Invalidation port = 9401
```

**B. Dumping Oracle Collaboration Suite Service Registry**

This script dumps the OID Virtual Services (Service Registry) section of the Oracle Collaboration Suite OID metadata. You must validate these results against the instructions in the Oracle Collaboration Suite Installation Guide.

```
ldapsearch -h ldap.mydomain.com -p 389 -s sub -v \
"(&(orclServiceType=*)(objectclass=orclVirtualService))"
```

## C. References

1. *Oracle Collaboration Suite 10.1.2 Documentation*  
<http://www.oracle.com/pls/cs101/homepage>
2. *Oracle Maximum Availability Architecture*  
<http://www.oracle.com/technology/ deploy/availability/htdocs/maa.htm>
3. *Oracle Collaboration Suite 10.1.2 Installation Guide for Linux*  
[http://download-west.oracle.com/docs/cd/B25553\\_01/install.1012/b25465/toc.htm](http://download-west.oracle.com/docs/cd/B25553_01/install.1012/b25465/toc.htm)
4. *Oracle Collaboration Suite 10.1.2 High Availability Guide*  
[http://www.oracle.com/pls/cs101/to\\_toc?pathname=collab.1012%2Fb25481%2Ftoc.htm&remark=portal+%28Books%29](http://www.oracle.com/pls/cs101/to_toc?pathname=collab.1012%2Fb25481%2Ftoc.htm&remark=portal+%28Books%29)
5. *BIG-IP Reference and Solution Manuals*  
<http://tech.f5.com/home/bigip/manuals/index.html#bigip>
6. *Configuring BigIp for Oracle Application Server High Availability*  
[http://www.oracle.com/technology/products/ias/hi\\_av/BigIP.pdf](http://www.oracle.com/technology/products/ias/hi_av/BigIP.pdf)
7. *Tested Load Balancers with Oracle Application Server 10g*  
[http://www.oracle.com/technology/products/ias/hi\\_av/Tested\\_LBR\\_FW\\_SSLLAcceler.html](http://www.oracle.com/technology/products/ias/hi_av/Tested_LBR_FW_SSLLAcceler.html)



Configuring Highly Available Oracle Collaboration Suite with F5 BIG-IP Application Traffic Manager, January 2006  
Author: Ray Dutcher, Oracle HA Systems Group; Randy Cleveland, F5 Networks  
Contributing Authors: Susan Kornberg and Pradeep Bhat, Oracle HA Systems Group; Mike Schrock, F5 Networks

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

Copyright © 2006, Oracle. All rights reserved.  
This document is provided for information purposes only and the contents hereof are subject to change without notice.  
This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.  
Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.