CoyotePoint
Systems Inc

DEPLOYMENT GUIDE

Microsoft Office Communications Server 2007

&

Coyote Point Equalizer Deployment Guide

Table of Contents

Unified Communications Application Delivery

Microsoft® unified communications technologies use the power of software to deliver complete communications--messaging, voice, and video--across the applications and devices that people use every day. Integrating the experiences associated with the telephone--phone calls, voice mail, and conferencing--into the work done on a computer--documents, spreadsheets, instant messaging, e-mail, calendars--has the power to fundamentally change the way the world works.
A user's client provides presence information (presence state) via a network connection to the presence service, which is stored in what constitutes a personal availability record and can be made available for distribution to other users to convey availability for communication.

Microsoft Office Communications Server 2007 is the first Microsoft product to combine enterprise-ready IM (instant messaging), presence, conferencing, and VoIP (Voice over IP) telephony in a fully integrated unified communications solution. Office communications Server 2007 provides richer presence capabilities, enhanced support for group IM, and improved deployment and management than its predecessor, Microsoft Office Live Communications Server 2005 SP1[1]. To existing features, such as federation and public IM connectivity, Office Communications Server 2007 includes real-time conferencing hosted on servers inside the organization's firewall and a full-featured, software-powered VoIP solution that can stand on its own or integrate easily with an existing PBX infrastructure.

Office Communications Server 2007 extends the architecture of Live Communications Server 2005 to include components that support VoIP and conferencing. The key architectural features include:

- Pool configurations
- Front-End Servers
- Conferencing components

- VoIP components
- Perimeter network configuration and components
- Conference protocols
- Conference call flow

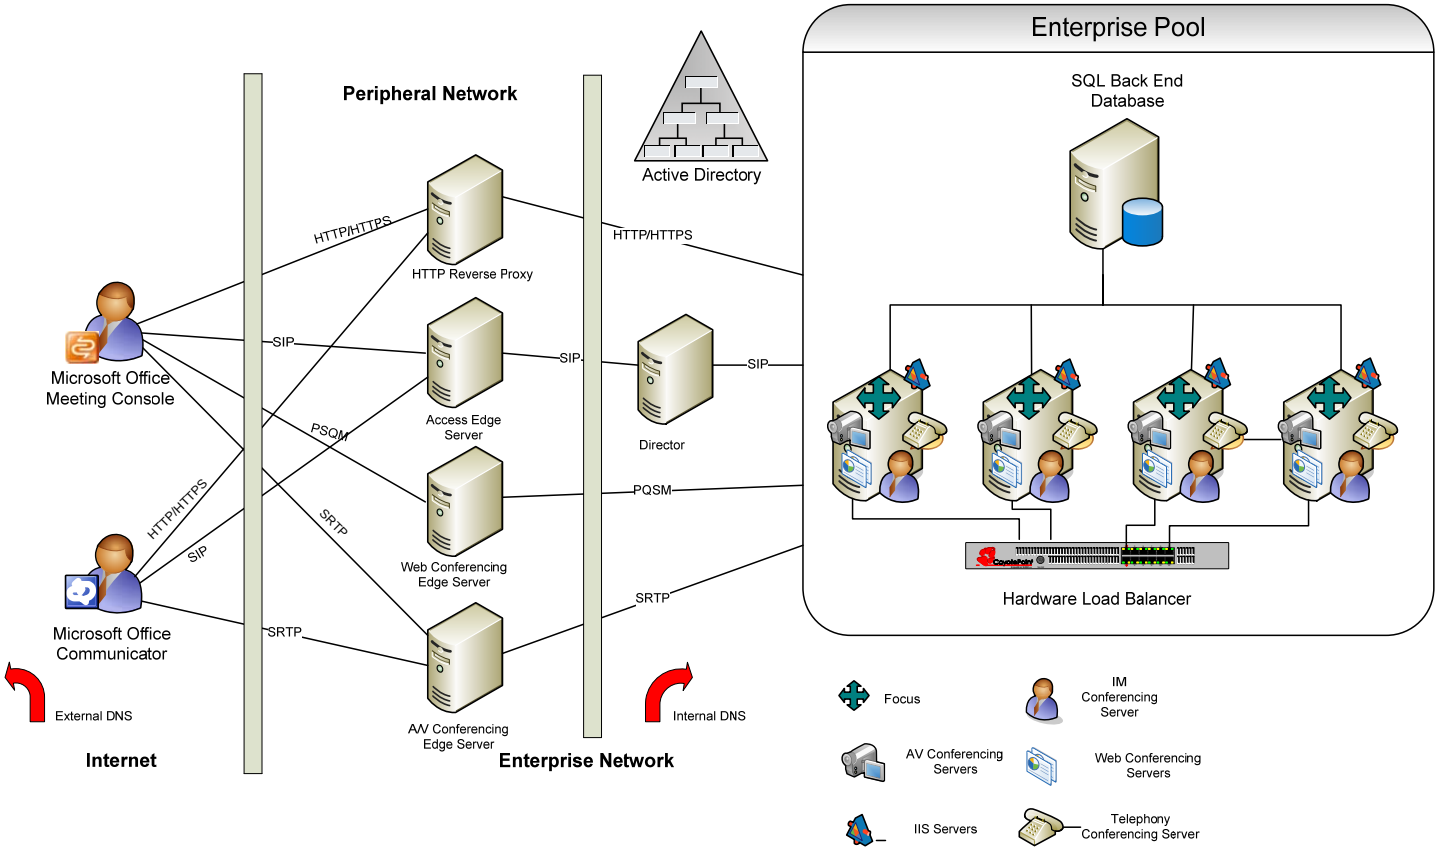

The diagram above shows an example of a complete OCS solution, with front-end servers in the Peripheral network providing access to the Enterprise Pool of OCS servers. In a simple enterprise configuration, a Coyote Point Equalizer is used to distribute traffic from the front-end servers to the server pool.

Coyote Point load balancers, when deployed in front of the Microsoft Office Communications Server 2007, increase application uptime and maximize server farm utilization. The load balancers receive all client requests, and distribute them efficiently to the “best” server among the available pool.

Equalizers consider server availability, load, response time, and other user-configured performance metrics when selecting a server for incoming client connections.

By performing sophisticated and customizable “health checks” to the servers and the Microsoft Office Communications Server 2007, Equalizer quickly identifies resource outages in real time and re-directs client connections to other available servers. Server capacity can be increased or

decreased on demand without impacting the applications and client connections. When demand grows, IT engineers can simply “slide” in new server resources and configure the Equalizer to use the new servers for client connections.

Coyote Point Equalizers are application aware and can inspect many types of application level content to perform intelligent switching of client requests to appropriate servers. Application switching eliminates the need to replicate content and application functions on all servers, and optimizes overall resource utilization, application performance and availability. Equalizers support balancing based on broad content types including URL, HTTP headers, HTTP cookies, SSL session IDs, and XML tags.

For implementations where session persistence across multiple TCP ports on the same server is a key requirement, a crucial benefit of Equalizer is its ability to ensure that a client stays with one real server so all real time information is preserved as the client continues to communicate across several application ports. Once a client connects to a server, successive transactions from the client are directed to the same server. The power of performance delivered by Equalizer ensures that your applications provide the best end-user response time and scalability at both Layer 4 and Layer 7, using sticky connections and cookie persistence.

Deployment Architecture

An Office Communications Server 2007 pool consists of one or more Front End Servers that provide IM, presence, and conferencing services and are connected to a SQL Server database for storing user and conference information.

Depending on the pool configuration, the database might reside on the same server. In addition, certain conferencing components might be deployed on the same physical computer, depending on the chosen pool configuration.

Office Communications Server 2007 offers three pool configurations: the Standard Edition configuration and the Consolidated (Figure 1) and Expanded Enterprise Edition (Figure 2) configurations.

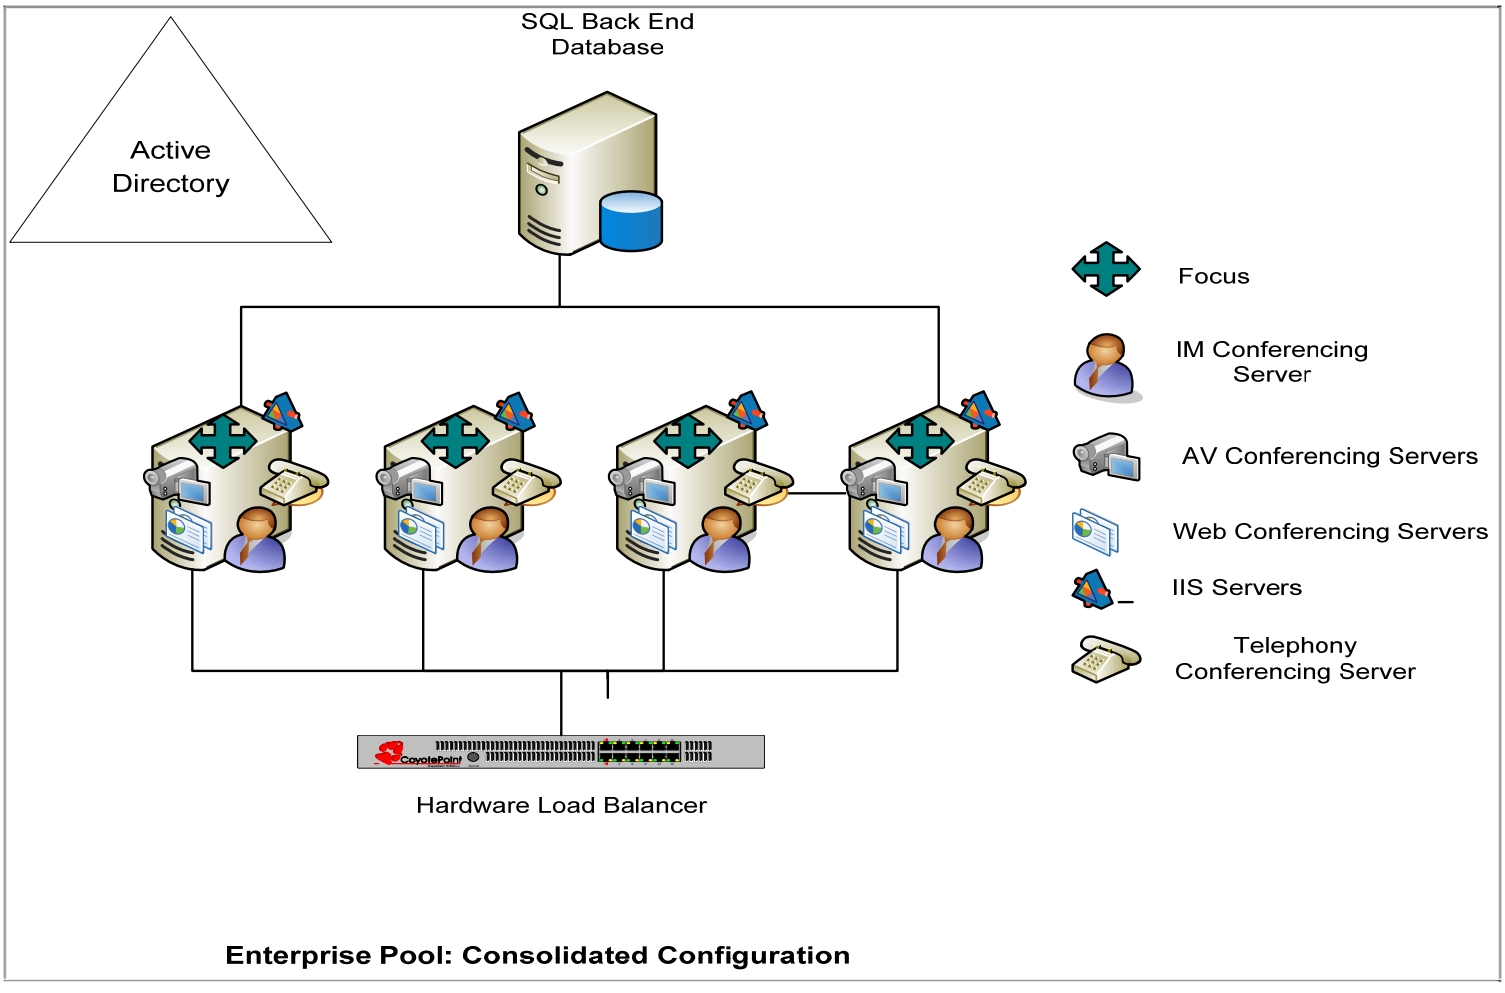

Figure 1

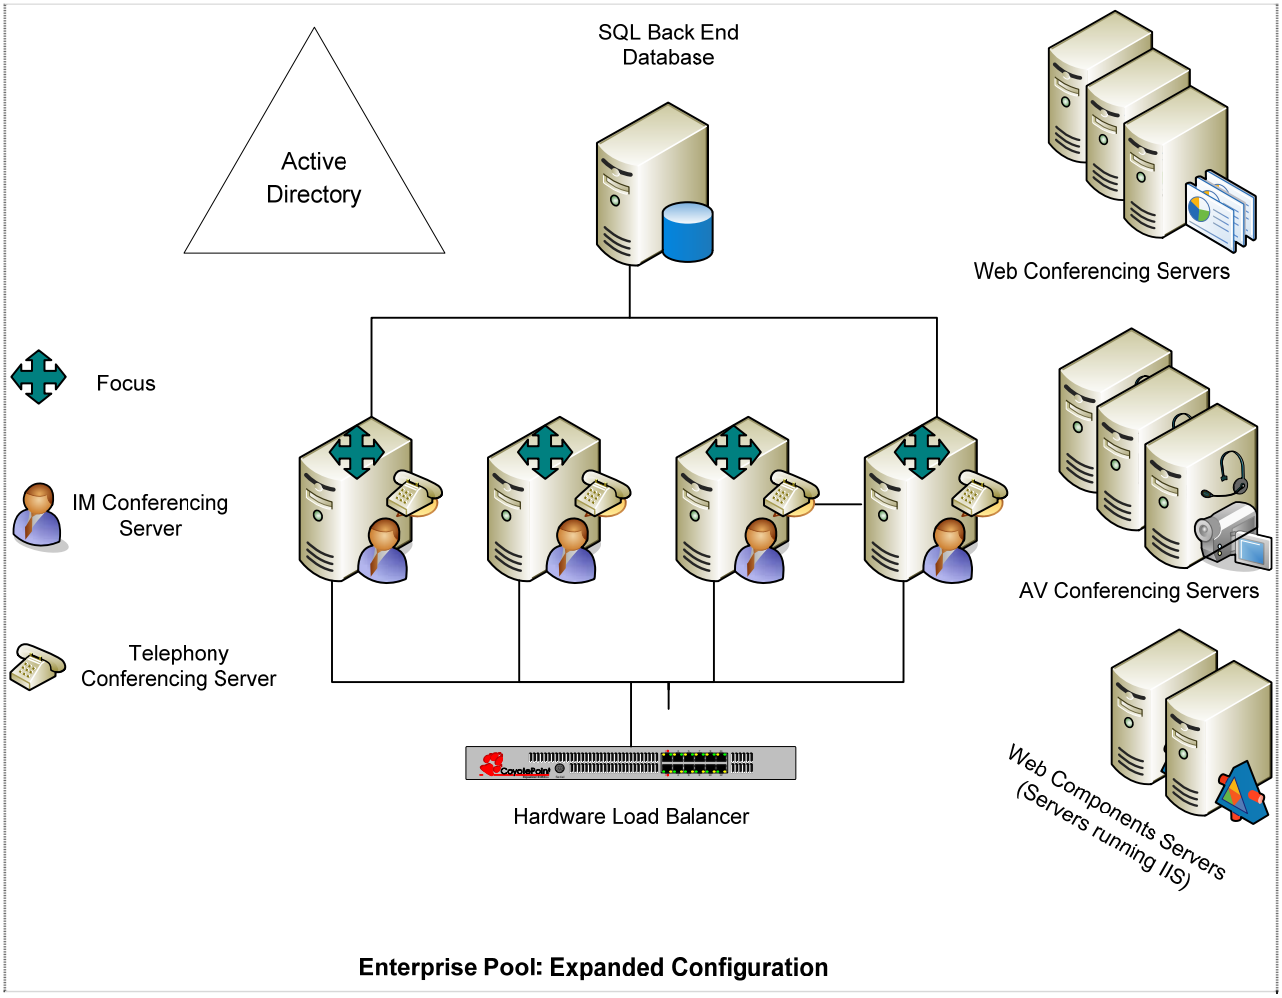

Figure 2

The Standard Edition configuration does not require load balancing as it only uses one server in the pool and will not be covered by this document.

Both the Enterprise Edition configurations consist of identical front end servers that are connected to a separate dedicated Microsoft SQL Server 2005 back-end database. (In an Enterprise pool, the back-end database must be on a dedicated computer, separate from all Enterprise Edition servers.)

General Requirements

An Office Communications Server 2007 Enterprise pool consisting of more than one Front End Server requires a hardware load balancer. If you are deploying a Standard Edition Server or a single

Enterprise Edition Front End Server, a load balancer is not required. A hardware load balancer is also required for arrays of Office Communications Server 2007 Edge Servers or an array of Standard Edition Servers configured as a Director. These requirements are summarized in the following table:

Deployment	Load Balancer Requirement
A Single Standard Edition Server	Load balancer not required
Enterprise Pool with multiple Front End Servers	Hardware load balancer required
Array of Directors	Hardware load balancer required
Array of Edge Servers	Hardware load balancer required

The following table summarizes the hardware load balancer ports that are required for Office Communications Server 2007.

Port Required	Virtual IP	Port Use
5060	Load balancer VIP used by the Front End Servers	Client to server SIP communication over TCP
5061	Load balancer VIP used by the Front End Servers	Client to Front End Server SIP communication over TLS SIP Communication between Front End Servers over MTLS
135	Load balancer VIP used by the Front End Servers	To move users and perform other "pool" level WMI operations over DCOM
444	Load balancer VIP used by the Front End Servers	Communication between the internal components that manage conferencing and the conferencing servers
443	Load balancer VIP used by the Web Components Server	HTTPS traffic to the pool URLs

The configurations provided in this document are configured for use to load balance groups of servers whether they are EE pools, access groups, or Director Servers. The provided configuration is provided for a single network configuration where the servers are not directly connected to the Equalizer and will require source NAT (Network Address Translation) to ensure return communication returns through the Equalizer. (Source NAT on Equalizer is enabled using the cluster **spoof** option.)

Equalizer Configuration

This installation will follow the load balancer requirements as determined by Microsoft in their guide titled 'OCS Planning guide' and detailed in step 7

The equalizer will be configured to support the enterprise pool, consolidated configuration diagram; this configuration is a basic cluster configuration for the Equalizer that will work with Enterprise Pool, Access Groups, and Director Pools by utilizing the standard ports used by these servers for the OCS 2007 applications detailed in the chart above. The version of Equalizer software tested is v8.5 loaded onto an Equalizer GX device.

It is assumed that the OCS configuration is installed, tested, and working properly according to Microsoft documentation. It is also assumed that Equalizer has been initially configured and licensed as explained in Chapters 2 and 4 of the Equalizer *Installation and Administration Guide*

The equalizer will have five clusters configured, one for each port utilised by OCS, each supporting four front end servers. Each cluster will have the same IP address, and this will be the resolvable IP address for OCS access. The Equalizer will have its own IP address for setup and maintenance. The IP numbering will be as follows;

Device	**IP Address**
Equalizer	192.168.10.100
Cluster 1 port 135	192.168.10.200
Cluster 2 port 443	192.168.10.200
Cluster 3 port 444	192.168.10.200
Cluster 4 port 5060	192.168.10.200
Cluster 5 port 5061	192.168.10.200
Server 1	192.168.10.101
Server 2	192.168.10.102
Server 3	192.168.10.103
Server 4	192.168.10.104

A diagram showing Equalizer's cluster and server configuration is shown on the next page.

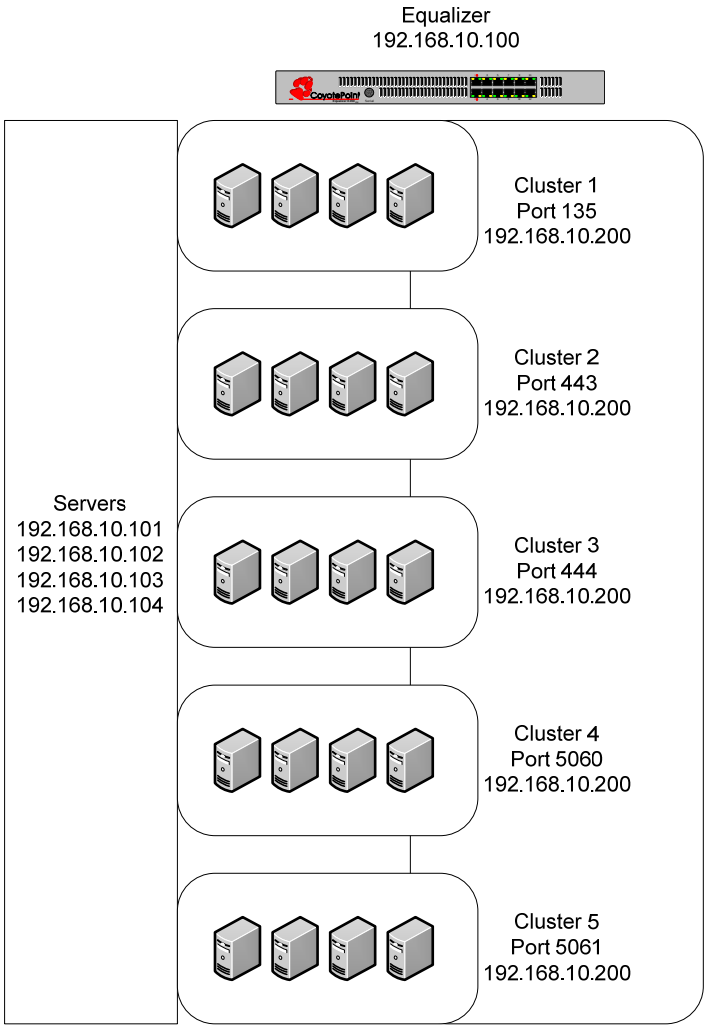

Cluster configuration

Create five L4 TCP clusters completing this information for each:

Cluster name:
Cluster IP address
Start port, end port
Idle timeout
Enable Spoof
Allow persistence
LB Algorithm

The following procedure shows you how to set up this configuration on Equalizer:

Step 1: Log into the Administrative Interface by pointing your browser at Equalizer's IP address. Add a new Layer 4 TCP cluster by doing the following:

a. Right-click on **Equalizer** in the left frame.
b. Select **Add Cluster** from the drop-down menu.
c. Select **Layer 4 TCP** and click the **Next** icon (**>**).
d. Set the **Start Port** to **135** (leave the end port blank). The **Add New Cluster** screen should now look like the screen shown below.
e. Click **Next** to confirm your settings.
f. Click **commit** to create the cluster.

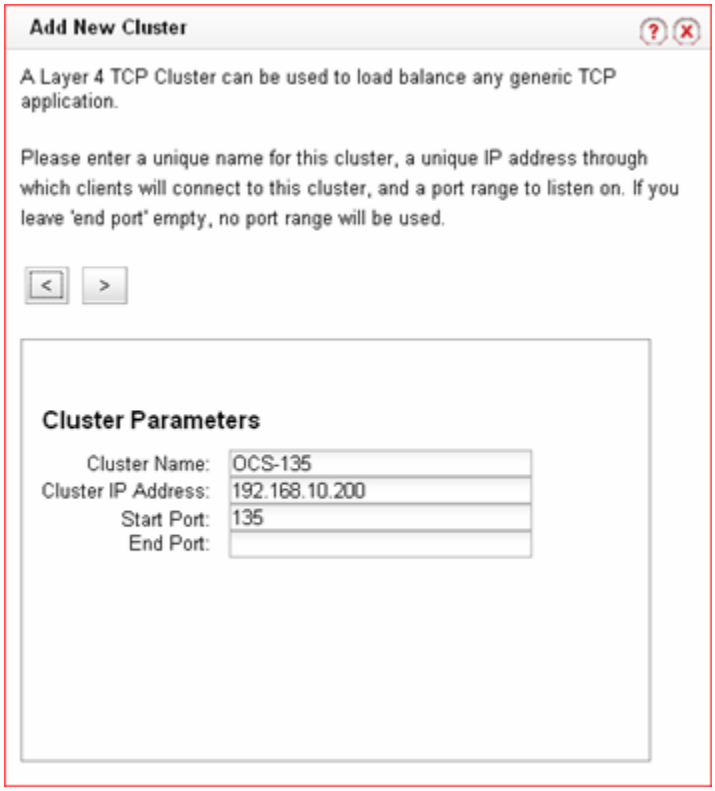

The new cluster appears in the left frame, and the right frame displays the **Configuration** tabs for the new cluster.

Step 2: On the **Required** tab

a) Set the **idle timeout** is set to a *minimum* of 1200 seconds to comply with the OCS requirements.
b) Ensure that the **spoof** check box is enabled (this is the default setting).
c) Click **commit**. The tab should now appear like the example below.

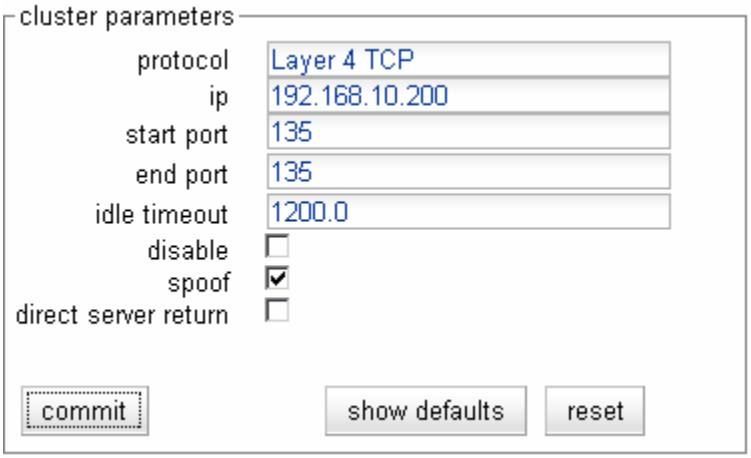

Step 3: Click on the **Persistence** tab and do the following:

a) Set the **sticky time** to 300 (seconds)**.** The **sticky time** is the number of seconds that Equalizer should "remember" connections from clients.
b) Enable the **inter-cluster sticky** checkbox. This is important, as it maintains port affinity across all the Layer 4 clusters you created.

Inter-cluster sticky is a Layer 4 persistence option that, when enabled, ensures that Equalizer attempts to direct requests from a particular client to the same server on another available port if the intended server port is unreachable. The Layer 4 clusters must have the same IP address, different ports, and a non-zero sticky time.

Inter-cluster sticky ensures that Equalizer directs requests from a particular client to the same server even if the incoming connection is to a different virtual cluster. When you enable inter cluster stickiness for a cluster, Equalizer checks the cluster for a sticky record as it receives each

connection request, just like it does for ordinary sticky connections. If Equalizer does not find a sticky record, Equalizer proceeds to check all of the other clusters that have the same IP address. If Equalizer still does not find a sticky record, it connects the user based on the current load balancing policy.

Step 4: Click on the **LB Algorithm** tab and select the **least connections** policy from the policy drop down. Click **commit**.

Least connections load balancing dispatches the highest percentage of requests to the server with the least number of active connections. Equalizer tries to avoid overloading the server so it checks the server's response time and server agent value. Least connections optimizes the balance of connections to servers in the cluster.

Step 5: Set up the servers in the cluster.

a) Right click on the cluster name in the left frame and select **Add Server** from the popup menu.

b) Type in the **Server Name**, **IP Address**, and accept the default **Port** (the same port you used for the cluster port).

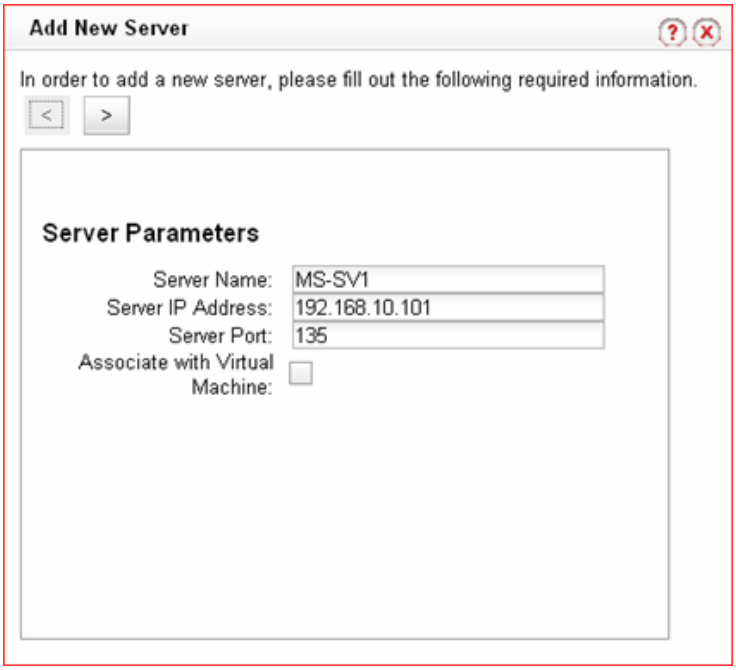

c) Click the Next button (**>**), and on the following screen click **commit** to create the server. The new server appears in the left frame, and the right frame displays the **Configuration** tabs for the new server.

d) **For servers using ports 135 and 444**, set the **probe port** to either **5060** or **5061** for determining server availability.

Attempting to monitor ports 135 and 444 on the servers will cause the load balancer to incorrectly detect these servers to be available because these ports are open even though Office Communications Server is not running.

e) Set **max connections** (the maximum number of permitted open connections for the server) to **60000** or less, if SNAT (**spoof**) is used.

Once this limit is reached, no more traffic is routed to the server until the number of open connections falls below this limit. (This limit is set by default to 0, which means that there is no maximum connections limit on the server.)

f) If your servers are all configured similarly and have similar processing capabilities, leave the **initial weight** set to the default of 100. Otherwise, set servers with lower capacities to lower numbers and more powerful servers to higher numbers (maximum: 200).

This number determines a starting point for the percentage of requests to route to each server and *is adjusted dynamically by Equalizer* as traffic is passed to the server and Equalizer examines server statistics (such as response time).

g) Click **commit** to save your settings. The tab should now look similar to this example:

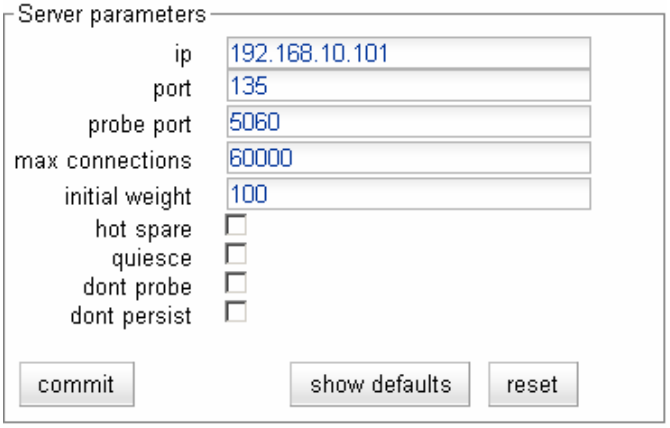

Step 6: Repeat **Step 5** for each additional server in the cluster.
Step 7: Repeat **Steps 4** through **Step 6** for each additional cluster in the configuration.

Step 8: Configure the routing tables on all servers so that Equalizer is the gateway for any outbound packets on the internal network. When **spoof** is enabled, servers in the cluster must use the Equalizer as the default gateway for routing. Configure the routing tables on all servers so that Equalizer is the gateway for any outbound packets on the internal network.

Enabling the **spoof** option causes Equalizer to use the client IP address as the source address when Equalizer routes a request to a server in a virtual cluster. This option is on by default. If you disable this option, the server receiving the request will see Equalizer's address as the source address -- the OCS servers need to see the client address. If the servers do not use Equalizer's internal address as the gateway when they send responses to clients, the reply packets will not be translated on their way to the client, causing the clients to reject the reply packets because they do not belong to an established connection. From the client side, it would look like the server was not responding. (Remember, the clients connect to Equalizer, not the servers, so in this configuration the servers need to respond through Equalizer.)

Step 9: Click on Equalizer in the left frame to display a cluster configuration summary. The final configuration should be similar to the screen below:

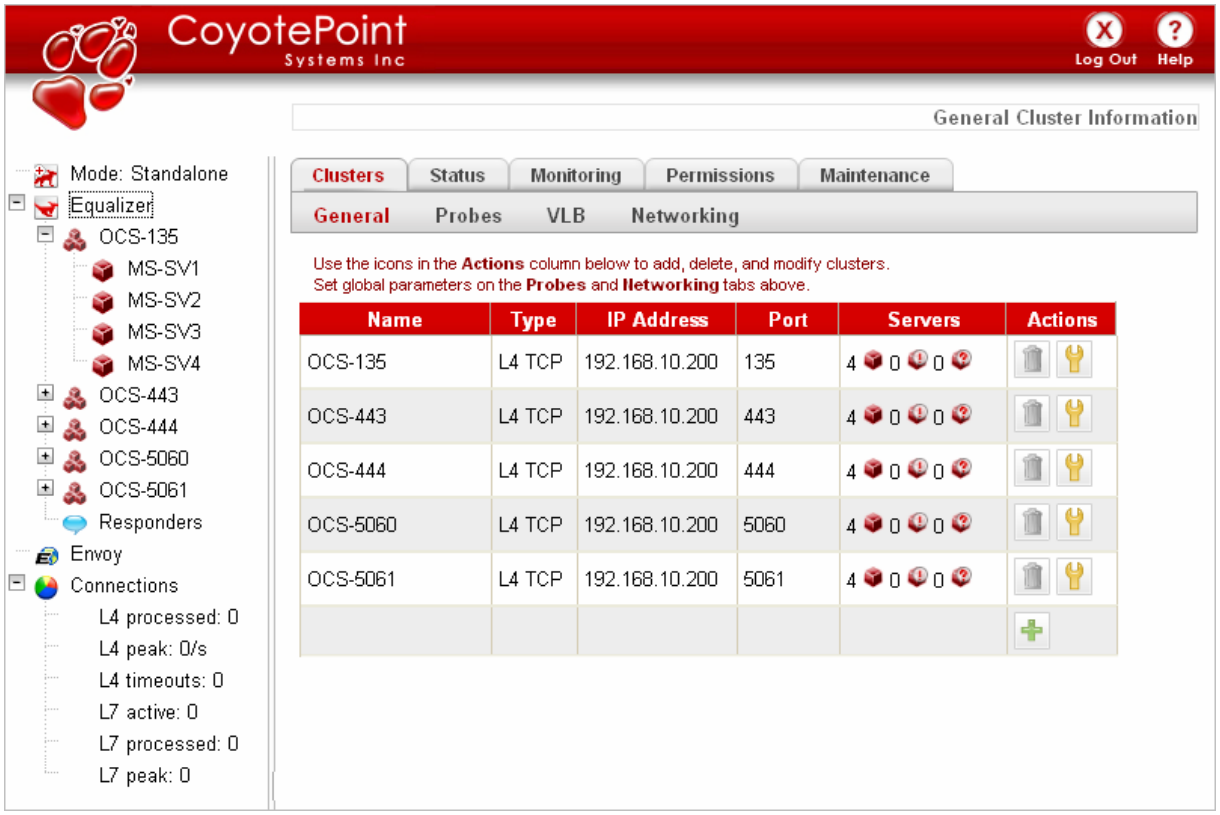

Testing Your Configuration

Once you have installed and configured Equalizer and your servers, perform tests to verify that your configuration is working properly:

1. Connect to the cluster IPs from multiple clients.

2. The numbers under **Connections** in the left frame of the Administrative Interface should increase as traffic flows through Equalizer. Click on **Connections** to look at detailed statistics.

3. Click **Equalizer > Cluster Summary** (shown below) and then click on each cluster name in the list displayed in the right frame. The numbers next to the servers in the clusters should increase as connections are load balanced to each server.

4. Check the logs on your application servers for errors.

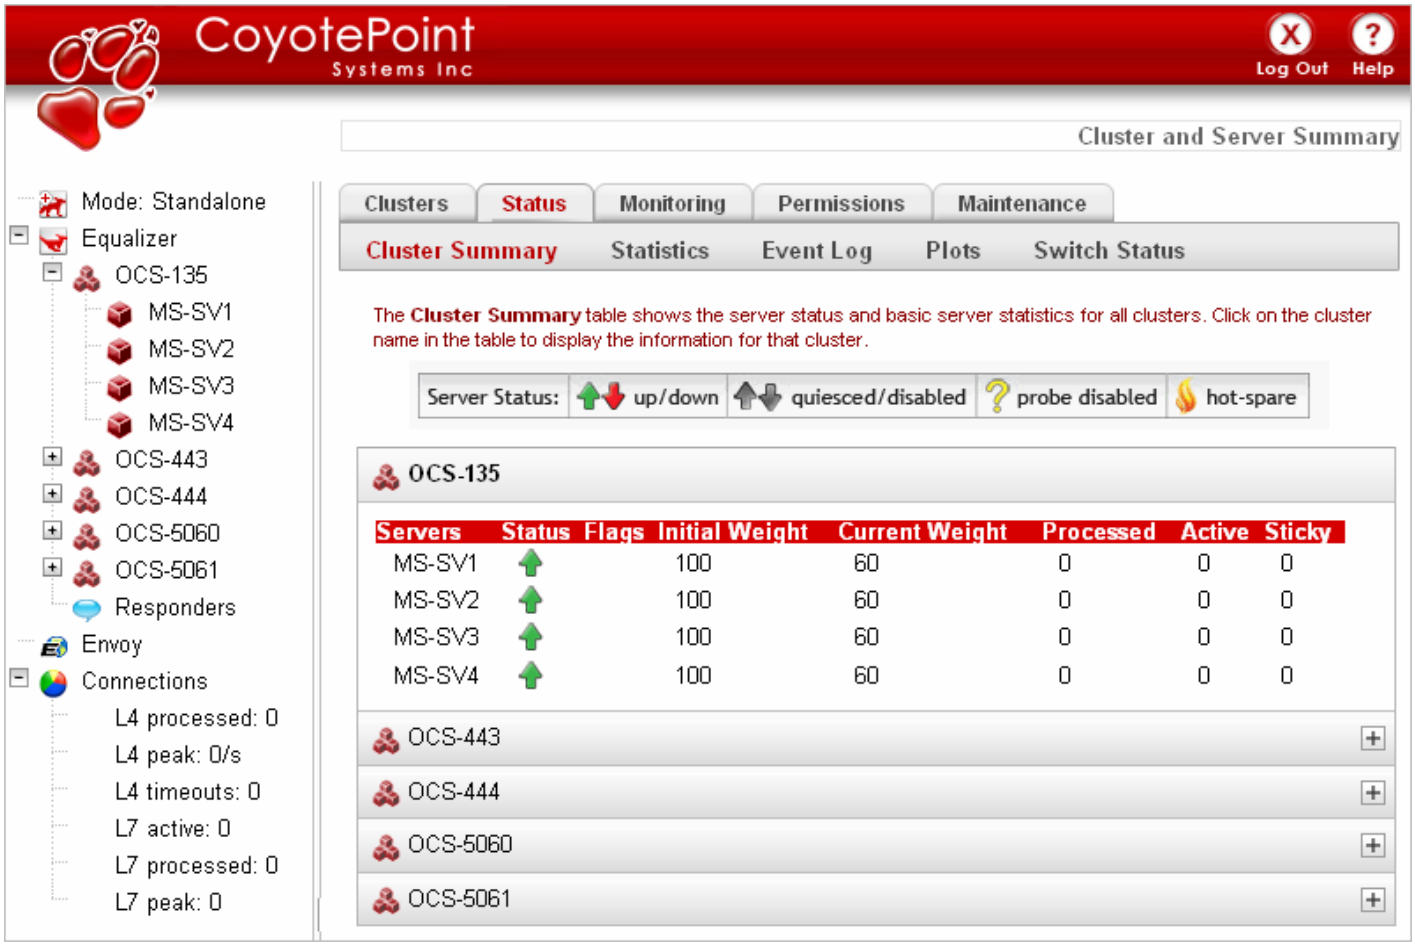

Using Equalizer in an Expanded OCS Configuration

Besides managing the OCS server pool, a number of other points in an expanded OCS configuration's Peripheral and Internal networks can benefit from Equalizer's application acceleration capabilities. These are illustrated in the following diagram (look for the boxes with black sides indicating Load Balancer placement). Equalizer can be deployed to load balance connections to multiple:

- Access edge servers
- A/V edge servers

- CWA application servers
- Directors

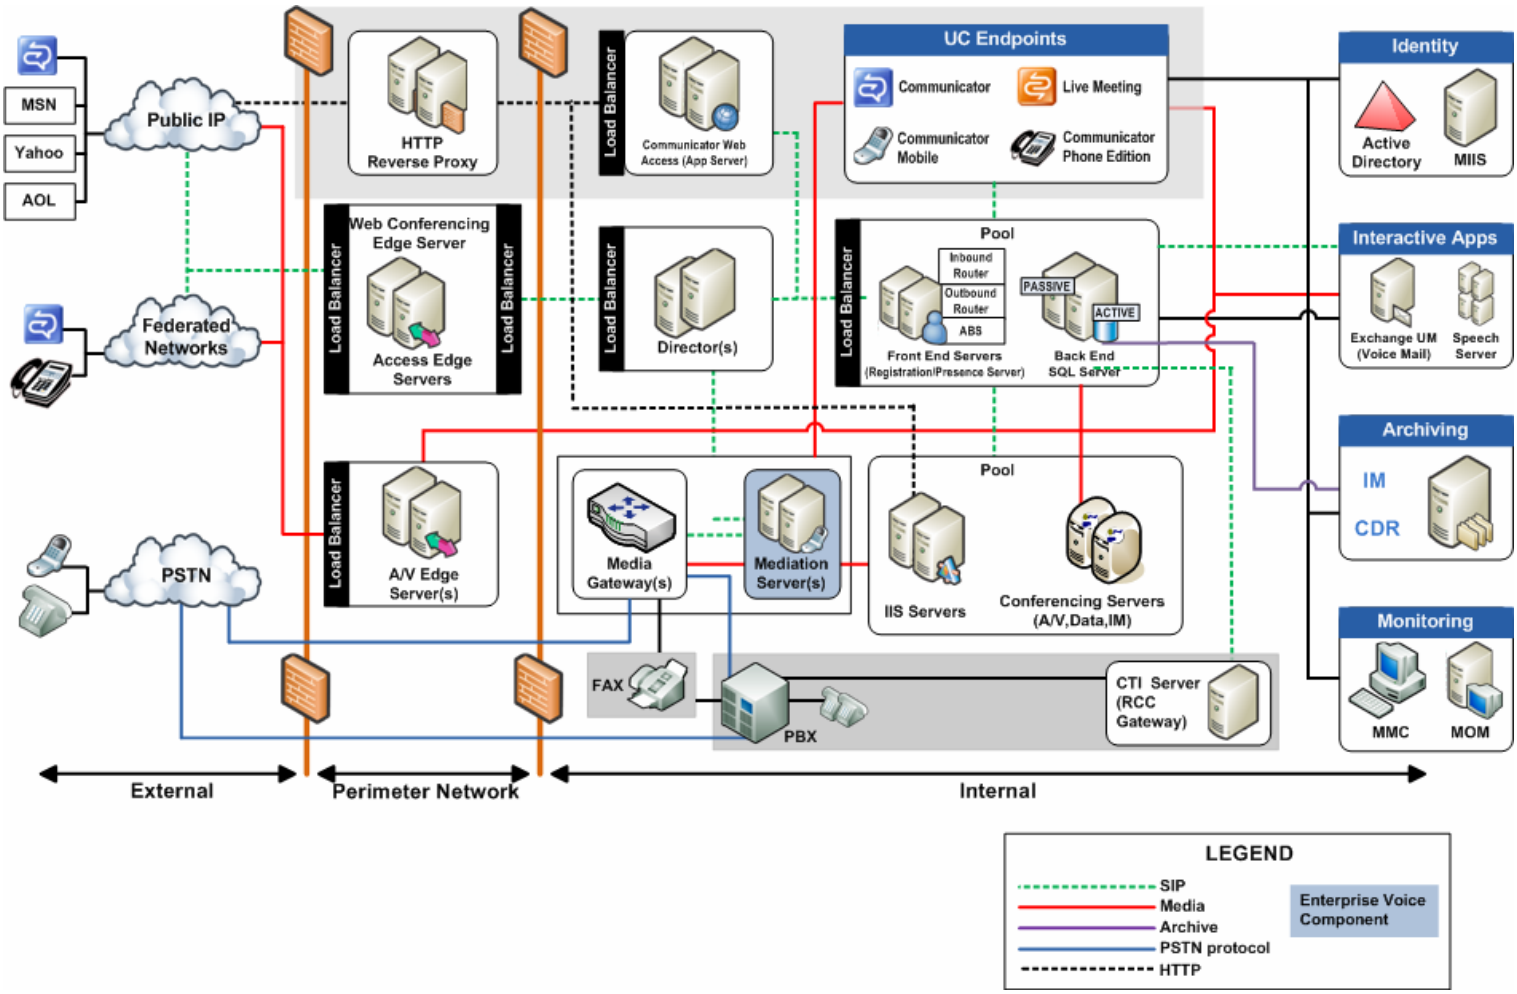

Summary

Microsoft states that hardware load balancing is a requirement for successful deployment of the Enterprise configurations of Office Communications Server. Coyote Point provides high application availability and load balancing technology optimized for the OCS environment enabling you to do just that.

This guide has presented a step-by-step example of how to use and configure Equalizer in an Enterprise Pool MS OCS environment. There are a number of other points in an expanded OCS configuration's Peripheral and Internal networks that can benefit from Equalizer's application acceleration capabilities. These include load balancing connections to multiple Access edge servers, A/V edge servers, CWA application servers, and Directors.

If you want more information on how Coyote Point can add value to your MS deployment plans please email info@coyotepoint.com or call 1-877-367-2696.

About Coyote Point

Coyote Point Systems Inc., the original load balancing technology pioneer, is a recognized leader in delivering affordable solutions for high application availability and accelerated application performance. With over a decade of experience, Coyote Point has delivered the eighth generation of its award-winning Equalizer Series platform. Equalizer Series load balancing and acceleration systems provide the industry's foremost combination of performance, affordability and ease of use. Over 8,500 customers worldwide have deployed Equalizer Series systems to ensure non-stop operations, accelerated delivery and on-demand scalability of business-critical applications and Websites. Coyote Point is headquartered in San Jose, CA and works with leading channel partners in the U.S., Canada, Europe, Asia, Australia and Africa.

Visit our website at http://www.coyotepoint.com.

