DEPLOYMENT GUIDE
Version 1.0

Deploying the BIG-IP LTM with Oracle Coherence*Extend

ORACLE WEBLOGIC READY

Table of Contents

Deploying the BIG-IP LTM with Oracle Coherence*Extend

Welcome to the F5 deployment guide for the BIG-IP Local Traffic Manager (LTM) and Oracle Coherence*Extend™. This guide describes how to configure the BIG-IP LTM for Oracle Coherence*Extend Proxy services when you are looking to create optimized Coherence client connections at a remote site.

Oracle Coherence is the industry's leading in-memory grid architecture, enabling companies to scale mission-critical applications by providing fast and reliable access to frequently used information. As a shared memory infrastructure, Oracle Coherence enables real-time data analysis, in-memory grid computations, parallel transaction and event processing, and grid computing. Expanding a Coherence cluster to a remote location requires the use of Coherence*Extend Proxy services.

For more information on Oracle Coherence*Extend, see ***http://www.oracle.com/us/products/middleware/coherence/index.html***.

The BIG-IP LTM, when installed at the remote location, provides high availability, load balancing, connection management, SSL offload, and improved response time for Coherence clients.

For more information about the BIG-IP system, visit ***http://www.f5.com/products/big-ip/***.

Using these technologies from F5 and Oracle together can provide enterprise class grid computing services for mission critical data.

Prerequisites and configuration notes

The following are prerequisites and configuration notes for this deployment:

- You must have 2 identical BIG-IP LTMs, configured for High Availability (recommended).
- You must be running BIG-IP TMOS software version 10.2.1 or later, and be running the same version on each unit.
- You must have administrative access to the Web management interface on the BIG-IP.
- You must be running Oracle Coherence version 3.7 or later.
- You must have administrative access to the Oracle Coherence servers, to be able to edit and control services.
- This guide assumes you have already physically installed the BIG-IPs, configured at least one VLAN and Self-IP, and configured the BIG-IP pair for High Availability. For information on how to do these tasks, see BIG-IP documentation on AskF5.com. It is assumed that the BIG-IP is on the same VLAN and IP subnet as the Oracle Coherence*Extend Proxy servers.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP LTM	v10.2, 10.2.1
Oracle Coherence	3.7

Document Version	Description
1.0	New guide

Configuration example

In this deployment guide, we use an example of a single Coherence cluster, with three Proxy servers, connected to a pair of LTMs in High Availability mode, with Coherence clients connected to the LTMs at the remote site. We show how to configure the BIG-IP LTM software to create SSL secured connections to the clients, and load balance the traffic to the Proxy servers.

The following is a simple, logical diagram of the network, servers, clients, and BIG-IP LTM in our example.

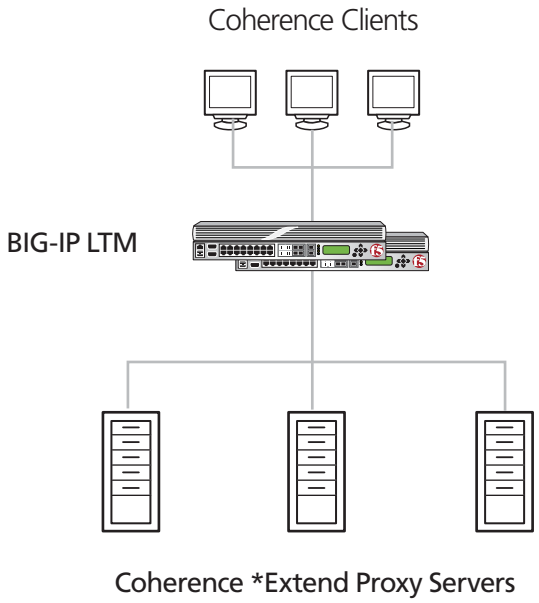

Figure 1 *Example Configuration*

Configuring the BIG-IP LTM

Use the following procedures for configuring the BIG-IP for Oracle Coherence*Extend.

Configuring the health monitor

The first task is to create the health monitors. Monitors are optional, but we highly recommend using the monitors below to verify that the nodes and services are available.

For this configuration, the type of monitor you create depends on the version of the BIG-IP LTM you are using:

- If you are using version 10.2.1 or later, create the monitor in *Creating the TCP monitor for the Proxy service (v10.2.1 and later)* on this page.
- If you are using 10.2 or earlier, create the monitor in *Adding an External monitor for the Proxy service*

Creating the TCP monitor for the Proxy service (v10.2.1 and later)

You can use a basic TCP health monitor to monitor the pool of Proxy servers. This type of monitor marks a proxy server up if the BIG-IP device is able to establish a TCP/IP connection with the server. While this is a generally good indication that a proxy server is functional, it does not guarantee that the proxy server can actually process client traffic.

The BIG-IP version 10.2.1 and later allows you to create a custom TCP health monitor that sends a "Coherence*Extend Ping Request" to a proxy server and validate that an appropriate response is returned. This ensures the proxy server is up and able to process client traffic.

Important

You must be using BIG-IP software version 10.2.1 and later for this monitor. If using an earlier version of BIG-IP software, please go to the next section "Adding an External monitor for the Proxy service."

To create the custom TCP monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
   In our example, we type **Extend_TCP_monitor**.
4. From the **Type** list, select **TCP**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a (1:3) +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.

6. In the **Send String** box, type the following binary string the BIG-IP will send the proxy to verify the service is running. You must type this exact string.

   \x07\x00\x03\x00\x00\x42\x00\x40

   This binary string is a Coherence*Extend "ping" message that checks to see if the service is able to process requests.

7. In the **Receive String** box, type the binary string the BIG-IP should expect to receive from the proxy when the service is running. The following binary string is the response expected from the binary Send String in step 6. You must type this exact string:

   \x09\x00\x04\x02\x00\x42\x00\x03\x64\x40

8. Click the **Finished** button.

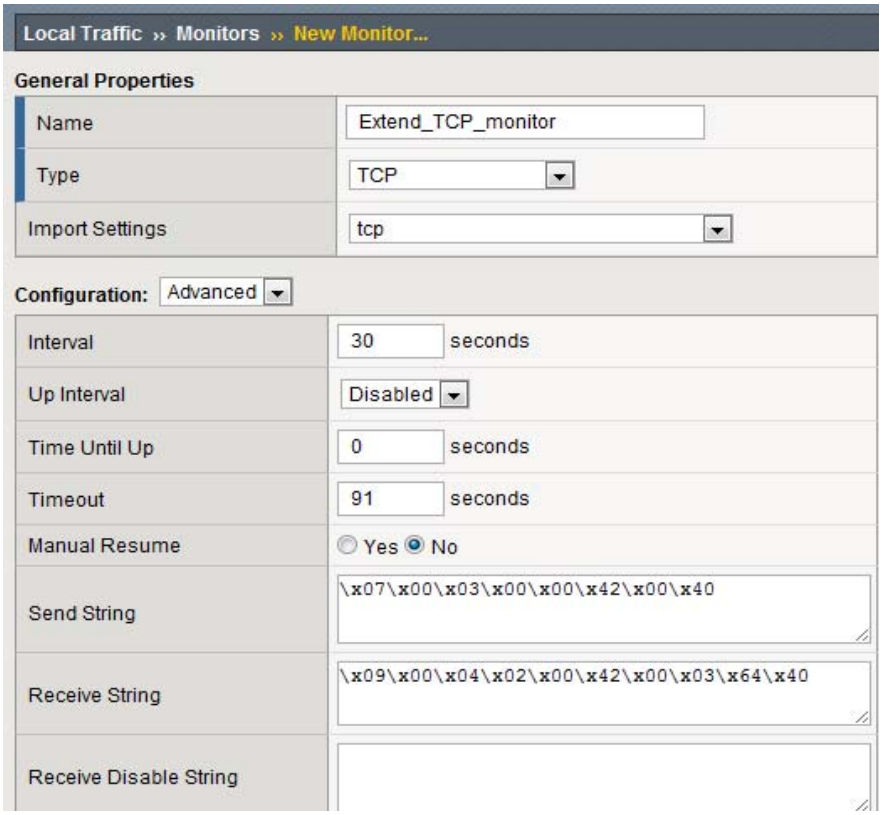

Figure 2 *TCP monitor configuration - truncated to show relevant settings*

Adding an External monitor for the Proxy service

To create the monitor for the Proxy service and BIG-IP LTM version 10.2 and earlier, you must configure an external health monitor that runs a shell script. This is a two part process, first you create the script called extend_ping in the /usr/bin/monitors directory of the BIG-IP device, and then you create the monitor that calls the script.

Important

This monitor is for BIG-IP version 10.2.0 and earlier only.

To create and import the script on the BIG-IP LTM

1. Copy and paste the following script into a file. Save the file as **extend_ping** in a location accessible by the BIG-IP system.

```
#! /bin/bash
################################################################################
###  EXTERNAL MONITOR FOR COHERENCE*EXTEND
###   INPUTS:
###     $1    The IP address of the pool member to test
###     $2    The port number of the pool member to test
###     $3+   White space delimited parms as listed in the monitor "args"
###   OUTPUTS:
###     If null is returned, the member is "down"
###     If any string of one or more characters is returned, the member is "up"
################################################################################

IP=${1:-"127.0.0.1"}
IP=${IP##*:} # This removes the leading ::ffff:
PORT=${2:-"80"}
TIMEOUT=${3:-1}
SLEEP=${4:-1}

PID_FILE="/var/run/extend_ping.$IP.$PORT.pid"
HEX_REQUEST="0700030000420040"
HEX_RESPONSE="09000402004200036440"

###
### Terminate existing process, if any
###
if [ -f $PID_FILE ]
then
   kill -9 `cat $PID_FILE` > /dev/null 2>&1
fi
echo "$$" > $PID_FILE

###
###  Ping the server and return a user friendly result
###
RESULT=`/bin/echo "$HEX_REQUEST" | /usr/bin/xxd -r -p | /usr/bin/nc -i \
    $SLEEP -w $TIMEOUT $IP $PORT | /usr/bin/xxd -p | /bin/grep \
    "$HEX_RESPONSE" 2> /dev/null`

if [ "$RESULT" != "" ] ; then
   /bin/echo "$IP:$PORT is \"UP\""
fi

rm -f $PID_FILE
```

2. Open a secure copy program. In this example, we use WinSCP, which is available as a free download from http://winscp.net/. For more options, see *http://support.f5.com/kb/en-us/solutions/public/0000/100/sol175.html*
3. In the **Host name** box, type the host name or IP address of your BIG-IP system.
4. In the **User name** and **Password** boxes, type the appropriate administrator log on information.
5. Click **Login**.
   The WinSCP client opens.
6. In the left pane, navigate to the location where you saved the script in step 1.
7. In the right pane, navigate to **/usr/bin/monitors**.
8. In the left pane, select the script and drag it to the right pane.
   You can now safely close WinSCP.

   The next task is to create the monitor that calls the script in the BIG-IP Configuration utility.

To create the monitor that calls the script

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
   In our example, we type **Extend_ping**.
4. From the **Type** list, select **External**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a (1:3) +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.
6. In the **External Program** box, type **`/usr/bin/monitors/extend_ping`**
7. Click the **Finished** button.

Creating the pools

The next step is to create a pool on the BIG-IP LTM system for the Proxy services.

To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button. The New Pool screen opens.

3. In the **Name** box, type a name for your pool. In our example, we type **Extend_pool**.
4. In the **Health Monitors** section, select the appropriate health monitor you created in *Configuring the health monitor*, on page 3 and then click the Add (**<<**) button:
5. From the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).
   In our example, we select **Least Connections (node)**.
6. In the New Members section, make sure the **New Address** option button is selected.
7. In the **Address** box, type the IP address of your first *Extend Proxy service. In our example, we type **10.10.10.101**.
8. In the **Service Port** box, type in the TCP port on which the Proxy service is running. In our example, we type **9099**.
9. Click the **Add** button to add the member to the list.
10. Repeat steps 7-9 for each Proxy service you have running. In our example, we have three Proxy servers, 10.10.10.101, .102, and .103.
11. Leave the rest of the settings at the default levels.
12. Click the **Finished** button.

Creating the profiles

The next step is to create the profiles. Although you may use the default profiles, we strongly recommend you create new profiles based on the default parent profiles. By creating new profiles, you may easily modify the profile settings specific to your deployment without altering default global behaviors.

For the Oracle Coherence*Extend Proxy configuration, we create three new profiles: two TCP profiles, and an SSL profile (optional).

Creating TCP profiles

In this section, we create the TCP profiles. We recommend creating **tcp-lan-optimized** and **tcp-wan-optimized** profiles.

Creating the WAN optimized TCP profile

The TCP WAN Profile is used to configure the TCP parameters for the Coherence clients, and can be tuned to your particular network. In our example, we use the existing TCP WAN parent profile.

To create a new WAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **Protocol** menu, select **TCP**.

2. Click the **Create** button. The New TCP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **tcpwan_extend**.
4. From the **Parent Profile** list, select **tcp-wan-optimized**.
5. Modify any of the other settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the LAN optimized TCP profile

The TCP LAN Profile is used to configure the TCP parameters for the Coherence*Extend Proxy servers, and can be tuned to your particular network. In our example, we use the existing TCP LAN parent profile.

To create a new LAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens by default.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **tcplan_extend**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. Modify any of the other settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the SSL Profile

Next, we configure the optional SSL profile so the BIG-IP LTM can offload SSL from the Coherence*Extend deployment.

◆ Note

This step is optional, but is a best practice to secure the connection between the Coherence clients and the BIG-IP.

SSL offload can be enabled on the BIG-IP LTM in 4 steps:

1. Import a server SSL certificate and key.
2. Create a client SSL profile.
3. Add the SSL profile to the Coherence*Extend virtual server.

4. Enable SSL in the Coherence*Extend client cache configuration file.

Importing the SSL certificate and key

To import a server SSL certificate and key, use the following procedure.

To import the certificate and key

1. On the Main tab, expand **Local Traffic**, then **SSL Certificates**.
2. Click the **Import** button. The SSL Certificate screen opens.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).
5. In the **Certificate** (or **Key**) **Name** box, type a unique name for the certificate or key.
6. In the **Certificate** (or **Key**) **Source** box, choose to either upload the file or paste the text.
7. Click **Import**.
8. If you imported the certificate, repeat this procedure for the key.

Creating the SSL profile

The next task is to create the SSL profile.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **SSL** menu, click **Client**.
2. Click the **Create** button. The New Client SSL Profile screen opens.
3. In the **Name** box, type a name. We type **Extend_client_ssl**.
4. In the Configuration section, click a check in the **Certificate** and **Key** Custom boxes.
5. From the **Certificate** list, select the name of the Certificate you imported.
6. From the **Key** list, select the key you imported.
7. Click the **Finished** button.

Creating the virtual server

The Virtual Server is the IP address to which clients will connect, and contains all of the configuration profiles and the Pool that it will use for client traffic.

A virtual server is a the object on the BIG-IP system that is represented by an IP address and port. Clients on an external network can send application traffic to a virtual server, which then directs the traffic according to your configuration instructions. The virtual server will then load balance traffic to members of the pool.

Virtual servers increase the availability of resources for processing client requests. In the case of Coherence*Extend, you will configure a virtual server that will direct traffic to the pool of proxy services you configured earlier.

To create the virtual servers

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type an appropriate name for this virtual server. We recommend using a name that includes the service, such as **Extend_vs**.
4. In the **Address** box, type the IP address to which the Coherence*Extend clients will connect. In our example, we use **10.196.21.3.**
5. In the **Service Port** box, type the appropriate Service Port. In our example, we type **9099**.

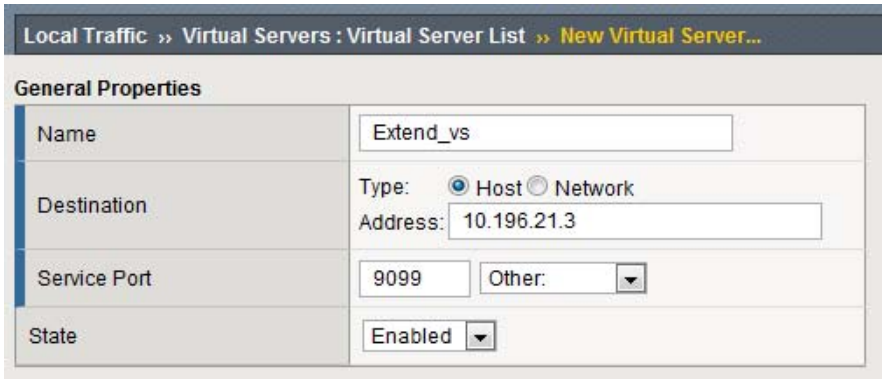

Figure 3 *General properties of the virtual server*

6. From the Configuration list, select **Advanced**.
   The Advanced configuration options appear.
7. Leave the **Type** list at the default setting: **Standard**.
8. From the **Protocol Profile (Client)** list, select the profile you created in *Creating the WAN optimized TCP profile*, on page 7. In our example, we select **tcpwan_extend**.

9. From the **Protocol Profile (Server)** list, select the profile you created in *Creating the LAN optimized TCP profile*, on page 8. In our example, we select **tcplan_extend**.
10. *Optional*: If you are offloading SSL from the Coherence*Extend clients, from the **SSL Profile (Client)** list, select the profile you created in *Creating the SSL profile*, on page 9. In our example, we select **Extend_ssl.**

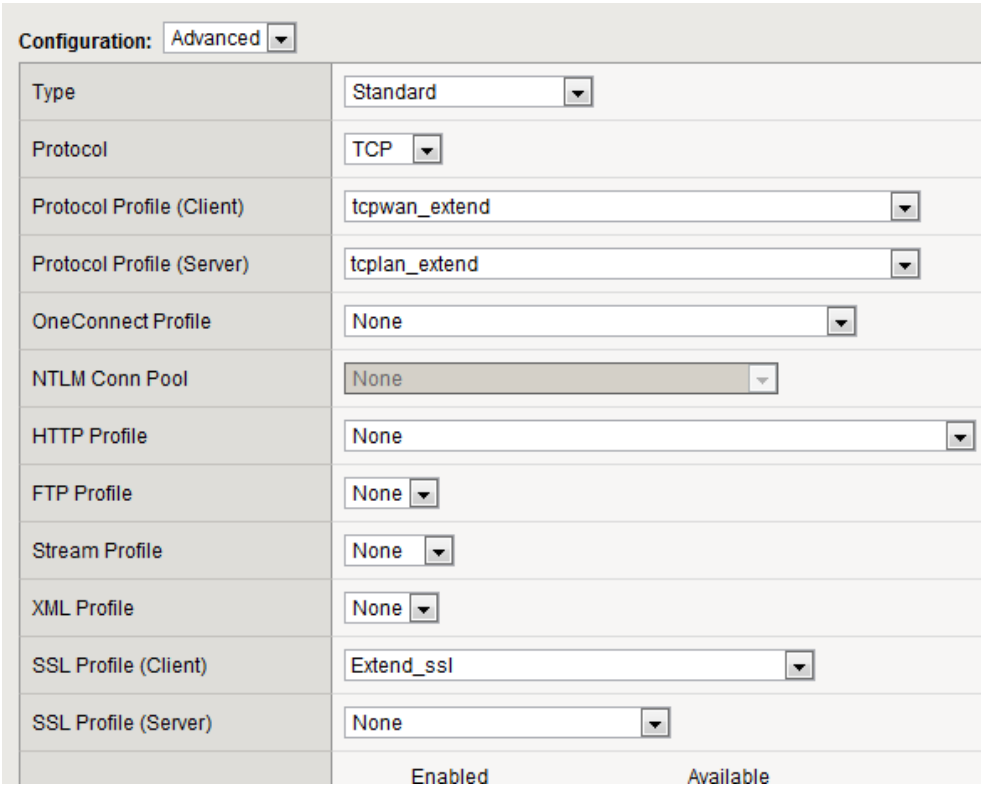

Figure 4 *Configuration section of the virtual server (truncated)*

11. From the **SNAT Pool** list, select **Automap**.
12. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the pools*, on page 6. In our example, we select **Extend_pool.**
13. Click the **Finished** button.

This completes the BIG-IP LTM configuration. Continue with the following section for modifications to the Oracle configuration.

Modifying the Oracle Coherence configuration

Now that the BIG-IP device has been configured, you must modify your Coherence*Extend cache configuration files.

To modify the server cache configuration files

1. Open the server cache configuration file for editing.
2. In the **proxy-scheme** element, locate **<load-balancer>**, and then type **client** as shown in the following example:

```
<proxy-scheme>
<service-name>ExtendTcpProxyService</service-name>
<acceptor-config>
   <tcp-acceptor>
      <local-address>
         <address>192.168.1.2</address>
         <port>9099</port>
      </local-address>
    </tcp-acceptor>
  </acceptor-config>
  <load-balancer>client</load-balancer>
  <autostart>true</autostart>
</proxy-scheme>
```

3. Repeat step 2 for all server cache configuration files.

Next, you must modify the client cache configuration file.

To modify the client cache configuration file

4. Open the client cache configuration file for editing.
5. In the **remote-cache-scheme** element, locate **<address>**, and then type the IP address of the virtual server you entered in step 4 of *Creating the virtual server*, on page 10. In our example, we type **10.196.21.3**.
6. Locate **<Port>**, and then type the port of the virtual server you entered in step 5 of *Creating the virtual server*, on page 10. In our example, we type **9099**.
7. Under <outgoing-message-handler>, locate **<heartbeat-interval>**. Type a heart-beat interval. In our example, we type **5s**. This causes the client to periodically send a heartbeat message over its TCP/IP connection at the configured interval. This is essential to prevent the BIG-IP device from disconnecting idle clients.

```
<remote-cache-scheme>
 <scheme-name>extend-direct</scheme-name>
 <service-name>ExtendTcpCacheService</service-name>
 <initiator-config>
    <tcp-initiator>
      <remote-addresses>
        <socket-address>
```

```
            <address>10.196.21.3</address>
            <port>9099</port>
          </socket-address>
        </remote-addresses>
      </tcp-initiator>
      <outgoing-message-handler>
        <heartbeat-interval>5s</heartbeat-interval>
      </outgoing-message-handler>
    </initiator-config>
  </remote-cache-scheme>
```

This completes the configuration.

References

F5 links:

www.f5.com

www.f5.com/oracle

http://www.f5.com/solutions/applications/oracle/middleware/

http://www.f5.com/products/big-ip/

Oracle links:

www.oracle.com

http://www.oracle.com/us/products/middleware/coherence/index.html

Oracle® Coherence Client Guide - Release 3.7; Part Number E18678-01

http://download.oracle.com/docs/cd/E18686_01/coh.37/e18678/toc.htm