ORACLE ENTERPRISE MANAGER 12c

An Oracle White Paper
March, 2012

Enterprise Manager 12c Cloud Control: Configuring OMS High Availability with F5 BIG-IP Local Traffic Manager

ORACLE

Executive Overview

Oracle Enterprise Manager is Oracle's integrated enterprise IT management product line and provides the industry's first complete cloud lifecycle management solution. Oracle Enterprise Manager's Business-Driven IT Management capabilities allow you to quickly set up, manage and support enterprise clouds and traditional Oracle IT environments from applications to disk. Enterprise Manager allows customers to achieve:

- *Best service levels for traditional and cloud applications* through management from a business perspective including Oracle Fusion Applications

- *Maximum return on IT management* investment through the best solutions for intelligent management of the Oracle stack and engineered systems

- *Unmatched customer support experience* through real-time integration of Oracle's knowledgebase with each customer environment

Oracle Maximum Availability Architecture (MAA) is the Oracle best practices blueprint for implementing Oracle high-availability technologies. Oracle Enterprise Manager is the management platform for Oracle solutions. This white paper has been jointly written by Oracle Corporation and F5 Networks and provides the detailed steps for implementation of an Oracle MAA solution for Oracle Enterprise Manager Cloud Control using BIG-IP from F5 Networks as the front end for the Cloud Control mid-tiers, known as the Oracle Management Service (OMS). The BIG-IP hardware platform can provide load balancing, high availability, service monitoring, TCP/IP enhancements, and application persistence for the Cloud Control environment as the front end for several Cloud Control services. Most of the procedures in this document are performed on the BIG-IP Local Traffic Manager (LTM), targeting different areas of the infrastructure where high availability is required to provide continuous access to the Cloud Control OMS application that has been deemed mission critical. This paper is designed to provide the Cloud Control Administrator with an introduction to the high availability and load balancing features available with F5 solutions. Step-by-step configuration instructions and screen shots are provided to make it easier to understand and implement BIG-IP as a critical component of the Cloud Control architecture. In general, assume that the following software versions are used in this white paper:

- BIG-IP Version 11.1.0, Build 1943.0 Final

- Cloud Control Release 12.1.0.1.0

Any distinction in release numbers is noted within the relevant discussions of this paper.

Note: This white paper assumes that you are familiar with BIG-IP from F5 Networks. See Appendix A for a quick terminology reference. For detailed information, see the BIG-IP Solutions Guide and BIG-IP Configuration Guide, and Chapter 21 in the Oracle Enterprise Manager Cloud Control Administrator's Guide.

About F5 BIG-IP and Oracle Enterprise Manager Cloud Control

The diagram below shows F5 and Oracle Enterprise Manager components in a Cloud Control Environment

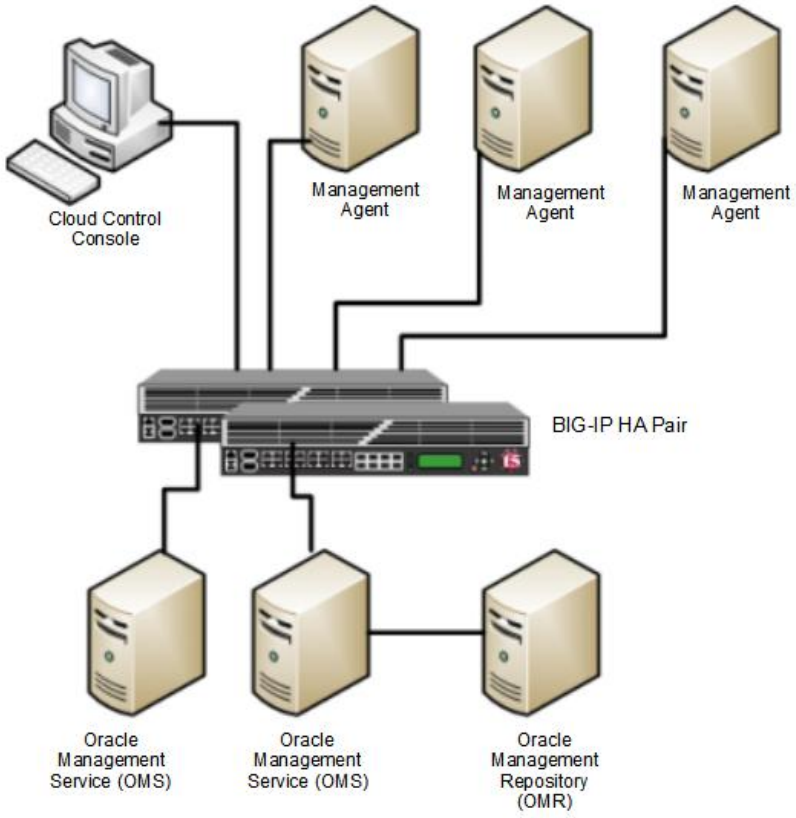

Cloud Control OMS Servers provide http or https access to a set of Cloud Control services, listed below, to the Cloud Control clients, including the Cloud Control console and Management Agents. When more than one Cloud Control OMS Server is deployed, F5 BIG IP can load balance requests for each service via virtual servers, with the Cloud Control clients making service requests using a virtual hostname.

The Cloud Control services that can be served by the F5 BIG IP in a multi-OMS setup are:

CLOUD CONTROL SERVICE	DESCRIPTION
Secure Console	HTTPS access to Cloud Control Console
Unsecure Console	HTTP access to Cloud Control Console
Secure Upload	Secure Agent to OMS communication
Agent Registration	Unsecure Agent to OMS communication

Each Cloud Control service that is managed by F5 BIG-IP requires that you configure the following F5 BIG-IP Local Traffic Manager objects:

- A health monitor for the service.

  The health monitor is the process by which BIG-IP determines that the service is up and running and can take connections.

- A TCP profile for the service.

  The TCP profile is used to tune the TCP/IP stack from BIG-IP for optimum performance.

- A Pool for the service.

  A Pool is a group of two or more OMS Cloud Control servers that are load balanced, with each pool running an instance of the different Cloud Control services.

- A Persistence profile for the service.

  The Persistence profile is used to link a client to the proper Cloud Control pool member for the duration of a connection. This is required for all Cloud Control services except Secure Upload.

- A Virtual Server for the service.

  A Virtual Server is a unique IP address and port that represents a pool of servers.

The remainder of this paper provides detailed instructions for configuring BIG-IP LTM for Cloud Control services. Each of the configuration discussions imparts:

- Operational best practices when using the F5 BIG-IP Web configuration utility to configure Oracle Enterprise Manager Cloud Control services.
- Screen shots of the BIG-IP Web interface that are based on BIG-IP Version 11.1.0 software.
- A Configuration Summary page naming all of the Cloud Control services and matching F5 configuration elements.

For additional information about configuring BIG-IP, see the BIG-IP documentation at http://www.f5.com.

Configuring an F5 BIG-IP LTM for Cloud Control Services

Use the instructions that follow to configure Oracle Enterprise Manager Cloud Control to work with the F5 BIG-IP LTM.

Prerequisites and Best Practice Recommendations

Use the following general guidelines when building your configuration.

Use BIG-IP Administrative Partitions

A feature of the BIG-IP software is the ability to use *Administrative Partitions* to allow multiple administrators or operators to manage the configuration. The best practice recommendation is to create a dedicated Administrative Partition on the BIG-IP for configuration access and use by the Cloud Control administrators. All the necessary F5 configuration elements for the MAA Cloud Control environment are located in the Administrative Partition. Additions, deletions, and changes to these pools created in this partition would not interfere with any other services provided by the BIG-IP.

For more information about configuring Administrative Partitions, see the BIG-IP documentation.

Use the Configuration Table and Standard Naming Conventions

To make the configuration consistent, easy to read, and easy to administer, this white paper uses a standard naming convention for the F5 configuration. Your organization may already use naming standards (which your Network Operations team can provide if necessary), or you can create naming conventions or adopt the ones used in this white paper.

The following table shows the naming conventions used by the MAA example described in this white paper.

BIG-IP CONFIGURATION OBJECT	CONVENTION
Health Monitors	mon_<service_label><port>
TCP Profiles	tcp_<service_label><port>
Pools	pool_<service_label><port>
Cookie Persistence Profile	cookie_<service_label><port>
Source IP Address Persistence Profile	sourceip_<service_label><port>
Virtual Server	vs_<service_label><port>

Using the Secure Console service as an example, we derived the service label "ccsc" and terminated each name with the TCP port number as a suffix.

In the following list of names, the TCP port number 7799 is the secure console port for cloud control, and 443 is used for the virtual server:

- `mon_ccsc7799`
- `tcp_ccsc7799`
- `sourceip_ccsc7799`
- `pool_ccsc7799`
- `vs_ccsc443`

Note: A Virtual Server port can be different to a pool port, as is the case in this example. Configuring port 443 for the secure console Virtual Server port allows https console access without specifying a port even though the Cloud Control secure console is running on port 7799. When configured using the following instructions, BIG-IP takes care of forwarding the request from the Virtual Server to the correct port on the OMS servers transparently.

The example values are reproduced in the table below, which provides a reference for all of the F5 configuration objects in this document. All of the names used in this white paper follow this convention, which is considered to be a best practice.

CLOUD CONTROL SERVICE	TCP PORT	MONITOR NAME	TCP PROFILE NAME	PERSISTENCE PROFILE	POOL NAME	VIRTUAL SERVER NAME	VIRTUAL SERVER PORT
Secure Console	7799	mon_ccsc7799	tcp_ccsc7799	sourceip_ccsc7799	pool_ccsc7799	vs_ccsc443	443
Unsecure Console	7788	mon_ccuc7788	tcp_ccuc7788	sourceip_ccuc7788	pool_ccuc7788	vs_ccuc7788	7788
Secure Upload	4900	mon_ccsu4900	tcp_ccsu4900	None	pool_ccsu4900	vs_ccsu4900	4900
Agent Registration	4889	mon_ccar4889	tcp_aaar4889	cookie_ccar4889	pool_ccar4889	vs_ccar4889	4889

Methodology

To configure BIG-IP for Cloud Control, it is necessary to create health monitor, load balancing pool, persistence profile and virtual server configuration objects for the Cloud Control services listed in the above table. The following sections describe how to create and configure each of the configuration objects, providing a reference table with the required settings for each of the Cloud Control services, followed by a detailed example, including screenshots, using the Secure Console service as an example. The steps in each section should be repeated for each of the Cloud Control services.

Create the Health Monitors

Create the following Monitors.

CLOUD CONTROL SERVICE	TCP PORT	MONITOR NAME	TYPE	INTERVAL	TIMEOUT	SEND STRING	RECEIVE STRING
Secure Console ***(when not using SSO)***	7799	mon_ccsc7799	https	5	16	GET /em/console/home HTTP/1.1\r\nHost: \r\nConnection: Close \r\n\r\n	/em/login.jsp
Secure Console ***(when using SSO)***	7799	mon_ccsc7799	https	5	16	GET /empbs/genwallet \r\n	GenWallet Servlet activated
Unsecure Console ***(when not using SSO)***	7788	mon_ccuc7788	http	5	16	GET /em/console/home HTTP/1.1\r\nHost: \r\nConnection: Close \r\n\r\n	/em/login.jsp
Unsecure Console ***(when using SSO)***	7788	mon_ccuc7788	http	5	16	GET /empbs/genwallet \r\n	GenWallet Servlet activated
Secure Upload	4900	mon_ccsu4900	https	60	181	GET /empbs/upload \r\n	Http Receiver Servlet active!
Agent Registration	4889	mon_ccar4889	http	60	181	GET /empbs/genwallet \r\n	GenWallet Servlet activated

Note: There are two entries for both the Secure Console and the Unsecure Console services. Configuration of the Secure Console and Unsecure console monitors differs depending on whether SSO has been configured for Enterprise Manager authentication. Only one monitor needs to be configured for each service, choose the relevant one for your environment.

The following steps should be followed for each Monitor that needs to be created:

1. On the **Main** tab, expand **Local Traffic**, and then click **Monitor**s.
2. On the Monitors screen, click **Create**.

   The New Monitor screen opens.
3. In the **Name** field, enter a unique name for the Monitor. For example: mon_ccsc7799
4. From the **Type** list, select the type for the Monitor. For example: HTTPS.

   The Monitor configuration options display.
5. From the **Configuration** list, select Advanced.
6. In the Configuration section, enter the appropriate values in Interval and Timeout fields:
   - **Interval** is the Health Monitor property that specifies the frequency at which the system issues the monitor check.
   - **Timeout** is the setting that allows the monitor to fail three times before marking a pool member as down. The recommendation is to set the BIG-IP LTM Health Monitor Timeout setting as (3 * "Interval") + 1, allowing at least a 1:3 +1 ratio between the interval and the timeout.

For example, set Interval to 5 and set Timeout to 16. Refer to the table above for the appropriate Interval and Timeout values for your monitor.

7. In the **Send String** field, add the Send String for the Monitor you are creating. For example:

   GET /em/console/home HTTP/1.1\r\nHost: \r\nConnection: Close \r\n\r\n

8. In the **Receive String** field, add the Receive String for the Monitor you are creating. For example:

   /em/login.jsp

9. In the **Alias Service Port** field, enter the port for the monitor you are creating, for example 7799.

   All other configuration settings are optional.

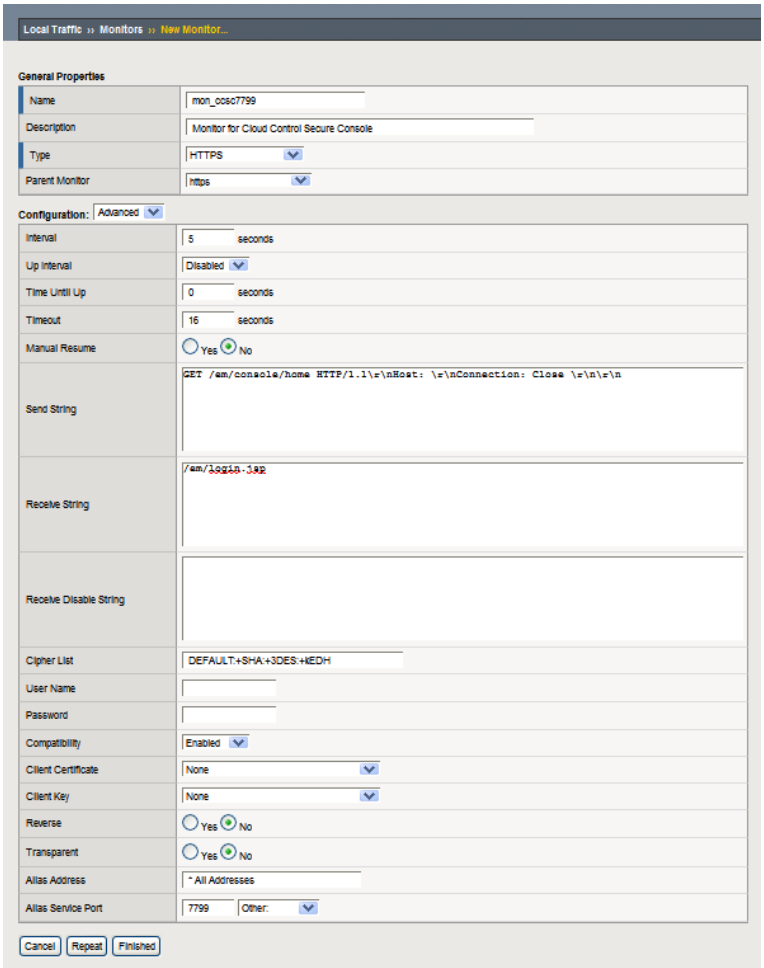

10. Click Finished.

Create the Cloud Control Pools

A BIG-IP LTM pool is a set of servers grouped together to receive traffic according to a load balancing method. A pool needs to be created for each of the Cloud Control services as per the following table:

CLOUD CONTROL SERVICE	POOL NAME	ASSOCIATED HEALTH MONITOR	Load Balancing	MEMBERS
Secure Console	pool_ccsc7799	mon_ccsc7799	Least Connections (member)	OMS Host A:7799 OMS Host B:7799
Unsecure Console	pool_ccuc7788	mon_ccuc7799	Least Connections (member)	OMS Host A:7799 OMS Host B:7799
Secure Upload	pool_ccsu4900	mon_ccsu4900	Least Connections (member)	OMS Host A:4900 OMS Host B:4900
Agent Registration	pool_ccar4889	mon_ccar4889	Least Connections (member)	OMS Host A:4889 OMS Host B:4889

The following steps should be followed for each Pool that needs to be created:

1. On the **Main** tab, expand **Local Traffic**, and then click **Pools**.

   The Pool screen opens.

2. In the upper right portion of the screen, click **Create**.

   The New Pool screen opens.

   Note: For more (optional) pool configuration settings, from the Configuration list, select Advanced. Configure these settings, as applicable, for your network.

3. In the **Name** field, enter a unique name for your pool.

   For example, enter **pool_ccsc7799**.

4. In the **Health Monitors** section, select the name of the monitor for the service that the pool is being created for, and click the Add (**<<**) button.

   In this example, we selected **mon_ccsc7799**.

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 this example, we selected **Least Connections (member)**.

6. Keep the **Priority Group** Activation value as **Disabled**.

7. In the **New Members** section, add each OMS host as a member, one at a time, by entering the OMS hostname in the **Node Name** field, the OMS IP address in the **Address** field and the port for the service that the pool is being created for in the **Service Port** field, then clicking **Add**.

8. Click **Finished**.

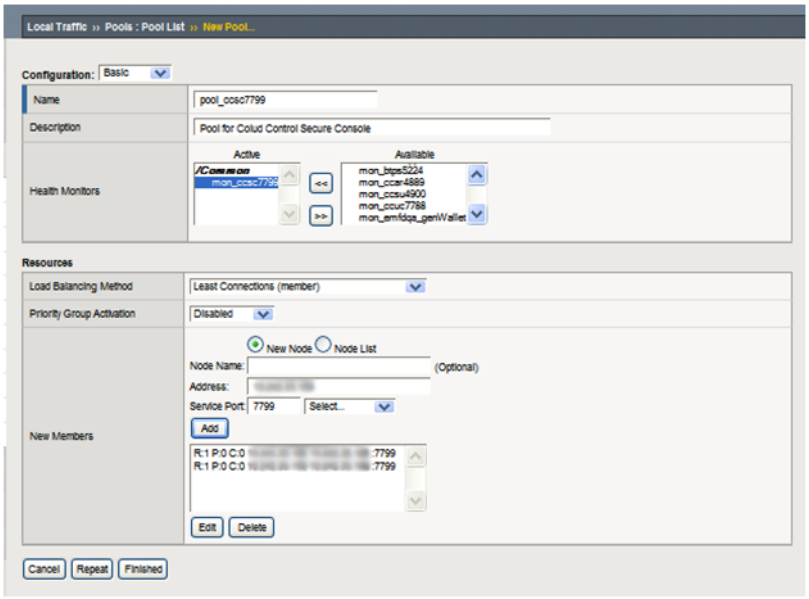

Create the TCP Profiles

In our example, we base each TCP profile on the default TCP profile, and keep all the options at their default settings. You can configure these options, as appropriate, for your network. A TCP profile needs to be created for each of the Cloud Control services as per the following table:

CLOUD CONTROL SERVICE	TCP PROFILE NAME
Secure Console	tcp_ccsc7799
Unsecure Console	tcp_ccuc7788
Secure Upload	tcp_ccsu4900
Agent Registration	tcp_ccar4889

The steps below should be followed for each TCP profile that needs to be created:

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.

   The HTTP Profiles screen opens.
3. On the Menu bar, from the **Protocol** menu, select **TCP**.
4. In the upper right portion of the screen, click **Create**.

   The New TCP Profile screen opens.
5. In the **Name** field, enter a unique name for this profile. For example: **tcp_ccsc7799**.

6. If needed, modify as applicable for your network. See the F5 online help for more information about the configuration options. Note that this example keeps the settings at their default levels.
7. Click **Finished**.

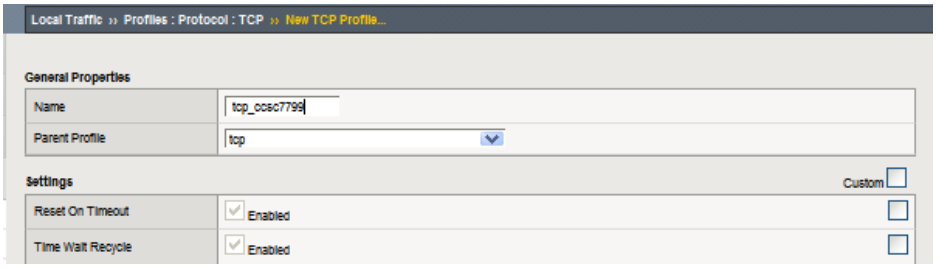

Create the Persistence Profiles

A persistence profile needs to be created for each of the Cloud Control services, except for the secure upload service, as per the following table:

CLOUD CONTROL SERVICE	F5 PERSISTENCE PROFILE NAME	TYPE	TIMEOUT	EXPIRATION
Secure Console	sourceip_ccsc7799	Source Address Affinity	3600	Not Applicable
Unsecure Console	sourceip_ccuc7788	Source Address Affinity	3600	Not Applicable
Agent Registration	cookie_ccar4889	Cookie	Not Applicable	3600

The following steps should be followed for each persistence profile that needs to be created**:**

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.

   The HTTP Profiles screen opens.
2. On the Menu bar, click **Persistence**.

   The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click **Create**.
4. The New Persistence Profile screen opens.
5. In the Name field, enter a unique name for this profile. For example, enter **sourceip_ccsc7799**.
6. If the persistence type for the service being created is 'Source Address Affinity', from the **Persistence Type list** select **Source Address Affinity**.
   - The configuration options for SourceIP persistence display.
   - Check the box next to the 'Timeout' field to allow the Timeout value to be overridden
   - Modify the **Timeout** value to **3600**.

If the persistence type for the service being created is 'Cookie', from the **Persistence Type list** select **Cookie**.

- The configuration options for Cookie persistence display.
- Check the box next to the 'Expiration' field to allow the Expiration value to be overridden
- Clear the clear the **Session Cookie** box.
- The expiration options appear.
- Provide the value **3600** in the **seconds** field.

7. Click **Finished**.

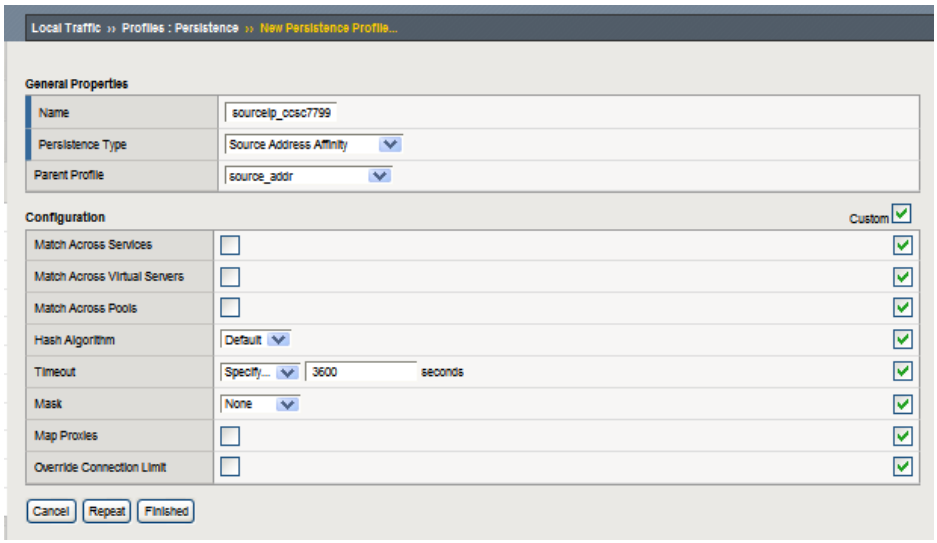

Note: For more information about creating or modifying profiles, or applying profiles in general, see the BIG-IP documentation.

Create a Redirect iRule for the Unsecure Console service

A Redirect iRule can be created to take incoming HTTP requests (non-secure) and redirect those requests to the correct HTTPS (secure) virtual server without user interaction. This Redirect iRule is used in the configuration of the Cloud Control unsecure console service virtual server, to redirect clients to the matching Cloud Control secure console service.

The following steps should be followed to create the Redirect iRule:

1. On the Main tab, expand **Local Traffic** and click **iRules**.
2. In the upper right portion of the iRule screen, click **Create**.
3. In the **Name** field on the New iRule screen, enter a name for your iRule. In our example, we use **ccuc_httptohttps**.
4. In the **Definition** section, copy and paste the following iRule:

```
when HTTP_REQUEST {
    HTTP::redirect "https://[getfield [HTTP::host] : 1][HTTP::uri]"
}
```

5. Click **Finished**.

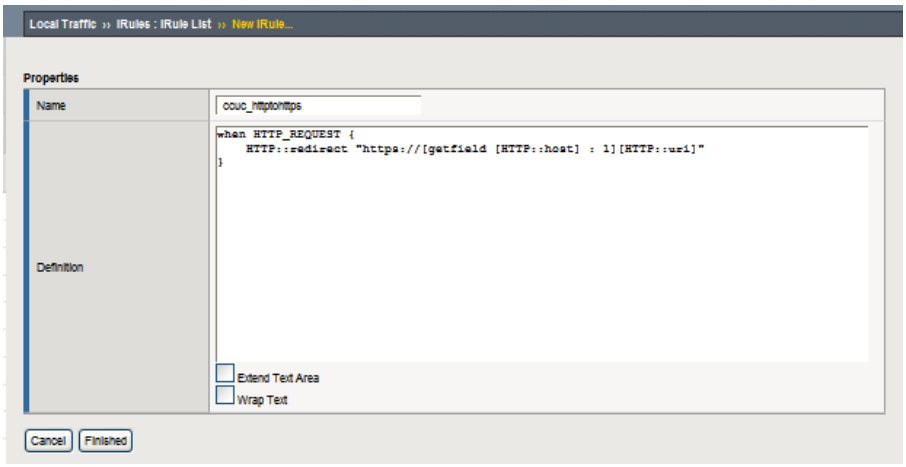

Create the Virtual Servers

The final step is to define virtual servers that references the profiles and pools created for each Cloud Control Service. A virtual server with its virtual address and port number, is the client addressable host name or IP address through which members of a load balancing pool are made available to a client. A virtual server needs to be created for each of the Cloud Control services as per the following table:

CLOUD CONTROL SERVICE	VIRTUAL SERVER NAME	VIRTUAL IP AND PORT	PROTOCOL PROFILE (CLIENT)	HTTP PROFILE	SNAT POOL	iRULE	DEFAULT POOL	DEFAULT PERSISTENCE PROFILE
Secure Console	vs_ccsc443	<virtual Host IP>:443	tcp_ccsc7799	None	Auto	None	pool_ccsc 7799	sourceip_ ccsc7799
Unsecure Console	vs_ccuc7788	<virtual Host IP>:7788	tcp_ccuc7788	http	Auto	ccuc_ httpto https	pool_ccuc 7788	sourceip_ ccuc7788
Secure Upload	vs_ccsu4900	<virtual Host IP>:4900	tcp_ccsu4900	None	Auto	None	pool_ccsu 4900	None
Agent Registration	vs_ccar4889	<virtual Host IP>:4889	tcp_ccar4889	http	Auto	None	pool_ccar 4889	cookie_ ccar4889

The following steps should be followed for each virtual server that needs to be created**:**

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.

   The Virtual Servers screen opens.

2. In the upper right portion of the screen, click the **Create** button.

The New Virtual Server screen opens.

3. In the **Name** field, enter a unique name for this virtual server.

   In this example, we entered **vs_ccsc443**.

4. In the **Destination** section, select the **Host** option button.
5. In the **Address** field, enter the IP address of this virtual server.
6. In the **Service Port** field, enter the Virtual IP Port for the service being created. For example, **443**.
7. From the Configuration list, select **Advanced**.

   The Advanced configuration options display.

8. Keep the **Type** list at the default setting: **Standard**.
9. From the **Protocol Profile (Client)** list select the name of the profile for the service being created.

   In this example, we selected **tcp_ccsc7799**.

10. Keep the **Protocol Profile (Server)** option at the default setting.
11. For the Agent Registration and Unsecure Console services only, select **http** from the **HTTP Profile** list.
12. Change the **SNAT Pool** setting to **Automap**.
13. For the Unsecure Console service only, in the **iRules** section, add the iRule created earlier by selecting it in the **Available** list and clicking **<<** to add it to the **Enabled** list.
14. In the Resources section, from the **Default Pool** list, select the pool created for the service that the virtual server is being created for.

    In this example, we selected **pool_ccsc7799**.

15. From the **Default Persistence Profile** list, select the persistence profile created for the service that the virtual server is being created for.

    In this example, we selected **sourceip_ccsc7799**.

16. Click **Finished**.

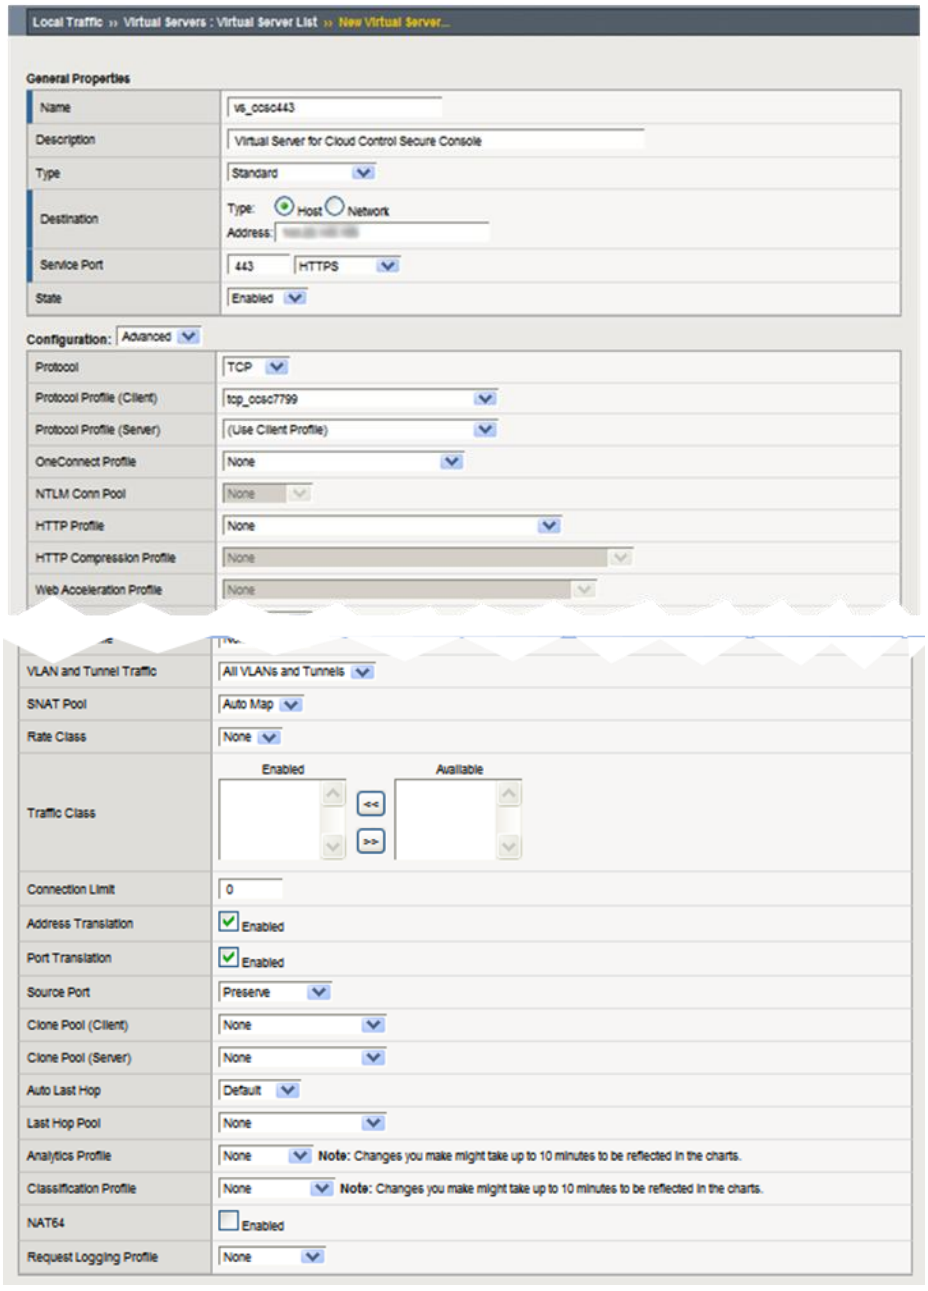
Local Traffic ›› Virtual Servers : Virtual Server List ›› New Virtual Server...
General Properties
Name
vs_ocsc443
Description
Virtual Server for Cloud Control Secure Console
Type
Standard
Destination
Type: Host Network
Address:
Service Port
443
HTTPS
State
Enabled
Configuration: Advanced
Protocol
TCP
Protocol Profile (Client)
tcp_ocsc7799
Protocol Profile (Server)
(Use Client Profile)
OneConnect Profile
None
NTLM Conn Pool
None
HTTP Profile
None
HTTP Compression Profile
None
Web Acceleration Profile
None
VLAN and Tunnel Traffic
All VLANs and Tunnels
SNAT Pool
Auto Map
Rate Class
None
Traffic Class
Enabled
Available
<<
>>
Connection Limit
0
Address Translation
Enabled
Port Translation
Enabled
Source Port
Preserve
Clone Pool (Client)
None
Clone Pool (Server)
None
Auto Last Hop
Default
Last Hop Pool
None
Analytics Profile
None
Note: Changes you make might take up to 10 minutes to be reflected in the charts.
Classification Profile
None
Note: Changes you make might take up to 10 minutes to be reflected in the charts.
NAT64
Enabled
Request Logging Profile
None

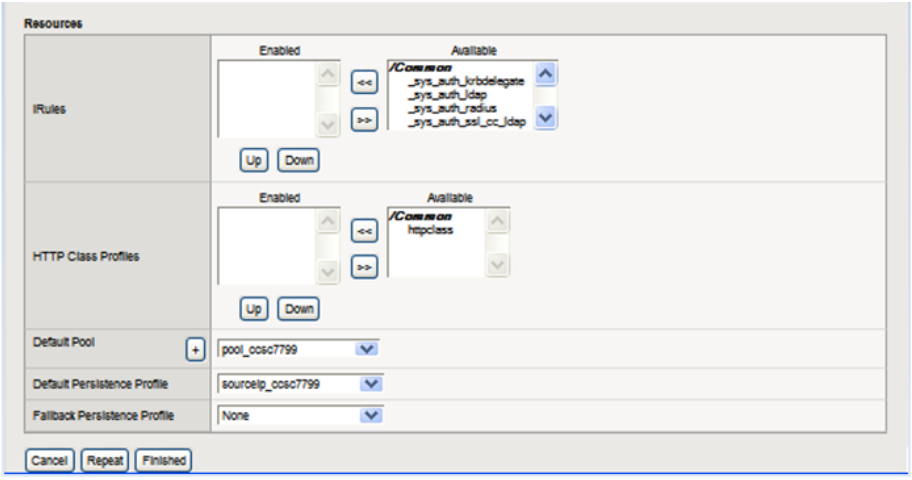

Configuring Enterprise Manager for Use with F5 BIG-IP LTM

Resecure Management Service

The management services must now be reconfigured so that the Management Service certificate uses the hostname associated with the Load Balancer. Steps 1 and 2 must be repeated for each configured OMS

1: Resecure OMS

In our example we issued the following command:

```
$ emctl secure oms -sysman_pwd xxxxxx -reg_pwd xxxxxx -host
slb.example.com -secure_port 4900 -slb_port 4900 -slb_console_port
443 -console -lock -lock_console
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Securing OMS... Started.
Securing OMS... Successful
Restart OMS
```

2: Restart the OMS

```
$ ./emctl stop oms -all
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Stopping WebTier...
WebTier Successfully Stopped
Stopping Oracle Management Server...
Oracle Management Server Successfully Stopped
AdminServer Successfully Stopped
Oracle Management Server is Down
$
$ ./emctl start oms
```

```
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Starting WebTier...
WebTier Successfully Started
Starting Oracle Management Server...
Oracle Management Server Successfully Started
Oracle Management Server is Up
```

3: Resecure all Management Agents

```
$ ./emctl secure agent -emdWalletSrcUrl
https://slb.example.com:4900/em
Oracle Enterprise Manager 12c Release 1 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Agent successfully stopped...   Done.
Securing agent...   Started.
Enter Agent Registration Password :
Agent successfully restarted...   Done.
EMD gensudoprops completed successfully
Securing agent...   Successful.
```

Verify Status of Management Service

The OMS configuration can be checked using the emctl status oms –details command. Following successful configuration this should show that the SLB or virtual hostname field has been set.

```
$ ./emctl status oms -details
Oracle Enterprise Manager Cloud Control 12c Release 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation.  All rights reserved.
Enter Enterprise Manager Root (SYSMAN) Password :
Console Server Host : omsa.example.com
HTTP Console Port   : 7788
HTTPS Console Port  : 7799
HTTP Upload Port    : 4889
HTTPS Upload Port   : 4900
SLB or virtual hostname: slb.example.com
HTTPS SLB Upload Port : 4900
HTTPS SLB Console Port : 443
Agent Upload is locked.
OMS Console is unlocked.
Active CA ID: 1
Console URL: https://slb.example.com:443/em
Upload URL: https://slb.example.com:4900/empbs/upload

WLS Domain Information
Domain Name      : GCDomain
Admin Server Host: omsa.xxx.xxx.xxx

Managed Server Information
Managed Server Instance Name: EMGC_OMS1
Managed Server Instance Host: omsa.xxx.xxx.xxx
```

Appendix A: F5 BIG-IP Local Traffic Manager Terms

This document assumes that you are familiar with F5 Networks BIG-IP. This section discusses the basic terminology. For a detailed discussion of these terms, see the BIG-IP Solutions Guide and the BIG-IP Configuration Guide.

Monitor

Monitors are used to verify the operational state of pool members. Monitors verify connections and services on nodes that are members of load-balancing pools. A monitor is designed to check the status of a service on an ongoing basis, at a set interval. If the service being checked does not respond within a specified timeout period, or the status of the service indicates that the performance has degraded, the BIG-IP system automatically takes it out of the pool and will choose the other members of the pool. When the node or service becomes available again, the monitor detects this and the member is automatically accessible to the pool and able to handle traffic. Monitors can be as simple as an ICMP ping to a server's IP address, to a TCP 3-way handshake to a service port, or as sophisticated as an HTTP Get Request with parameters, or SSL session negotiation. F5 monitors can also be custom programmed for specific needs.

Pool

A *pool* is a set of servers grouped together to receive traffic on a specific TCP port using a load balancing method. Each pool can have its own unique characteristic for a persistence definition and the load-balancing algorithm used. The preferred setting of the load balance algorithm for all Cloud Control pools is Least Connections (Member). 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, and ultimately to one of the pool members.

Member

A member of the pool is defined as a node, as a destination for traffic, with an IP address and a port definition, expressed as a.b.c.d:nn, or 192.168.1.200:80 for a Web server with IP address 192.168.1.200 and listening on port 80. There must be at least two members in every pool to provide high availability. If one of the pool members is unavailable or offline, traffic is sent to the remaining member or members.

Virtual Server

A *virtual server* with its virtual IP Address and port number is the client addressable hostname or IP address through which members of a load balancing pool are made available to a client. After a virtual server receives a request, it directs the request to a member of the pool based on a chosen load balancing method. After a virtual server receives traffic, either directly or through a rule, the virtual server can optionally perform a number of different operations, such as inserting or modifying a header into an HTTP request, setting a persistence record, or redirecting the request to another site or fallback destination. Before creating a virtual server, you must configure a load balancing pool of the actual physical devices (members) 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 criterion, then you can create a rule specifying the criteria, and BIG-IP would forward the traffic to a pool matching the rule's criteria. A virtual server is configured to a specific port or to accept "ANY" ports. A given F5 BIG-IP device may contain one or more virtual servers.

Profile

A profile is an F5 object that contains user-configurable settings for controlling the behavior of a particular type of network traffic, such as TCP or HTTP connections. BIG-IP version 9.0 and later uses profiles. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient. It also allows for different characteristics to be matched to specific clients or applications. For example, one HTTP profile could be configured for Internet Explorer browsers, a different profile for Mozilla browsers, and yet another profile for hand held mobile browsers. You would have complete control over all the HTTP options in each profile, to match the characteristics of these different Web browser types.

Although it is possible to use the default profiles, the best practice recommendation is to create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to this deployment, and ensures that you do not accidentally overwrite the default profile.

Persistence

Certain types of applications may require the same client returning to the same pool member, this is called persistence, or "stickiness". It can be configured using a persistence profile, and applied to the virtual server. For ORACLE Cloud Control services, persistence needs to be configured for every service, except for the Secure Upload service.

Rule

A rule is a user-written script that uses criteria to choose among one or more pools. In the BIG-IP software, it is called an iRule and provides a powerful and more granular level of control over traffic management. For an incoming request to a virtual server, the iRule is evaluated and selects the pool to which a request will be sent. For more information about F5 iRules, see the F5 DevCentral Web site

Enterprise Manager 12c Cloud Control:

Configuring OMS High Availability with F5 BIG-IP Local Traffic Manager
March, 2012
Authors: Mark McGill (Oracle), David Parker-Bastable (Oracle)

Contributing Authors: Farouk Abushaban (Oracle), James Viscusi (Oracle), Chris Akker (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

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together