



EQUALIZER

Equalizer Installation and Administration Guide

Version 7.2
May 2005



Coyote Point Systems, Inc.
12 South First Street
Suite 616
San Jose, California 95113

Copyright © 1997-2005 Coyote Point Systems, Inc.

All Rights Reserved. Printed in the USA.

Equalizer is a trademark of Coyote Point Systems Incorporated. All other brand or product names are trademarks or registered trademarks of their respective companies or organizations.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

See Appendix F for complete License and Warranty information for this product.



Preface	ix
In This Guide	ix
Typographical Conventions	x
1 Overview	1
Introducing Equalizer	1
Overview of Equalizer	1
Intelligent Load Balancing	1
Load Balancing UDP Services	2
Maintaining Persistent Sessions	3
Layer 7 Load Balancing and Server Selection	4
Geographic Load Balancing	5
Configuring the Equalizer Network	8
Equalizer's Network Ports	8
Using Equalizer as a Gateway Between Networks	10
Using Equalizer in a Single Network Environment	11
Using a Second Equalizer as a Backup Unit	12
Using Reserved IP Addresses	14
Equalizer Configuration Worksheets	16
Standard Configuration Worksheet	16
Special Configuration Worksheet for Using Reserved IP Addresses ..	18
2 Installing Equalizer	19
Before You Install Equalizer	19
Stepping Through the Hardware Installation	19
3 Configuring Equalizer Hardware	21
Setting Up a Terminal or Terminal Emulator for Equalizer	21
Performing the Initial Configuration	21
Starting to Configure Equalizer	22
Configuring the Network Parameters	23

Committing Changes to the Configuration Parameters	24
Setting the Time Zone	25
Setting the Date and Time	25
Changing Equalizer's Console Password	25
Changing the Administration Interface Password	25
Upgrading Equalizer Software	25
Shutting Down Equalizer	26
Configuring Geographic Load Balancing	26
Configuring the Authoritative Name Server to Query Envoy	26
Using Geographic Load Balancing with Firewalled Networks	26
Configuring Servers	27
Configuring a Second Equalizer As a Backup	27
Testing Your Configuration	28
4 Accessing Browser Controls	29
Introducing the Equalizer Administration Interface	29
Accessing the Equalizer Administration Interface	29
Logging In	29
Navigating Through the Interface	30
5 Configuring Equalizer Operation	35
Licensing Your Equalizer	35
Setting Up a Failover Configuration	36
Enabling Outbound NAT	39
Enabling Passive FTP Connections	41
Managing Stale Connections	41
Enabling Sticky Network Aggregation	42
Configuring Custom Event Handling	44
Forwarding Equalizer Log Information	44
Specifying a Command to Run When a Particular Event Occurs	44
Changing Other System Parameters	45
Changing the Administration Passwords	47
Saving or Restoring Your Configuration	47
Saving Your Configuration	47
Backing Up Your Configuration	48

Restoring a Saved Configuration	48
Shutting Down Equalizer	49
Rebooting Down Equalizer	50
6 Monitoring Equalizer Operation	51
Displaying Equalizer Information	51
Displaying the Virtual Cluster Summary	53
Displaying the System Event Log	54
Displaying Cluster Information	56
Plotting Cluster Performance History	57
Displaying Server Information	58
Plotting Server Performance History	59
Displaying Geographic Cluster Parameters	61
Plotting Geographic Cluster Performance History	62
Displaying Site Information	63
Plotting Site Performance History	64
7 Administering Virtual Clusters	65
Working with Virtual Clusters	65
Adding a Virtual Cluster	65
Deleting a Virtual Cluster	70
Configuring a Cluster's Load-Balancing Options	70
Providing FTP Services on a Virtual Cluster	73
Configuring a Cluster to Use Server Agents	73
Enabling Persistent Sessions	74
Using Active Content Verification (ACV)	76
Using Secure Server Certificates for HTTPS Clusters	79
Managing Servers	81
Adding a Server to a Cluster	81
Deleting a Server	83
Adjusting a Server's Static Weight	83
Shutting Down a Server Gracefully	85
8 Match Rules.....	87
Overview of Match Rules	87

General Match Expressions and Match Bodies	88
Match Expressions	88
Match Bodies	89
Match Rule Example	89
Constructing Match Rules	90
Viewing the Default Match Rule	90
Defining a Match Rule	91
Modifying a Match Rule	93
Removing a Match Rule	94
Match Functions	94
Common Match Functions	94
HTTP Protocol and Request URI Match Functions	95
HTTP Header Matching Functions	95
HTTPS Specific Match Functions	96
9 Administering Geographic Clusters	99
Geographic Load Balancing with Envoy	99
Installing and Configuring Envoy	102
Installing Envoy	102
Configuring the Authoritative Name Server to Query Envoy	102
Using Envoy with Firewalled Networks	103
Working with Geographic Clusters	104
Adding a Geographic Cluster	104
Configuring a Geographic Cluster's Load-Balancing Options	105
Deleting a Geographic Cluster	105
Working with Sites	106
Adding a Site to a Geographic Cluster	106
Adjusting a Site's Static Weight	107
Deleting a Site from a Geographic Cluster	107
A Using Server Agents	109
Introducing Server Agents	109
Custom Server Agents	109
B Using Reserved IP Addresses	111
C Regular Expression Format	113

Terms	113
Learning About Atoms	113
Creating a Bracket Expression	114
Matching Expressions	115
D Mini SendMail	117
Syntax	117
Flags	117
Example	117
E Troubleshooting	119
Equalizer Doesn't Boot	119
Clients Time Out While Trying to Contact a Virtual Cluster	119
Backup Equalizer Continues to Boot	120
Can't View Equalizer Administration Pages	120
Equalizer Administration Page Takes a Long Time to Display	120
Equalizer Doesn't Respond to Pings to the Admin Address	120
Browser Hangs When Trying to Connect Via FTP to an FTP Cluster	121
Return Packets from the Server Aren't Routing Correctly	121
Web Server Cannot Tell Whether Incoming Requests Originate Externally or Internally	121
F License and Warranty	123
Glossary	127
Index	135



The *Equalizer Installation and Administration Guide* is intended for people who are installing, configuring, or administering an Equalizer™ system.

In This Guide

This guide contains the following chapters and appendices:

- Chapter 1, *Overview*, contains detailed descriptions of Equalizer concepts and terminology. This chapter includes information to help you plan your Equalizer configuration. If you are setting up Equalizer for the first time, be sure to read the *Overview* chapter before attempting to install and configure your system.
- Chapter 2, *Installing Equalizer*, provides comprehensive instructions for installing Equalizer.
- Chapter 3, *Configuring Equalizer Hardware*, instructs you in setting up Equalizer to work with your networks and servers.
- Chapter 4, *Accessing Browser Controls*, discusses how to use Equalizer's HTML-based administration interface to check the current Equalizer status and to change settings within Equalizer.
- Chapter 5, *Configuring Equalizer Operation*, provides information on modifying Equalizer's configuration through the Equalizer Administration Interface, including setting up a failover configuration.
- Chapter 6, *Monitoring Equalizer Operation*, describes how to view information, statistics, and graphical displays about Equalizer's operation.
- Chapter 7, *Administering Virtual Clusters*, tells you how to add and remove virtual clusters and servers, changing load balancing options, and shutting down servers.
- Chapter 8, *Match Rules*, shows you to create match rules that distribute requests based on a request's attributes.
- Chapter 9, *Administering Geographic Clusters*, shows you how to use the Envoy add-in to add and remove geographic clusters and sites and change geographic load balancing and targeting options.
- Appendix A, *Using Server Agents*, describes how to develop custom server agents.
- Appendix B, *Using Reserved IP Addresses*, describes how to configure Equalizer to distribute requests to servers assigned IP addresses on reserved, non-routable networks.
- Appendix C, *Regular Expression Format*, discusses Equalizer's regular expressions, components, formats, and usage.
- Appendix D, *Mini_SendMail*, describes and documents the mini_sendmail program and its flags.

- Appendix E, *Troubleshooting*, helps you to diagnose Equalizer installation and configuration problems.
- Appendix H, *License and Warranty*, contains the complete License and Warranty information.
- The glossary defines the technology-specific terms used throughout this book.

Typographical Conventions

The following typographical conventions appear throughout this guide:

Italics indicates the introduction of new terms, is used to emphasize text, and indicates variables.

Boldface text highlights field, key, or button names in instructions.

`Courier` text denotes commands, file names, directory names, keywords, and syntax from text.

1. Numbered lists show steps that you must complete in the numbered order.
- ✓ Bulleted lists identify items that you should verify or procedures you should use to resolve particular problems.

Note – Notes highlight important information and special considerations.

Caution – Caution notes warn when an action could result in loss of data or damage to your equipment.



An attention icon emphasizes information critical to Equalizer operation.



Introducing Equalizer

This chapter provides an overview of Equalizer's features and discusses some common configurations.

Overview of Equalizer

Equalizer™ is a high-performance content switch that features:

- Intelligent load balancing based on multiple, user-configurable criteria.
- Real-time server and cluster performance monitoring.
- Server and cluster administration from a single interface.
- Session persistence using cookies or IP addresses
- Hot-backup configurations (requires a second Equalizer) featuring no single point of failure.
- Layer 7 (L7) content-sensitive routing.
- Geographic load balancing (requires the optional Envoy add-in).

See Coyote Point's website, www.coyotepoint.com, for a current list of products and their features.

Intelligent Load Balancing

Equalizer functions as a gateway to one or more sets of servers known as *virtual clusters*. When a client submits a request to a site that Equalizer manages, Equalizer identifies the virtual cluster for which the request is intended, determines the server in the cluster that will be best able to handle the request, and forwards the request to that server for processing.

To route the request, Equalizer modifies the header of the request packet and forwards the modified packet to the selected server. When operating in Layer 7, Equalizer can evaluate and, in some cases, modify the contents of both the request and response headers.

To determine the best server to route a request to, Equalizer uses intelligent load-balancing algorithms that take into account the configuration options set for the cluster and servers, real-time server status information, Layer 7 rules, and information from the request itself.

Load Balancing Configuration

When you configure your virtual cluster, you can select one of the following load-balancing algorithms to control how Equalizer balances the load across your servers: round robin, static weight, adaptive, fastest response, least connections, or server agent.

When you configure the servers in a virtual cluster, you assign a static weight between 20 and 200 for each server. When you select one of the adaptive load-balancing algorithms, Equalizer uses the

servers' static weights as a starting point to determine the percentage of requests to route to each server. Each server handles a percentage of the total load based on its fraction of the total weights in the server cluster. Equalizer dynamically adjusts server weights according to real-time conditions to ensure that Equalizer routes requests to the server that is best able to respond. A server with a weight of zero (0) is considered down or unavailable: Equalizer does not route new requests to servers in this state.

Real-Time Server Status Information

Equalizer can gather real-time information about a server's status using Server Agents and Active Content Verification.

You can install *server agents* on each server to provide Equalizer with periodic performance statistics. This enables Equalizer to adjust the dynamic weights of the servers in a cluster according to their actual performance characteristics. If the server is overloaded and you have enabled adaptive load balancing, Equalizer responds by reducing the server's dynamic weight so that the server receives fewer requests. Coyote Point provides APIs useful for creating these agents. For more information see "Using Server Agents" on page 109.

Equalizer's Active Content Verification (ACV) provides a way to check the validity of a server's response using most network services that support a text-based request/response protocol, such as HTTP. When you enable ACV for a cluster, Equalizer requests data from each server in the cluster (using an *ACV Probe string*) and verifies the returned data (against an *ACV Response string*). However, you cannot use ACV with UDP-based services. If Equalizer receives no response or the response string is not in the response, the verification fails and Equalizer stops routing new requests to that server.

Network Address Translation and Spoofing

Equalizer's NAT subsystem distributes incoming Layer 4 or Layer 7 (with spoofing) client requests among the available servers. The NAT subsystem records the existence of the request, selects the best available server, rewrites the TCP/UDP and IP headers of the request packet, and then forwards the translated packet to the selected server. Because the servers are configured to use Equalizer to gateway all packets, Equalizer performs the reverse translation as the server response packets leave the cluster. For more information about configuring spoofing see "Adding a Virtual Cluster" on page 65.

When IP spoofing is enabled, the servers see their client's actual IP address. However any response must be gatewayed through the Equalizer because clients will only recognize the Equalizer's address—they did not communicate directly with the server.

When Equalizer receives an incoming packet that is not destined for a virtual cluster address, Equalizer passes the packet through unaltered. Similarly, when Equalizer receives an outgoing packet that is not a response to an existing virtual cluster connection, Equalizer passes the packet through to the external network.

Load Balancing UDP Services

You can configure Equalizer Virtual Clusters to provide load balancing and server failure detection for many UDP (User Datagram Protocol) based services. UDP load balancing is ideal for stateless protocols such as DNS and RADIUS, can load-balance WAP (Wireless Application Protocol) gateways, and can even load-balance certain types of NFS server cluster that provide a single-system image.

Equalizer does not support Active Content Verification for UDP clusters.

Maintaining Persistent Sessions

Maintaining persistent sessions is useful when state information is shared between a client and server. For example, a web-based shopping cart application may depend on persistent session information between the client and server. Maintaining persistent sessions is necessary because of the information shared between a client and server: the details in the shopping cart potentially need to persist across many individual TCP connections. Equalizer supports two mechanisms for maintaining persistent sessions: Cookie-based and IP-address-based persistence.

Cookie-Based Persistence

Equalizer can use cookie-based persistence for HTTP and HTTPS clusters that support Layer 7 load balancing. In cookie-based persistence, Equalizer “stuffs” a cookie into the server’s response header on its way back to the client. This cookie uniquely identifies the server to which the client was just connected. The client includes (sends) the cookie in subsequent requests to the Equalizer. Equalizer uses the information in the cookie to route the requests back to the same server.

Equalizer can direct requests from a particular client to the same server, even if the connection is to a different virtual cluster. For example, if a user switches from an HTTP cluster to an HTTPS cluster, the persistent cookie will still be valid if the HTTPS cluster contains a server with the same IP address.

If the server with which a client has a persistent sessions is unavailable, Equalizer automatically selects a different server. Then, the client must establish a new session; Equalizer stuffs a new cookie in the next response.

IP-Address Based Persistence

For generic TCP and UDP clusters that support Layer 4 load balancing, Equalizer supports IP-address based persistent sessions. When Equalizer enables its *sticky connections* feature, Equalizer identifies clients by their IP addresses when they connect to a cluster. Equalizer routes requests received from a particular client during a specified period of time to the same server in the cluster.

A *sticky timer* measures the amount of time that has passed since there was a connection from a particular IP address to a specific cluster. The sticky time period begins to expire as soon as there are no longer any active connections between the client and the selected cluster. Equalizer resets the timer whenever a new connection occurs. If the client does not establish any new connections to the same cluster, the timer continues to run until the sticky time period expires. At expiration, Equalizer handles any new connection from that client like any other incoming connection and routes to an available server based on the selected load-balancing criteria.

To correctly handle sticky connections from ISPs that use multiple proxy servers to direct user connections, Equalizer supports *sticky network aggregation* with which only the network portion of a client’s IP address maintains a sticky connection. Sticky network aggregation directs the user to the same server no matter which proxy he or she connects through.

You can also configure Equalizer to ensure that it 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 incoming request.

Layer 7 Load Balancing and Server Selection

Equalizer's support for Layer 7 content-sensitive load balancing enables administrators to define rules for routing HTTP and HTTPS requests, depending on the content of the request. Layer 7 load balancing routes requests based on information from the application layer. This provides access to the actual data payloads of the TCP/UDP packets exchanged between a client and server. For example, by examining the payloads, a program can base load-balancing decisions for HTTP requests on information in client request headers and methods, server response headers, and page data.

Equalizer's Layer 7 load balancing allows administrators to define rules in the administration interface for routing HTTP and HTTPS requests according to the request content. These rules are called *match* rules. For example, you can use Layer 7 rules to specify routing preferences such as,

```
“Send all requests for graphics files to servers A, B and E;  
send all requests for Perl scripts to servers C and D; and  
send all other requests to server Z.”
```

This enables administrators to create extremely flexible cluster configurations. Administrators can use Layer 7 technology to implement client-server persistence based on HTTP cookies.

For HTTP requests, Layer 7 load balancing can make decisions based on the following:

- HTTP protocol version
- Host name
- Pathname of the request
- Filename of the request
- Pattern matches against arbitrary HTTP request headers

Go to “Match Functions” on page 94 for a complete list of match functions.

For HTTPS requests, load balancing decisions can be based on the SSL protocol level the client uses to connect.

Geographic Load Balancing

The optional Envoy add-on supports geographic load balancing, which enables requests to be automatically distributed across Equalizer sites in different physical locations. An Equalizer *site* is a cluster of servers under a single Equalizer's control. A *geographic cluster* is a collection of sites that provide a common service, such as Web sites. The various sites in a geographic cluster can be hundreds or even thousands of miles apart. For example, a geographic cluster might contain two sites, one in the eastern U.S. and one on the U.S.'s west coast (Figure 1).

Geographic load balancing can dramatically improve reliability by ensuring that your service remains available even if a site-wide failure occurs. Equalizer can also improve performance by routing requests to the location with the least network latency.




Figure 1 Geographic cluster with two sites

Geographic Load Balancing Routing

Envoy routes each incoming request to the site best able to handle it. If a site is unavailable or overloaded, Envoy routes requests to the other sites in the geographic cluster. When you enable geographic load balancing, Envoy directs incoming client requests to one of the sites in the geographic cluster based on the following criteria:

- **Availability:** If a site is unavailable due to network outage, server failure, or any other reason, Equalizer stops directing requests to that site.
- **Performance:** Envoy tracks the load and performance at each site and uses this information to determine the site that can process the request most efficiently.
- **Distance:** Envoy notes the site that is *closest* to the client (in network terms) and offers the least network latency.

Distributing the Geographic Load

Envoy uses the Domain Name System (DNS) protocol¹ to perform its geographic load distribution. DNS translates fully-qualified domain names such as `www.coyotepoint.com` into the IP addresses that identify hosts on the Internet. Envoy configures the authoritative name server for the domain to query the Equalizers in the geographic cluster to resolve the domain name. When Envoy receives a resolution request, it uses the load-balancing algorithms configured for the geographic

1. For more information about DNS, see Paul Albitz and Cricket Liu, *DNS and BIND*, 3rd ed. (O'Reilly & Associates, 1998).

cluster to determine the site that is best able to process the request and then returns the address of the selected site.

For example, the geographic cluster `www.coyotepoint.com` might have three sites (see Figure 2)—one on the east coast of the U.S., one in the west coast of the U.S., and one in Europe—each with an Equalizer (each with the Envoy add-on) and clusters of servers.




Figure 2 Three-site geographic cluster configuration

When a client in California attempts to connect to `www.coyotepoint.com`:

1. The client queries its local name server to resolve the domain name (see Figure 3).




Figure 3 Client queries its local DNS for `coyotepoint.com`

- The local name server queries the authoritative name server for `coyotepoint.com` (see Figure 4).




Figure 4 Client's local DNS queries the authoritative name server for `coyotepoint.com`

- The authoritative name server provides a list of Envoy-enabled Equalizers and returns this list to the client's local DNS (see Figure 5).

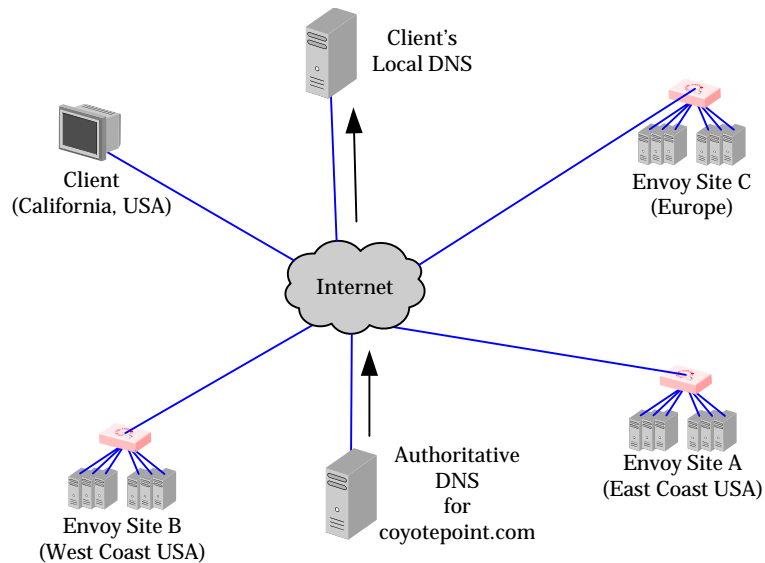


Figure 5 The authoritative name server for `coyotepoint.com` returns a list of delegates

4. The client's DNS selects one of the Equalizers in the list and queries it. If the queried site doesn't respond, the client tries each of the other sites.
5. Envoy returns the returns the IP Address of the virtual cluster best able to handle the client's request.

Configuring the Equalizer Network

Equalizer is a versatile traffic management solution. It works in a single or dual network mode. If you have a second unit, you can use it as a hot-backup unit. Equalizer also works with servers placed on a reserved, non-routable network and allows for IP address aliasing.

You can use Equalizer in a number of configurations. Before you install Equalizer, you need to determine where it will fit into your network and how you will configure it. This section describes some configuration choices. The following section provides a worksheet to help you plan your configuration.

Equalizer's Network Ports

All Equalizers have two types of network ports: *external* and *server*. The external port is always a single port labeled **External** or **Ext**. The server ports are labeled **Int** on dual-port models or labeled with numbers on switch-based models. Depending on the Equalizer switch-based model, there may be four or more of these ports.




Figure 6 Equalizer E450si

Equalizer's External Port

The external port is connected to the network to which the client machines and possibly the Internet or an Intranet are connected. This *external network* receives the client request packets that Equalizer distributes across the available servers. Equalizer also uses the external network to transmit response packets to clients. This port is only used for dual network (external and internal) configurations and single network configurations on dual-port models. It is not used for single network configurations on multi-server port models, see "Using Equalizer in a Single Network Environment" on page 11 for more information.

Hosts or routers on the external network can have routes to the internal network that are gatewayed through Equalizer's external address. Equalizer's external address is also its *administration address*, the IP address used to connect to Equalizer's browser-based administration interface.

Note – When using dual-port Equalizers in single network mode, use the external port to connect to the network to which the client machines, Intranet, or Internet are connected.

Equalizer's Server Port

Servers that process the incoming requests connect to the server ports: either directly or through a network device such as a switch. These physical servers provide services on specific IP addresses and ports and are organized into clusters. Equalizer's load-balancing subsystem translates client request packets and then forwards them to the selected server. When a server machine sends a response packet back to a client, Equalizer processes it and forwards it to the appropriate client across the external network.

When using Equalizer with NAT in layer 4 or spoofing in layer 7, you must configure the servers' routing tables so that Equalizer is the gateway for any outbound packets that leave the internal network. 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.). If you are using Equalizer without spoofing, you do not need to use Equalizer as a gateway.

When using Equalizer in single network mode, the client machines, Intranet, or Internet must connect to one of the server ports. In this instance one of the server ports is the external port.

Using Equalizer as a Gateway Between Networks

The most common Equalizer configuration is to have Equalizer function as the gateway between two separate networks—the internal network where the servers reside and the external network on which clients and the Internet or an Intranet reside. Figure 7 shows this configuration in detail.

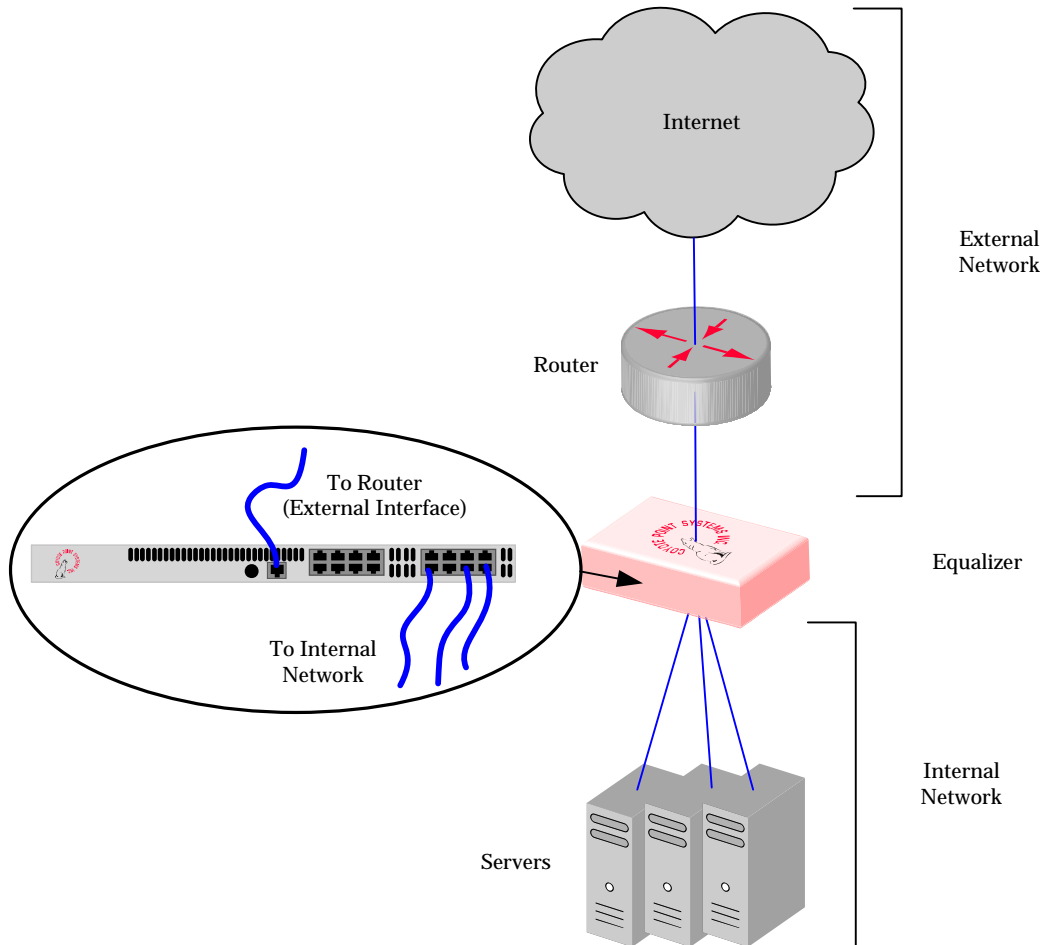


Figure 7 Sample two-network configuration

Using Equalizer in a Single Network Environment

If you do not want to split your network into internal and external networks, you can configure Equalizer to use a single-network mode, effectively placing both the clients and servers on the same network. Figure 8 on page 11 shows this configuration in detail. Certain protocols that use dynamic port mapping or multiple TCP/UDP ports work best in a single network environment. For example, use a single network configuration if you need your servers on your internal network to communicate with a Windows file server or a machine running pcAnywhere™.

You implement single-network configurations differently depending on the Equalizer model.

For switch-based Equalizer models, connect one of Equalizer's server ports to the network and do not use the external port. Servers connect to the other server ports as usual. You must configure servers, which must have valid network addresses on the external network, to use Equalizer's internal address as the gateway for outbound packets. You do not configure an IP address on the external port when using a single network configuration.




Figure 8 Sample single network configuration for a switch-based Equalizer

For dual-port Equalizer models, the reverse is true. You leave the server (INT) port disconnected and connect the external (EXT) port to a switch that maintains the connections to the servers and to the external network.

Most operating systems allow you to specify a network route (gateway) for packets destined for specific hosts. If you want your virtual clusters to accept connections from clients on the same network as the servers, you must configure the servers to route packets destined for these clients through Equalizer. Clients on the local network that do not have such routes configured connect to the server's IP address and not through a virtual cluster (that is, not routed through Equalizer).

Using a Second Equalizer as a Backup Unit

You can configure a second Equalizer as a backup unit that will take over in case of failure. This is known as a *hot-backup* configuration. The two Equalizers are siblings, the *primary* unit and the *backup* unit. If the primary Equalizer stops functioning, the backup unit adopts the primary unit's IP addresses (clusters) and begins servicing connections. In a failover configuration, the servers in a virtual cluster use a separate *failover alias* as their default gateway, rather than the IP address of the cluster or external port on a particular Equalizer. The failover alias migrates between the primary and backup unit as needed, automatically ensuring that the servers have a valid gateway in the event of a failure.

In a hot-backup configuration, both the primary and backup Equalizers are connected to the same networks; the backup unit's cluster and external ports must be connected to the same hubs or switches to which the primary Equalizer's ports are connected. Figure 9 shows a sample failover configuration.

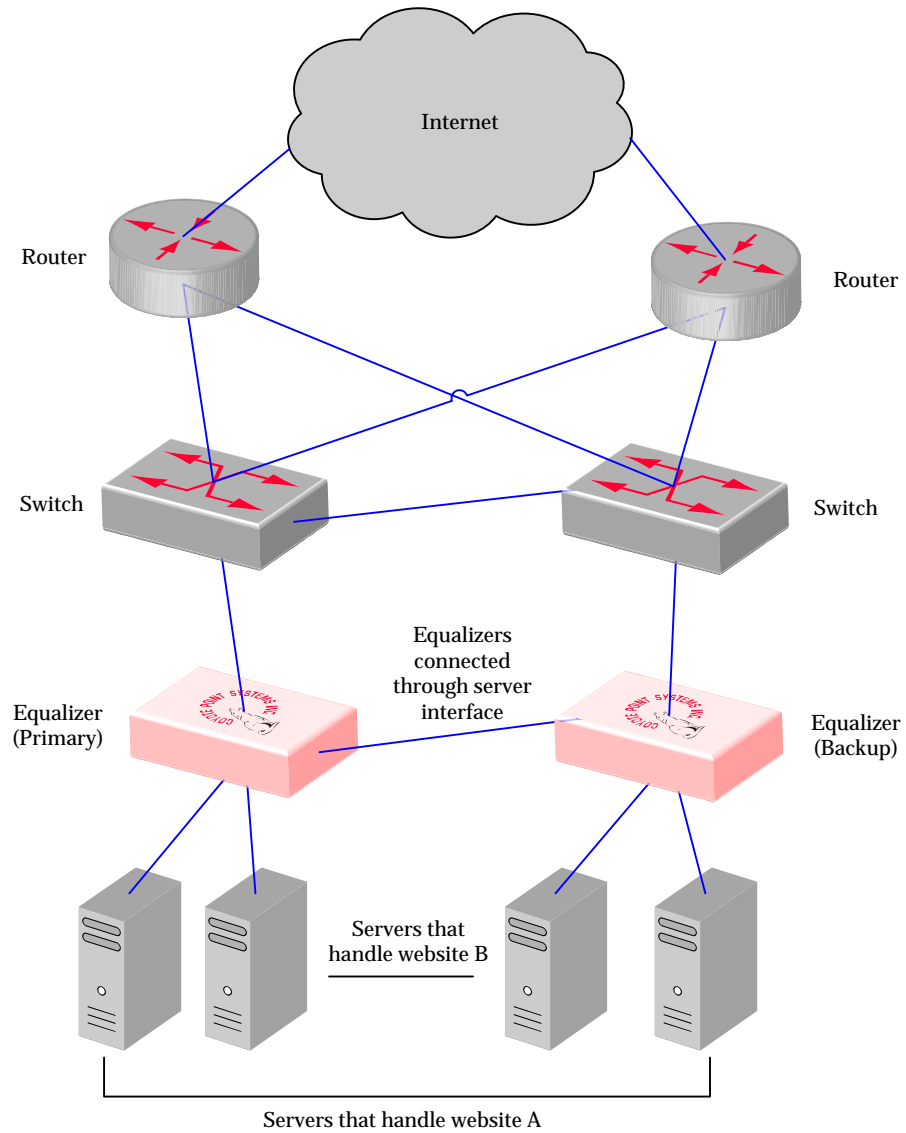


Figure 9 Sample failover configuration

In the sample failover configuration, there is no single point of failure. If a router goes down, the other router takes over or if a link fails, requests are routed through another link.

Figure 10 shows a sample of the cabling of one of the Equalizers shown in Figure 9.




Figure 10 Cabling example from the sample failover configuration

The backup-unit Equalizer monitors all traffic to and from the primary unit; both Equalizers periodically exchange status messages over the local area network. The sibling Equalizers also exchange current configuration information. When you update the configuration on either machine, the configuration on its sibling is automatically updated.

Should either Equalizer fail to respond to a status message probe, the survivor begins a diagnostic cycle and attempts to contact its sibling via the other network ports. If these attempts fail, the sibling is considered to be *down*.

When the backup Equalizer determines that its sibling is down, it initiates a failover process:

1. The backup Equalizer configures the virtual cluster aliases on the external port and sends out “gratuitous ARP” packets that instruct any external-network hosts to replace ARP table entries that point to the physical address of the failed Equalizer with the physical address of the backup unit.
2. The backup Equalizer configures a *failover gateway alias* on the port that is local to the servers.
 - In a non-hot-backup configuration, the servers use the IP address of the cluster or external port as their default gateway.
 - In a hot-backup environment, the gateway address can migrate between the primary and backup unit. This requires an additional address.
3. The Equalizer kernel changes from BACKUP mode to PRIMARY mode. The PRIMARY-mode Equalizer performs gateway routing of packets between its cluster and external ports, address translation, and load balancing.

When a failed unit is brought back online, it begins to exchange status messages with its sibling. Once both Equalizers have synchronized, the newly-started unit assumes the backup role.

Using Reserved IP Addresses

In environments in which conserving IP addresses is important, using reserved IP addresses can minimize the number of “real” IP addresses needed. Equalizer supports placing servers on *reserved*, non-routable networks such as the class A network 10.0.0.0 and the class C network 192.168.2.0.

For example, an ISP hosting several hundred unique web sites replicated on three servers might not want to assign real IP addresses for all of them because each virtual cluster would consume four addresses: three on the back-end servers and one for the virtual cluster. In this case, the ISP might use 10.0.0.0 (the now-defunct Arpanet) as the internal network and assign virtual server addresses out of this network for the servers.

Figure 11 shows a reserved network configuration in detail.

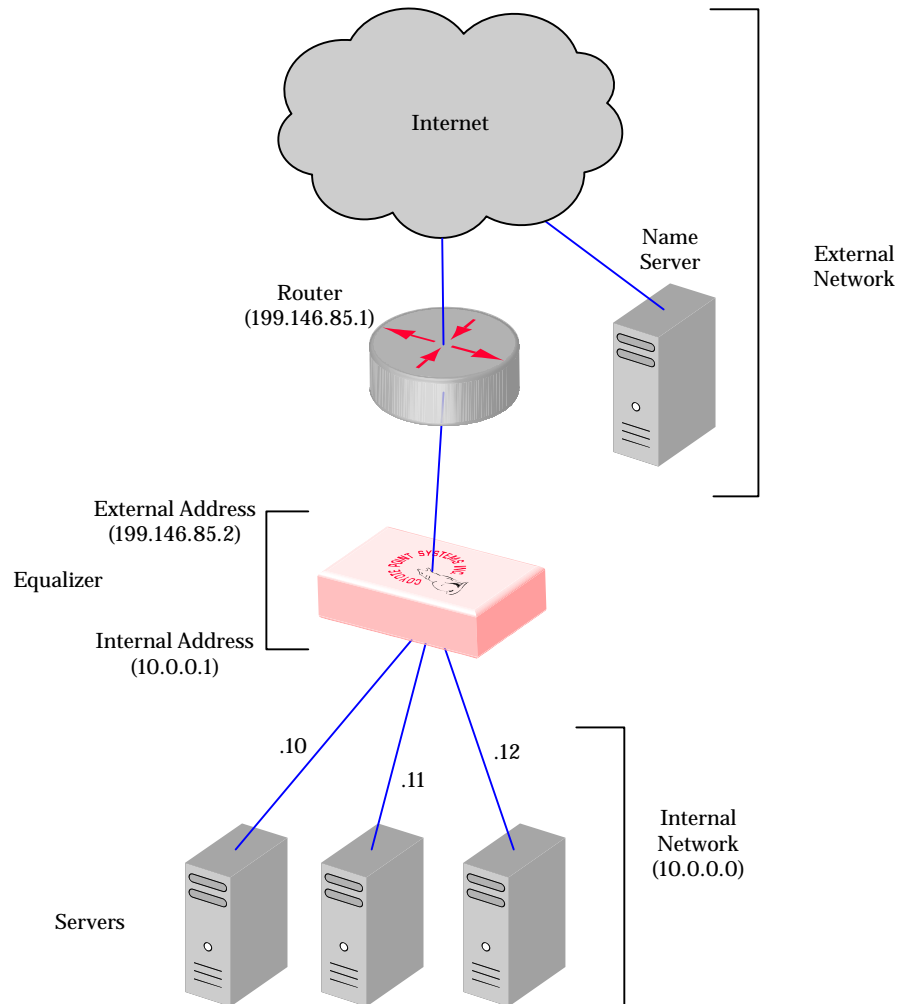


Figure 11 Reserved internal network configuration

If servers placed on a non-routable network need to communicate with hosts on the Internet for any reason (such as performing DNS resolution or sending e-mail), you need to configure Equalizer to perform *outbound NAT*. When you enable outbound NAT, Equalizer translates connections originating from the servers on the reserved network so that external hosts will not see packets originating from non-routable addresses. If you use a failover configuration, you must enable outbound NAT on both Equalizers. For more information, see “Setting Up a Failover Configuration” on page 36.

Note – Due to the additional overhead introduced by enabling outbound NAT, use reserved internal networks with caution.

Equalizer Configuration Worksheets

This section includes two configuration worksheets: one that you can use to prepare to install and configure Equalizer and one that contains questions to answer if you plan to use reserved IP addresses when you set up Equalizer.

Standard Configuration Worksheet

Before you install and configure Equalizer, write down the answers to all the following questions:

1. **What is your physical network layout?**

Will all your servers, Equalizer, and your Internet router reside on a single network? Or will you use a two-network configuration and split your network into multiple subnets? If you use two-network configuration, Equalizer will function as the gateway between them and must be connected to both.

If you don't have a subnet or separate network available to devote to Equalizer's internal network, you can use a single-network topology. For information about using Equalizer with a single network, refer to "Using Equalizer in a Single Network Environment" on page 11.

2. **Which network will be used as the external network?**

Equalizer's external port is connected to this network, which is connected to the Internet.

Example 1: Single Network

For the class C network 199 . 146 . 85 . 0 with a default netmask of 255 . 255 . 255 . 0, the external network would be 199 . 146 . 85 . 0. (See Figure 8 on page 11.)

Example 2: Two Class C Networks

If you use two class C networks, 199 . 146 . 85 . 0 and 199 . 145 . 90 . 0, and choose the first as the external network, the external network would be 199 . 146 . 85 . 0, with a netmask of 255 . 255 . 255 . 0. (See Figure 7 on page 10.)

3. **What is Equalizer's address on the external network?**

You can assign any suitable IP address on your external network as Equalizer's external/administration address. To administer Equalizer, enter this address in your browser's URL field.

Example 1: Single Network

Equalizer Administration Address: 199 . 146 . 85 . 2. (See Figure 8 on page 11.)

Example 2: Two Class C Networks

Equalizer Administration Address: 199 . 145 . 85 . 2. (See Figure 7 on page 10.)

4. What network will be used as the internal network?

This is the network on which the physical servers will reside. If you use separate external and internal networks, the internal network is connected to Equalizer's server port. You should configure routers within your site's network (the external network) to use Equalizer's external port as the gateway to the internal network.

Example 1: Single Network - Switch-based Equalizer (more than two ports)

External port (labeled Ext): Not Used. (See Figure 8 on page 11.)

Example 2: Single Network - Dual-port Equalizer

Server port (labeled Int): Not Used. (See Figure 8 on page 11.)

5. What is Equalizer's address on the internal network?

Typically, assign the lowest numbered address on the internal network as Equalizer's address. The back-end servers will use this address as their default gateway.

Example 1: Single Network

Equalizer Internal Network Address: Not applicable. (See Figure 8 on page 11.)

Example 2: Two Class C Networks

Equalizer Internal Network Address: 199 . 146 . 90 . 1. (See Figure 7 on page 10.)

6. How many physical server machines will you be configuring? What are their IP addresses on the internal network?

If you plan to use IP aliases on the server hosts (virtual hosting), decide the addresses that will be configured on each of the server machines. All server IP addresses and aliases must be unique; you can configure a particular server IP address or alias only one server machine.

7. What virtual cluster addresses will you be configuring?

Address the virtual cluster addresses on the external network.

For example, 199 . 146 . 85 . 4 : HTTP is a virtual cluster on port 80, and 199 . 146 . 85 . 4 : FTP is a virtual cluster on port 21.

8. What is the address of your internet router on the external network?

Equalizer uses this gateway when transmitting packets to hosts that are not on the internal or external networks.

9. What is the IP address of the name server that Equalizer will use?

If you configure a name server, Equalizer displays virtual cluster and server addresses by name rather than by IP address. If no name server is available, set the name server address to 0 . 0 . 0 . 0.

10. Where are your Name Servers?

If you configure Equalizer to use Envoy, determine the DNS servers in your organization that you need to configure to refer fully qualified domain lookups to your Equalizer machine(s).

Special Configuration Worksheet for Using Reserved IP Addresses

Equalizer supports placing servers on reserved, non-routable networks such as the class A network 10.0.0.0 and the class C network 192.168.2.0. In environments in which conservation of IP addresses is important, using reserved IP addresses can minimize the number of “real” IP addresses needed. However, due to the additional overhead introduced by enabling outbound NAT, approach using reserved internal networks with caution. For more information about using reserved IP addresses, see Appendix B.

Before you install and configure Equalizer using reserved IP addresses, write down the answers to both of the following questions:

1. What is the reserved network to be used for the internal network?

Equalizer uses this set of addresses to balance the load across the servers. (Equalizer uses the internal network to forward connections to the HTTP daemons running on the servers.)

Example:

10.0.0.0 (netmask 255.0.0.0) or 192.168.2.0 (netmask 255.255.255.0)

2. What is Equalizer's address on the internal network?

This is the address that the servers will use as their default gateway. This address must be on the reserved network (see the prior step).

Example:

10.0.0.1 or 192.168.2.1



Before You Install Equalizer

The first step in setting up Equalizer is to connect it to the local area network and a power source. Once you have installed Equalizer, you need to configure it as described in Chapter 3, “Configuring Equalizer Hardware”.

Stepping Through the Hardware Installation

To install Equalizer, follow these steps:

1. Carefully remove the Equalizer rack-mount enclosure and cables from the shipping container.
Save the original packaging in case you need to ship the Equalizer for any reason, such as sending it in for warranty service. (The Equalizer chassis does not contain any parts that you can service. If you open the chassis or attempt to make repairs, you may void your warranty.)
2. Place the Equalizer in its intended position in an EIA equipment rack or on a flat surface. To learn about environmental limits and power requirements, refer to Coyote Point’s website for the technical specifications for your Equalizer equipment.
3. Using the supplied serial cable, connect a serial terminal or a workstation running terminal emulator software to the port labeled *Serial* on the front panel of the Equalizer (see hardware appendix).
4. Connect Equalizer to the network with a quality category 5 network cable:
 - a. To use Equalizer as an intermediary between an external and internal network, connect Equalizer to the external network using the RJ-45 network connector marked *Ext* and connect Equalizer to the internal network using one or more of the numbered internal network connectors.
 - b. For a single-network topology with a switch-based Equalizer (more than two ports), connect Equalizer to the external network using one of the numbered RJ-45 network connectors on the front panel of the Equalizer and connect Equalizer to the internal network using one or more of the other numbered network connectors.
 - c. For a single-network topology with a dual-port Equalizer, connect Equalizer using the RJ-45 network connectors labeled *Ext* on the front panel of the Equalizer to a switch connected to both the external network and the internal network.
5. Connect Equalizer to an appropriate power source using the supplied power cord, which plugs into the 3-pin connector on the rear of the Equalizer enclosure. This system uses an auto-sensing power supply that can operate at 50Hz or 60Hz, 110-240 VAC input.
6. Turn on the power using the switch on the rear panel.

Once you have installed and started Equalizer, follow the directions in Chapter 3, “Configuring Equalizer Hardware” to configure the hardware for your network.



Setting Up a Terminal or Terminal Emulator for Equalizer

After the installation of the Equalizer hardware, you need to use a terminal or terminal emulator to complete the hardware configuration. To set up a terminal or terminal emulator for Equalizer, use the following settings:

- 9600 baud
- 8 data bits
- no parity
- one stop bit
- VT100 emulation

If you use the Windows built-in terminal emulator, HyperTerminal, you also need to enable both keyboard application mode and cursor keypad mode.

If your terminal software supports it, set it to ignore hang-ups on the serial line. This allows a single terminal session to continue running even if Equalizer restarts.

Coyote Point recommends using Tera Term (<http://hp.vector.co.jp/authors/VA002416/teraterm.html>) to configure the Equalizer hardware.

Performing the Initial Configuration

This section describes the configuration processes you should perform after installation of Equalizer hardware. The configuration follows these processes:

- Configuring Equalizer for your network
- Configuring the servers on your network to use Equalizer
- Testing the configuration to verify that the system is working properly

As Equalizer boots, the terminal displays a series of device probe and startup messages. Normally, you can ignore these diagnostic messages. However, if you do not configure the terminal emulation software to ignore hang-ups, the terminal session might exit twice during the boot process. If this happens, restart the terminal session.

Use the Equalizer Configuration Utility to specify the following:

- **Hostname:** The DNS hostname that is assigned to Equalizer (optional).
- **Network Interfaces:** The IP addresses of Equalizer on the external and internal networks and the netmasks associated with these networks.

- **Default Router:** The IP address of the router that Equalizer will use to forward outbound packets. The router is on the external network.
- **DNS Server:** The name server Equalizer uses.
- Current date, time, and time zone.
- Passwords for the Equalizer console and administration interface.

Starting to Configure Equalizer

To begin configuration, follow these steps:

1. When the boot process is complete, press **ENTER** on the terminal keyboard to display the login prompt.
2. When the login prompt appears, enter `eqadmin <ENTER>`.
3. When the password prompt appears, enter the password that Coyote Point gave you. Once you enter the password, Equalizer automatically launches the Equalizer Configuration Utility (see Figure 12), which provides a character-based interface for setting and changing Equalizer configuration parameters.
4. If you cannot clearly see the display or you do not see a menu similar to that shown in Figure 12, press `<ESC>` and make sure that your terminal emulator is set for VT100 emulation.

```
+----- Equalizer Configuration Menu -----+
| Equalizer main configuration menu. Select one of the options below using |
| the arrow keys or typing the first letter of the option you wish to invoke. |
| Invoke an option by pressing Enter. |
| Tab to [Exit Install] to exit this utility |
+-----+
| 1 Usage          Quick start - How to use this menu system |
| 2 Doc            Configuring networks, setting parameters, etc |
| 3 Interfaces     Set networking parameters |
| 4 Time Zone      Set the system's time zone. |
| 5 Clock          Set the system's time. |
| 6 Password       Set browser administration tool "touch" password. |
| 7 Console        Set console password. |
| 8 Commit         Commit changes & reboot |
| 9 Shutdown       Shutdown system prior to power-down. (does not commit) |
| 10 Upgrade       Install new software |
+-----+
|                                     [Select]   Exit Configuration |
+-----+
```

Figure 12 Equalizer Configuration Utility: Main Menu

5. To select a menu item within the configuration utility, press one or more arrow keys until you highlight the desired item. If the arrow keys do not operate within your terminal emulator, you can use **CTRL-n** to select the *next* menu item or **CTRL-p** to select the *previous* menu item. Press the **Tab** key to highlight one of the menu actions (such as Select or Cancel) displayed at the bottom of the window. Then press **ENTER** to continue.

Configuring the Network Parameters

To configure the Hostname, Network Interfaces, Default Router, and DNS, use the following steps. Even if you are using your Equalizer in a single network configuration, you need to enter information for both the external and internal (server) interfaces.

1. In the Equalizer Configuration Menu window, select option 3, **Interfaces**, and press **ENTER**. Equalizer displays the Configure Network Interfaces window (see Figure 13).

```

+----- Configure network interfaces -----+
| Configure each of the network interfaces listed below
| Assign an IP address on the external network to the external
| interface and an IP address on the internal network to the
| internal interface. The internal network is the network the
| servers are attached to, the external network is the network which is
| closest to the internet router. Assign the appropriate netmask to each
| interface, as well as a fully qualified hostname. Set the default gateway
| to the IP address of the router on the external network.
+-----+
|                               fxp0  External Ethernet interface
|                               fxp1  Internal Ethernet interface
+-----+
|
|                               [Configure]  Back
+-----+

```

Figure 13 Equalizer Configuration Utility: Sample Interfaces

2. Press one or more arrow keys until you highlight **External Ethernet interface**; then press **ENTER**. The Equalizer Configuration Utility displays the Network Configuration window (see Figure 14).

Regardless of the configuration (single or dual), you need to configure the **Host**, **Domain**, **Gateway**, and **Name Server** fields through the External Interface. The Internal Ethernet interface configuration does not include these fields.

```

+----- Network Configuration -----+
| Host:                               Domain:
+-----+                               +-----+
| Eq_ext.customer.com                 |customer.com
+-----+                               +-----+
| Gateway:                             Name server:
+-----+                               +-----+
| 10.0.0.1                             |
+-----+                               +-----+
|
|   +----- Configuration for Interface fxp0 -----+
|   | IP Address:                               Netmask:
|   | +-----+                               +-----+
|   | |10.0.0.220                               |255.255.255.0
|   | +-----+                               +-----+
|   | Extra options to ifconfig:
|   | +-----+
|   | |
|   | +-----+
+-----+
|
|   +-----+                               +-----+
|   | OK |                               | CANCEL |
+-----+

```

Figure 14 Equalizer Configuration Utility: Network Configuration

3. In the **Host** field, enter the name for the Equalizer on your network. (You can press the Tab key to move among the fields in this screen.)

4. In the **Domain** field, enter the domain name for the Equalizer. (for example, for the fully qualified name, `equalizer.mynet.com`, you would enter “equalizer” in the **Host** field and “mynet.com” in the **Domain** field.
5. In the **Gateway** field, enter the IP address of the router on the external network. This router is the gateway for all the packets Equalizer sends to the outside world through the external network. For example, if your external network router is located at IP address `192.22.33.1`, enter `192.22.33.1` in the Gateway field.
6. In the **Name Server** field, enter the IP address of the domain name server that Equalizer will use. To indicate that no name server is available, enter `0.0.0.0`.
7. If you will be using the external port (that is, using either a dual-network configuration for a switch-based Equalizer or any configuration on a two-port Equalizer) you need to assign an IP address to the external interface. In the IP address and Netmask fields, respectively, specify the IP address and netmask for the external interface. Use the address and netmask from your configuration worksheet. For more information, see “Equalizer Configuration Worksheets” on page 16.

For single network configurations using a switch-based Equalizer, configure an IP address for the external interface of `0.0.0.0` to disable the port.

8. When you’re finished, highlight **OK**. Then press **ENTER**.
Follow the next two steps only if you are using a switch-based Equalizer or a two-port Equalizer in a dual-network mode.
9. To specify the internal interface parameters, select **Internal Ethernet interface**. Then press **ENTER**.
10. Specify the **IP Address** and **Netmask**. For example, if the internal interface will have the address `192.22.34.2`, enter `192.22.34.2` in the **IP Address** field. Enter the **Netmask** specified in the configuration worksheet. If you do not enter an address, the default is `0.0.0.0`, which disables the server ports.
11. Highlight **OK**. Then press **ENTER**.
12. Highlight **Back**. Then press **ENTER** to return to the main configuration menu.

For the new settings to take effect, you must commit these changes and reboot Equalizer.

Committing Changes to the Configuration Parameters

For the changes in the Interfaces panel to take effect, you must commit the changes to the networking parameters and reboot Equalizer. To commit the network configuration changes, follow these steps:

1. In the Equalizer Configuration Menu window, select option 8, **Commit**; then press **ENTER**.
2. When the boot process is complete, check network connectivity on both the internal and external interfaces by pinging the assigned addresses. Ping the external address from a host on the external network, and ping the internal address from a host in the internal network.

Setting the Time Zone

To set the current time zone, follow these steps:

1. In the Equalizer Configuration Menu window, select option 4, **Time Zone**, and press **ENTER**.
2. Use the menus to specify your time zone.
3. Highlight **OK**; then press **ENTER**.

Setting the Date and Time

To set the current date and time, follow these steps:

1. In the Equalizer Configuration Menu window, select option 5, **Time**; then press **ENTER**.
2. Specify the current date and time, based on a 24-hour clock, in the format MM/DD/YY HH:MM.
3. Highlight **OK**; then press **ENTER**.

Changing Equalizer's Console Password

Use the console password to access this configuration utility. Your password can include any combination of printable characters (except spaces). To change Equalizer's console password, follow these steps:

1. In the Equalizer Configuration Menu window, select option 7, **Console**. Then press **ENTER**.
2. Type the new password. When prompted, enter the password again to confirm the change. The new password takes effect immediately.

Changing the Administration Interface Password

The administration interface password is the *edit mode* password for the HTML-based administration interface. Your password can include any combination of printable characters (except spaces) and can be no more than 20 characters in length. To change the administration password, follow these steps:

1. In the Equalizer Configuration Menu window, select option 6, **Password**, and press **ENTER**.
2. Type the new password. When prompted, enter the password again to confirm the change. The new password takes effect immediately.

Upgrading Equalizer Software

Use the Equalizer configuration utility to install the latest Equalizer software from Coyote Point. To install an upgrade:

Note – Before you can upgrade your Equalizer, you must license it. See “Licensing Your Equalizer” on page 35 for more information.

1. In the Equalizer Configuration Menu window, select option 10, **Upgrade**, and press **ENTER**.
2. Highlight **OK**; then press **ENTER**. This starts the Equalizer upgrade script from the Coyote Point FTP site.

firewall protects one or more of your sites, you must configure the firewall to permit Equalizer packets to pass through.

To use geographic load balancing with firewalled networks, you need to configure the firewalls so that the following occurs:

- Equalizer sites communicate with each other on UDP ports 5300 and 5301. The firewall must allow traffic on these ports to pass between Envoy sites.
- Equalizer sites and clients can exchange packets on UDP port 53. The firewall must allow traffic on this port to flow freely between an Equalizer server and any Internet clients so that clients trying to resolve hostnames via the Equalizer DNS server can exchange packets with Equalizer sites.

Equalizer sites can send ICMP echo request packets through the firewall and receive ICMP echo response packets from clients outside the firewall. (When a client attempts a DNS resolution, Equalizer sites send an ICMP echo request (ping) packet to the client; the client might respond with an ICMP echo response packet.)

Configuring Servers

To use Equalizer, you must configure your servers so that Equalizer gateways the packets the servers send to their clients. If you do not adjust the routing on your servers, a client may not receive a response when it attempts to contact a virtual cluster. Then, the connection will time out.

When you configure the servers, the *default* route gateway depends on your Equalizer configuration:

- If you use a two-network configuration, the gateway for the default route should be Equalizer's internal address regardless of the Equalizer model.
- If you use a single-network configuration on switch-based Equalizers, the gateway for the default route should be Equalizer's internal address.
- If you use a single-network configuration on dual-port Equalizers, the gateway for the default route should be Equalizer's external address.
- If you use a failover configuration, set the default route to the failover alias. For more information, see "Setting Up a Failover Configuration" on page 36.

To verify that you have configured a server's routing correctly, run `tracert` on the server with a destination address outside the internal network to ensure that Equalizer gets used as a gateway.

The way that you configure a server depends on the server's operating system. Configure each server from the system console, not through a telnet session.

Configuring a Second Equalizer As a Backup

You can configure a second Equalizer as a hot backup (or hot spare) so that if the Equalizer that currently handles requests (the *primary unit*) fails, the *backup unit* automatically takes over.

Both the primary and backup units are configured to default to either primary or backup role. When a failed unit comes back online, it assumes the backup role, even if it is designated the default primary. To set up a failover configuration, use the Equalizer Administration Interface. For more information, see "Setting Up a Failover Configuration" on page 36.

Testing Your Configuration

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

To perform these tests, you need the following:

- A test machine on the internal network—the same physical network as the servers (for example, one of the server machines).
- If you have a two-network configuration, a test machine on the external network.
- A client machine somewhere on the Internet, to simulate a “real-world” client. This machine should be set up so that the only way it can communicate with your servers or Equalizer is through your Internet router.

Then follow these steps:

1. From the internal-network test machine, ping the physical IP address of each server. You should be able to successfully ping all of the servers from the test machine.
2. From the internal-network test machine, ping the server aliases on each of the servers. You should be able to successfully ping all of the servers from the test machine using their aliases.
3. From the internal test machine and each of the servers, ping the Equalizer address that you use as the default gateway on your servers. (If you use a two-network topology, this will be Equalizer’s internal address or failover alias.)
4. From the internal-network test machine, telnet to the server aliases on service ports of running daemons. You should be able to telnet successfully to the server aliases.
5. If you use a two-network configuration: From the external-network test machine, ping a physical server IP address using `ping -R` to trace the route of the ping. The Equalizer IP address should appear in the list of interfaces that the ping packet traverses. You can also use the `tracert` tool to perform this test.
6. If you use a two-network configuration: After you have configured a virtual cluster and added servers, telnet to each of the virtual clusters configured on the Equalizer from an external-network test machine. (See Chapter 5, “Configuring Equalizer Operation” for information about how to add clusters and servers.)

When you telnet to a virtual cluster from the external test machine, Equalizer should connect you to one of the servers configured in the cluster. Repeatedly connect to the same virtual cluster to make sure that Equalizer routes the connections to different servers in the cluster. (Equalizer does not necessarily select the servers in a round-robin fashion. Equalizer uses an adaptive algorithm to determine the server that gets the next connection.)

Note – You also can use a client tool such as a Web browser to perform this test.

7. From the Internet-client test machine, connect to each virtual cluster.

For help in resolving configuration problems, see Appendix G, “Troubleshooting.”



Introducing the Equalizer Administration Interface

You use Equalizer’s HTML-based administration interface for routine monitoring and administrative tasks. Access the administration interface from a Javascript-enabled web browser to perform the following actions:

- Monitor the status of Equalizer and the configured clusters and servers
- View cluster and server performance statistics graphically
- Add virtual clusters
- Modify cluster parameters
- Delete clusters
- Add servers to a cluster
- Adjust server static weights
- Delete servers
- Shut down a server gracefully
- Shut down Equalizer

Accessing the Equalizer Administration Interface

You must access the Equalizer Administration Interface through a Javascript-enabled browser.

The Equalizer Administration Interface supports the following two user modes:



View, which enables you to *view*, but not edit, Equalizer configuration and status information. For more information about view mode, see Chapter 6, “Monitoring Equalizer Operation” which starts on page 51.



Edit, which enables you to *view* all the Equalizer configuration and status information and, most important, *edit* the configuration. For more information about edit mode, read through Chapter 5, “Configuring Equalizer Operation” which starts on page 35.

Logging In

To access the administration interface and log into Equalizer, follow these steps:

1. Launch a Javascript-enabled web browser.
2. From the browser, load the URL that corresponds to Equalizer's external address.

For example, if the external address is 199 . 146 . 85 . 2, open the Equalizer Administration Interface by typing `http://199.146.85.2/` in the appropriate location in the browser. If you are using a redundant pair of Equalizers, use the failover alias to ensure that the browser connects to the Equalizer that has the primary role.

Equalizer displays the login screen (see Figure 16):

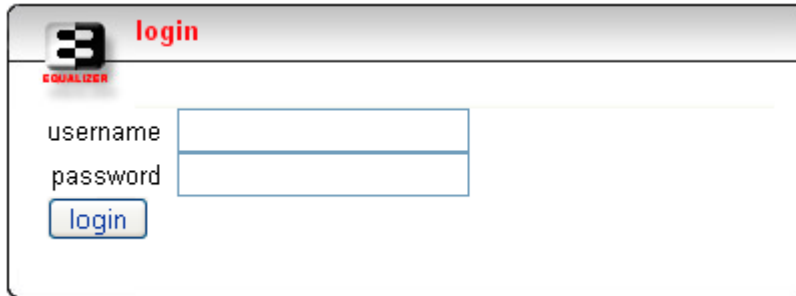


Figure 16 Log In screen

3. Enter the appropriate user name and password; then click the **login** button.

To obtain the initial user name and password combinations for view and edit access, see the password sheet that Coyote Point provided separately.

Note – If you have lost or forgotten the edit mode password, you can set it through the Equalizer Configuration Utility. For more information refer to “Changing the Administration Interface Password” which starts on page 25.

Navigating Through the Interface

The Equalizer Administration Interface (see Figure 17) provides two navigation mechanisms: links and menus.

You can access status information and current parameters of any of the items in the hierarchical list in the left frame by clicking the name of the item you want to view. The hierarchical list contains all

the currently configured clusters, servers, geographic clusters, and sites. Equalizer displays the status information and current parameters in the right frame.

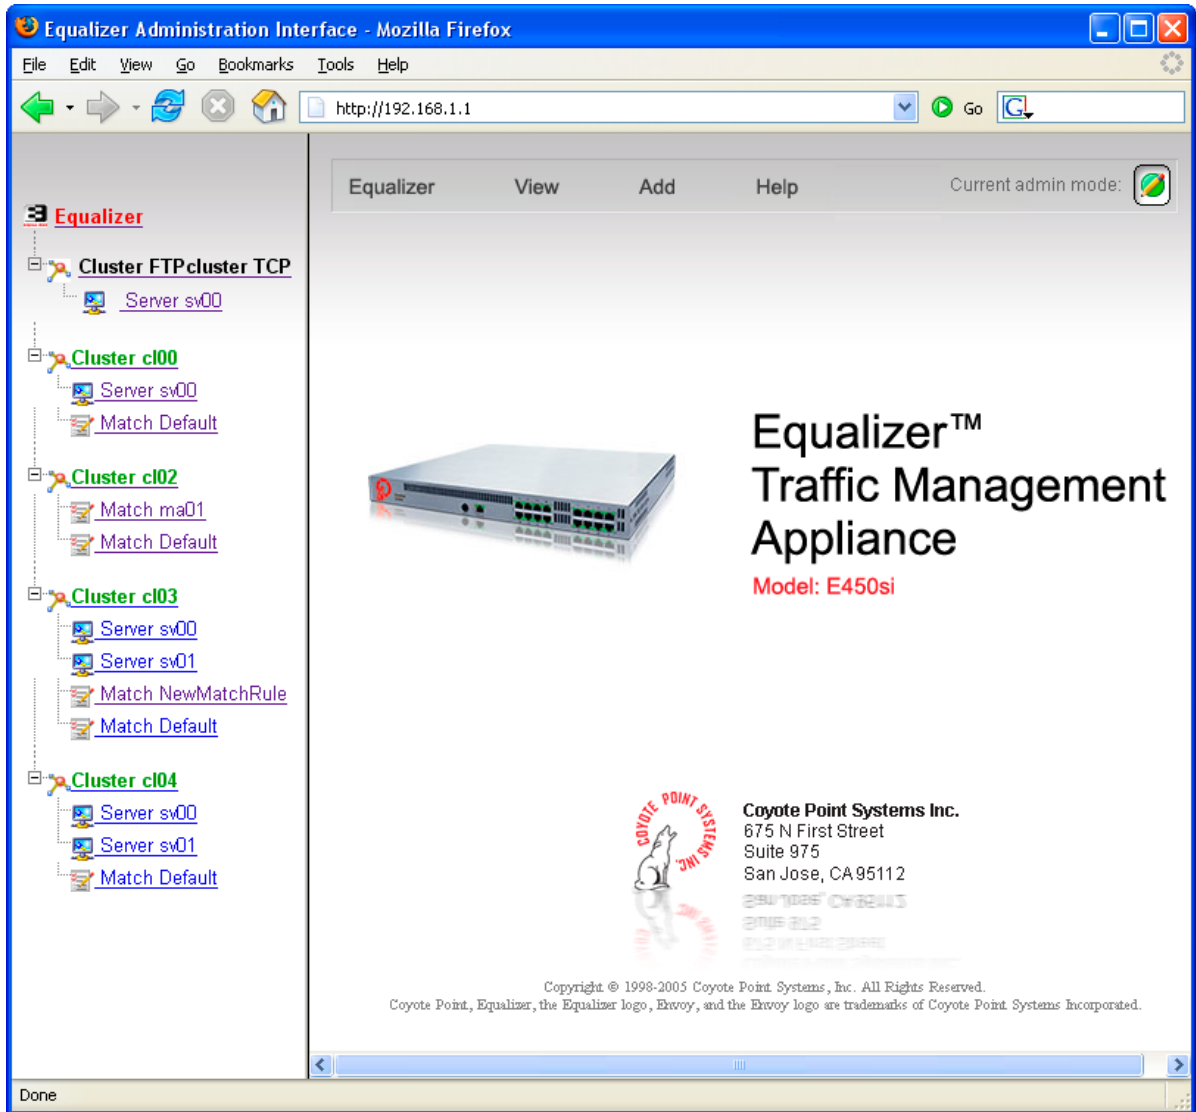


Figure 17 Equalizer’s Administration Interface

Using the Main Menu Bar

Use the menus in the main menu bar (see Figure 19) in the top frame and the local menus on the parameters pages to access Equalizer’s reporting options, modify the configuration, or view help information.




Figure 18 Main menu bar

- **Equalizer**, which contains the following commands for shutting down the Equalizer, logging out of the Administration Interface, and modifying Equalizer global parameters:

Shut Down Equalizer, from which you can perform a clean shutdown of the Equalizer system so you can safely turn off the power. Note that this option works only when you are logged in under edit mode. If you try to do this while you are logged in under view mode, Equalizer displays an error message.

Reboot Equalizer, from which you can cause Equalizer to reboot itself. If you try to do this while you are logged in under view mode, Equalizer displays an error message.

Log Out, from which you can log out and end the administration session.

Configure, which displays a submenu with six options for modifying (under Edit mode) the Equalizer global parameters: Change Passwords, Events, Failover, Backup/Restore Configuration, Manage Licenses, and System Parameters.

- **View**, which provides access to the following global status information:
 - Equalizer Status*, to display the Equalizer software and hardware information, basic configuration, and recent statistics.
 - Cluster Summary*, to display summary information for all the configured clusters.
 - Event Log*, to display the Equalizer event log. When you have finished viewing the event log, moving to another location automatically closes the event log.
- **Add**, which provides the following commands for adding clusters under Edit mode:
 - Virtual Cluster*, to add a new virtual cluster.
 - Geographic Cluster*, to add a new geographic cluster to a site. This command is only available if Envoy is installed.
- **Help**, which provides access to the following information about using Equalizer:
 - View Guide*, which displays the PDF file that contains this Installation and Administration guide.
 - Context Help*, which displays the relevant section in the PDF file corresponding to the current activity in the right frame.
 - About Equalizer*, which displays version and copyright information for Equalizer.

The icon displayed in the top right corner of the administration interface indicates the current user mode: View or Edit. When you are logged in under view mode, the configuration functions, such as adding a server or modifying a cluster's parameters are not available.

Accessing Local Menus

You can access local menus (see Figure 19) from the parameters screens that appear when you click an item in the left frame. Typically, you can find local menus in the upper right corner of the page. To activate a local menu, roll over it with your mouse. With local menus, you can view and change information about the currently-viewed item. For example, when you are in edit mode, the local menu in the Server Parameters page enables you to change the server's parameters, plot the server's history, or even delete the server.

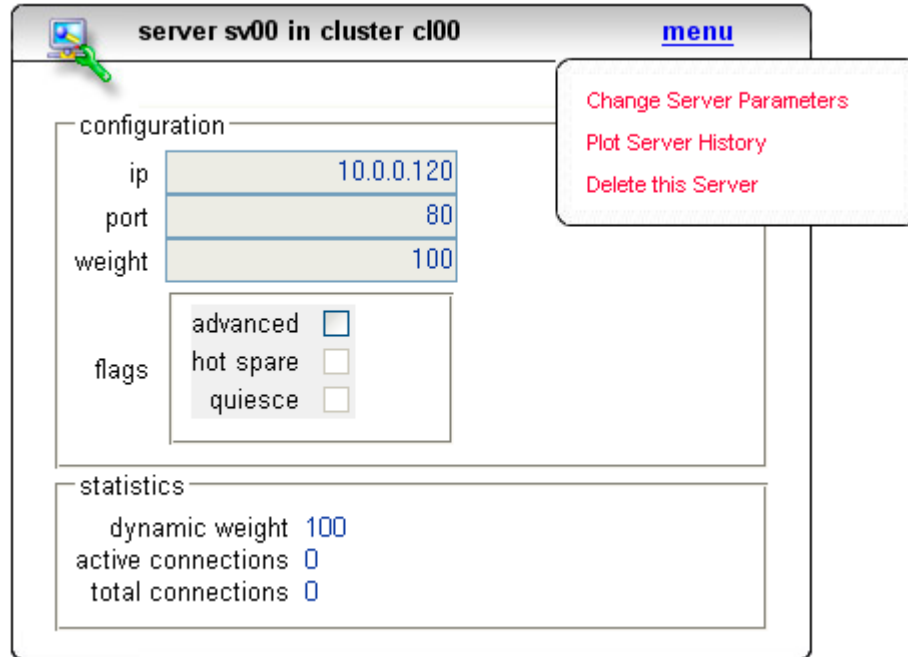



Figure 19 Local menu from the Server Parameters page

5 Configuring Equalizer Operation



You can modify Equalizer's configuration through the Equalizer Administration Interface and perform the following actions:

- License your Equalizer
- Set up a failover configuration with two Equalizers
- Enable outbound network address translation for reserved networks
- Enable passive FTP connections
- Configure stale connection handling
- Enable sticky network aggregation
- Configure custom event handling
- Set the administration passwords

Licensing Your Equalizer

You must register and license your Equalizer before performing any other configuration using the Equalizer Administration Interface. You need a license to enable your Equalizer's features.

1. Go to <http://www.coyotepoint.com/register.htm>.

To complete the registration, you'll need the serial number located on the Equalizer unit as well as the System ID. You can retrieve the System ID using the Administration interface. From the Equalizer menu, choose **Configure > Manage Licenses** to see the System ID. Simply copy and paste this ID into the registration form.

2. If you are not already, log into the Administration interface in edit mode.
3. Choose **Configure > Manage Licenses** from the Equalizer menu.

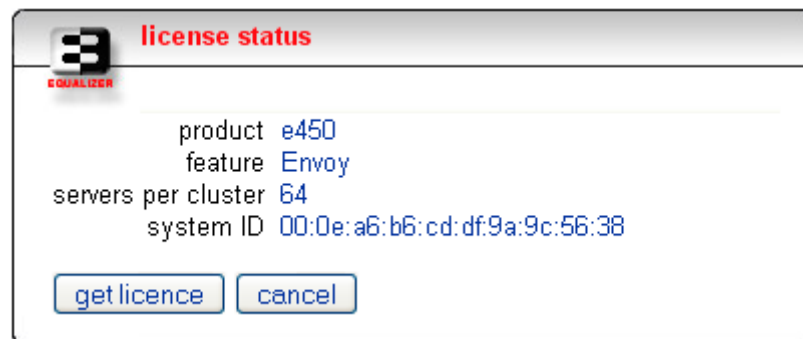


Figure 20 License status screen

4. Click the **get license** button to send a request for a license.

If Coyote Point's license server validates the request, it returns a license for the Equalizer to use. When you get a license, the Unlicensed Error warning will disappear from the left frame and your Equalizer's features are enabled.

Setting Up a Failover Configuration

You can use a second Equalizer as a hot backup. Then, the backup will automatically take over if the Equalizer that is currently handling requests fails. To use a second Equalizer as a hot backup, you need to install the backup Equalizer so its network interfaces correspond to the network interfaces of its sibling:

- You must plug the external interface of the backup unit into the same hub or switch into which the external interface of the primary unit is plugged.
- You must plug the server (or internal) interface of the backup unit into the same hub or switch into which the server interface of the primary unit is plugged.

Note – You should never create a loop between the external and internal interfaces.

- For failover configuration between two switch models, connect a cable from one Equalizer's switch interface to the others.

You must designate one of the Equalizers as the *preferred primary*. When you boot both Equalizers at the same time, the default primary Equalizer is activated. If the primary Equalizer fails, the backup takes over. When you bring the failed unit back online, it assumes the backup role until another failure occurs or you reboot its sibling.

When an Equalizer is brought online, the Equalizer checks to make sure that the network interfaces are in a valid state (that is, link active). If the appropriate interfaces are not valid, the Equalizer sits in a loop waiting for this situation to resolve itself (and sends comments to the console). When the appropriate interfaces are valid, the Equalizer tries to make contact with its sibling. If they establish contact, a negotiation ensues in which one system becomes the primary unit and the other becomes the backup unit (the default primary unit does not necessarily become the primary unit). Generally, the first system to start running the failover process becomes the primary unit.

If, at any point, either Equalizer loses contact with its sibling, that Equalizer attempts to resolve the issue by testing its own interfaces. If all interfaces test OK, the backup Equalizer tries to contact its sibling. If it fails three times to contact its sibling, the backup starts the process to become the primary. The backup checks that no other system has configured the gateway IP address or virtual cluster addresses. When it resolves these tests, the Equalizer assumes those IP addresses and starts handling traffic.

A partition occurs when both systems are unable to communicate with each other and both Equalizers enter primary mode. When this partition is healed and both units are running in primary mode, the two systems resolve this dispute by choosing one system to reboot itself. Generally, this

means that the system that is configured as the default backup will reboot; upon coming back up, it will enter backup mode.

Note – Any switch, such as one from Cisco or Dell, that comes with Spanning Tree enabled by default can cause a communication problem in a failover configuration when one or both of the Equalizers are dual-port models. This problem occurs at bootup because the switch disables its ports for roughly 30 seconds to listen to BPDU (bridge protocol data unit) traffic. The 30-second pause causes both Equalizers to attempt to become the primary unit; the default backup continually reboots.

To repair this condition, either disable Spanning Tree or enable PortFast for the ports connected to the Equalizers. This enables the ports to act as normal hubs and accept all traffic immediately.

To set up a failover configuration (see Figure 21), perform the following procedure:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Failover** from the Equalizer menu in the main menu bar. The failover configuration screen appears in the right frame.

failover configuration

configure failover aliases

internal address

internal netmask

single network mode

failover timing

receive timeout

connection timeout

probe interval

failover peers

Create entries for each system (peer) in a failover group.

peer

peer name

internal address

flags preferred primary

defaults commit delete cancel

Figure 21 Failover configuration screen

3. Specify the **address** and **netmask** for the failover aliases. The alias address is a is a unique IP address on the internal network (for a dual-network configuration) or external network (for a

single-network configuration). The Equalizer that is running in primary mode assumes this address; the servers should use this address as their default gateway.

4. To modify the Equalizer you connect to when you logged in, select create new.
5. Specify the requested IP address(es) for the peer.
6. Check **preferred primary** if the peer is the preferred primary.
7. Configure a second peer.
8. Click the **commit** button to save the parameters
9. When you finish the failover configuration, you must reboot Equalizer and its peer. Select Reboot from the Equalizer menu. Rebooting the default primary Equalizer first ensures that it assumes the primary role. Reboot the second unit immediately after rebooting the primary.

As the Equalizers reboot, observe the terminal connection. The console messages should indicate that each Equalizer has successfully assumed the primary or backup role.

You should accept the default failover timing parameters. These parameters affect how the siblings try to test each other. The **receive timeout** is the time in seconds that Equalizer allows to receive a response from its sibling before it timeouts. The **connection timeout** is the time allowed to establish a TCP connection with its sibling. When either of these timeouts occur, that counts as one of the three failed attempts that occur before the backup becomes the primary. The last failover timing value is **probe interval**. This is the number of seconds that each sibling tries to exchange status information. Normally the default values are the best to use, however, if you notice in the log files contain too many false positives (messages that Equalizer has regained contact with its sibling) you may want to increase the values.

Enabling Outbound NAT

If you use a reserved network configuration and the servers on the non-routable network must be able to communicate with hosts on the Internet, you must configure Equalizer to perform outbound network address translation (NAT). When outbound NAT is enabled, Equalizer translates connections originating from the servers on the reserved network so that external hosts won't see packets originating from non-routable addresses.

Note – If you use outbound NAT in a failover configuration, you should enable outbound NAT on both units in case a failover actually occurs.

To enable Equalizer to perform outbound NAT, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select Equalizer in the left frame. Equalizer displays the Equalizer status screen in the right frame.

3. Select **Change Equalizer Parameters** from the local menu. Equalizer displays the **modify system parameters** screen in the right frame.

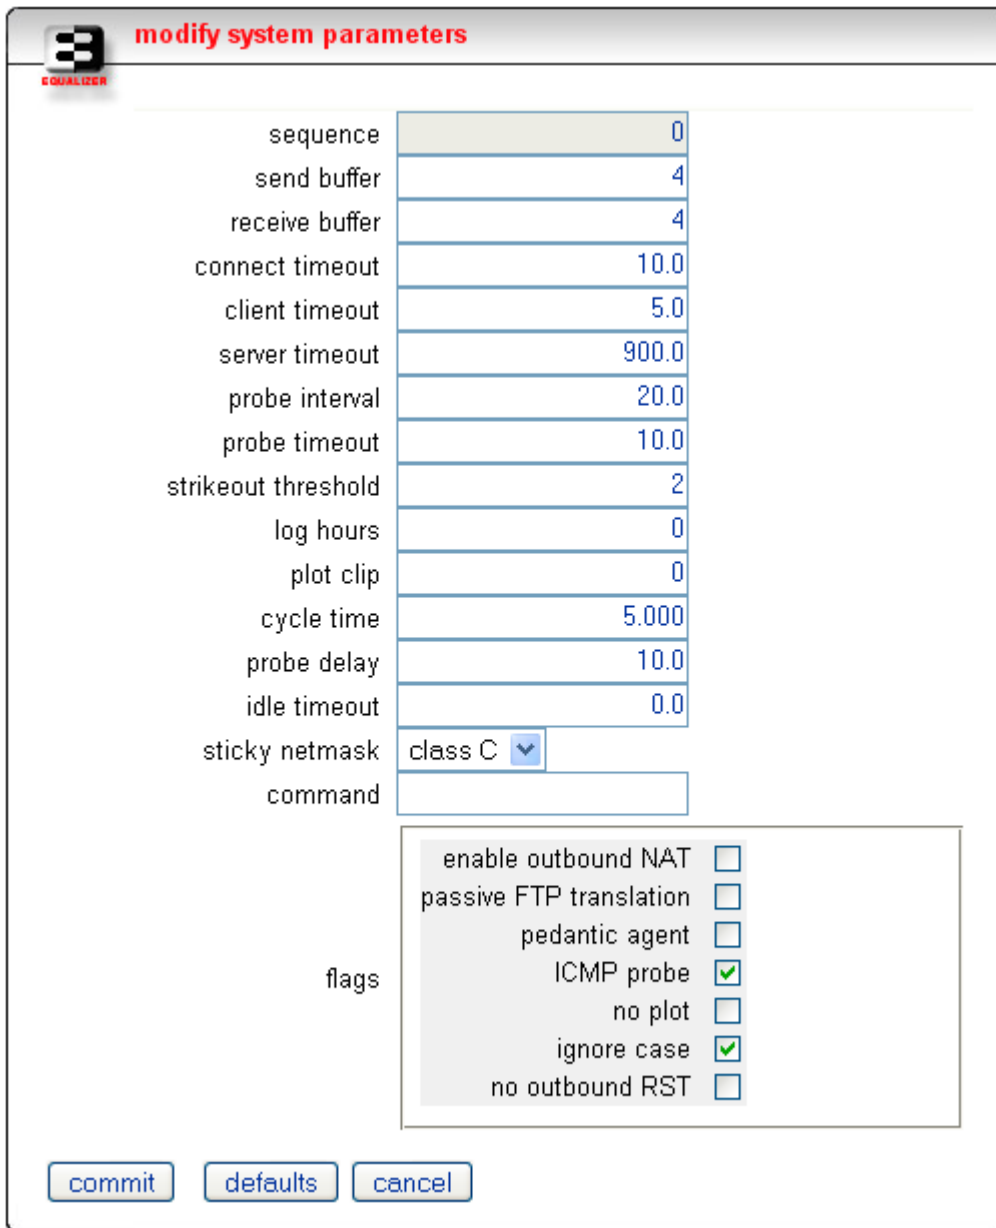


Figure 22 Enabling outbound NAT

4. Check the **enable outbound NAT** checkbox (it's in the flags section of the screen).
5. Click the **commit** button.

Enabling Passive FTP Connections

If your servers are on a network the outside world cannot reach, consider enabling Equalizer's passive FTP translation option. Then, Equalizer rewrites outgoing FTP PASV control messages from the servers so they contain the IP address of the virtual cluster rather than that of the server.

To enable passive FTP translation, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select Equalizer in the left frame. Equalizer displays the Equalizer status screen in the right frame.
3. Select **Change Equalizer Parameters** from the local menu. Equalizer displays the **modify system parameters** screen in the right frame
4. Check the **passive FTP translation** checkbox.
5. Click the **commit** button.

Managing Stale Connections

The stale connection timeout is the length of time that a partially open or closed connection is maintained. If a client fails to complete the TCP connection termination handshake sequence or sends a SYN packet but does not respond to the server's SYN/ACK, Equalizer marks the connection as incomplete. Equalizer reclaims connections in the incomplete state when the stale connection timeout expires. When Equalizer reclaims a connection, Equalizer sends an RST (reset) command to the server, enabling the server to free any resources associated with the connection. Stale connections apply to Layer 4 (L4) only.

If you change the stale timeout setting while partially established connections are currently in the queue, those connections will be affected by the new setting.

Note – Reducing the stale connection timeout can be an effective way to counter the effects of SYN flood attacks on server resources. A stale connection timeout of 10 seconds would be an appropriate value for a site under SYN flood attack.

To set the stale connection timeout, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select Equalizer in the left frame. Equalizer displays the Equalizer status screen in the right frame.
3. Select **Change Equalizer Parameters** from the local menu. Equalizer displays the **modify system parameters** screen in the right frame
4. Enter a value, in seconds, for **stale timeout**.
5. Click the **commit** button.

Enabling Sticky Network Aggregation

Sticky network aggregation enables Equalizer to correctly handle sticky connections from ISPs that use multiple proxy servers to direct user connections. When you enable sticky network aggregation, all the connections coming from a particular network are directed to the same server. (Typically, all the servers in a proxy farm are on the same network.)

When you enable sticky network aggregation, Equalizer routes all the connections from a particular network to the same server. The netmask value indicates which portion of the address Equalizer should use to identify particular networks. The mask corresponds to the number of bits in the network portion of the address:

- 8 bits corresponds to a Class A network
- 16 bits corresponds to a Class B network
- 24 bits corresponds to a Class C network

In previous versions of Equalizer, enabling sticky network aggregation was the equivalent of setting the sticky network aggregation mask to 24 bits (that is, Equalizer routed all connections from the same class C network to the same server).

Sticky network aggregation is applicable only for Layer 4 load balancing of generic TCP and UDP clusters.

Note – A potential drawback of using sticky network aggregation is that all users connecting through a particular proxy farm might be directed to the same server. In practice, this has not been a problem. Equalizer's load-balancing algorithms direct other visitors to different servers to keep the load balanced.

To enable sticky network aggregation, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select Equalizer in the left frame. Equalizer displays the Equalizer status screen in the right frame.

3. Select **Change Equalizer Parameters** from the local menu. Equalizer displays the **modify system parameters** screen in the right frame.

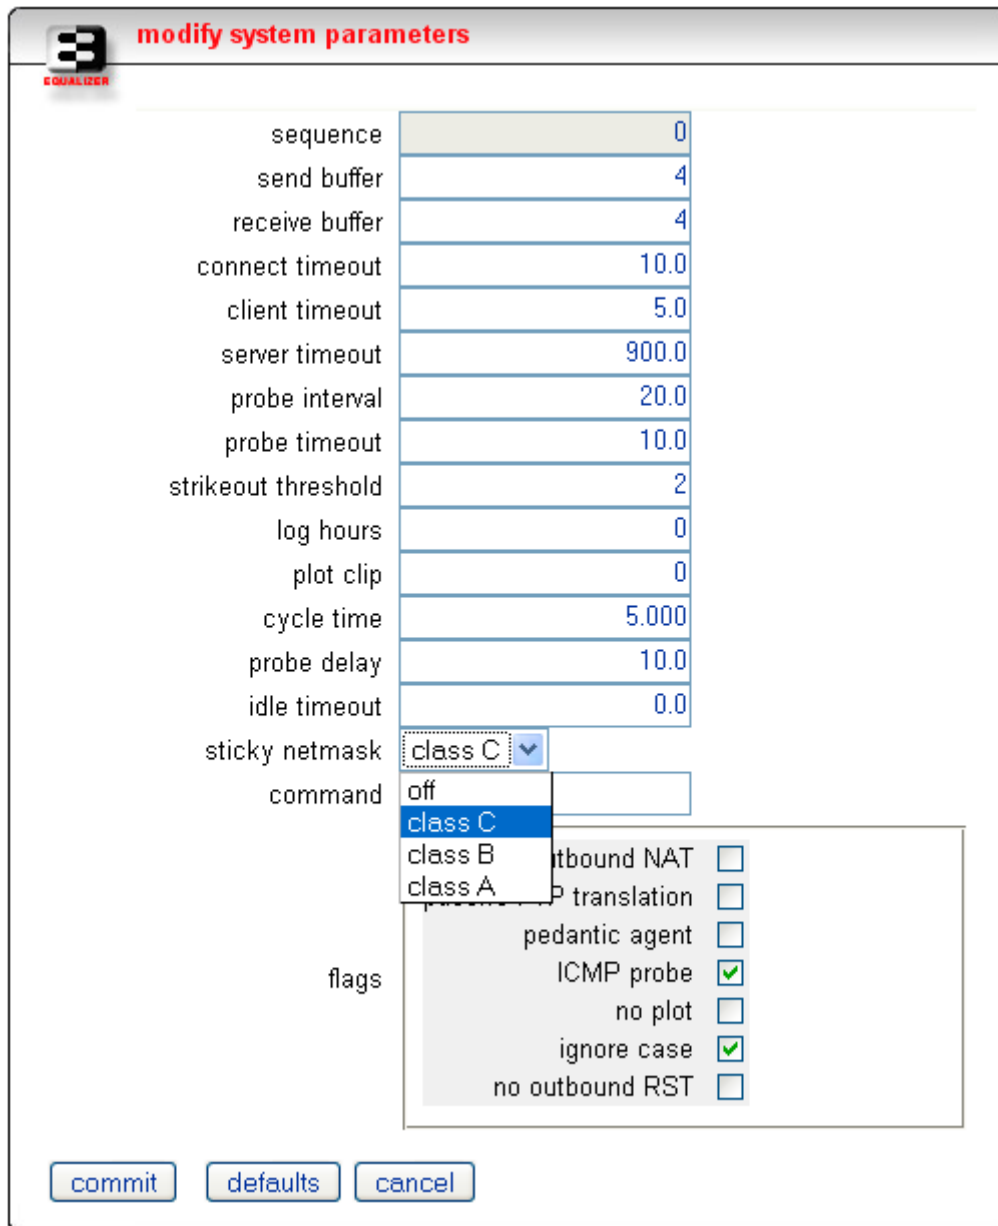


Figure 23 Enabling sticky network aggregation

4. Enable sticky network aggregation by selecting a **sticky netmask** from the pull-down menu.
5. Click the **commit** button.

Note – You must configure the sticky network aggregation mask identically for each Equalizer in a failover pair.

Configuring Custom Event Handling

You can configure Equalizer to perform certain actions when a server fails or other critical events occur. This is done through the Change Equalizer Parameters screen. You can forward Equalizer log information to another machine or specify a command to run when a particular event occurs.

Forwarding Equalizer Log Information

You can forward Equalizer's internal log information to another machine that is running a syslog daemon.

To specify a syslog host to which you will forward the log (see Figure 24), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Events** from the Equalizer menu in the main menu bar. The **event configuration** screen appears in the right frame.

Figure 24 Specifying a syslog host

3. Check the **use remote syslog** checkbox.
4. In the **syslog host** field, enter the hostname (not the IP address) of the machine to which you want to forward syslog messages.
5. Click the **commit** button.

Specifying a Command to Run When a Particular Event Occurs

You can configure Equalizer to run a command that you specify (such as sending an e-mail or running a custom shell script) whenever server events occur. The following events trigger the specified command:

- Failure of a server
- Restoration of a failed server
- Failure of a server agent
- Restoration of a server agent

- Failover in a high-availability Equalizer pair

For example, to send e-mail to `admin@yourdomain.com` whenever Equalizer detects a server failure, you could enter the following command:

```
/usr/local/sbin/mini_sendmail -f equalizer -s
your_local_SMTP_server_IP
```

Equalizer sends a string that describes the server event to the program specified as standard input.

For more information about the `mini_sendmail` command, which is unique to Equalizer, refer to Appendix F, “Mini SendMail” on page 117.

To specify a command to run, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Events** from the Equalizer menu in the main menu bar. The **event configuration** screen appears in the right frame.
3. In the **command to run on server event** field (see Figure 25), enter the command that you want to run Equalizer detects a server event.

Figure 25 Specifying an event-triggered command

4. Click the **commit** button.

Note – Any program that is specified in the command and that is to run for a server event must complete its work and terminate within one or two seconds to avoid interrupting Equalizer’s server failure detection facility.

Changing Other System Parameters

The **modify system parameters** screen displays information that affects Equalizer’s operation. Only modifiable fields not described in previous sections are described here.

- **send buffer** applies to L7 clusters and is the amount of memory in kilobytes reserved by each L7 proxy process to store outgoing data before it is placed on the network interface.

- **receive buffer** applies to L7 clusters and is the amount of memory in kilobytes reserved by each L7 proxy process to store data that has been received on an interface before it is processed by an L7 proxy process.
- **connect timeout** applies to L7 clusters and is the time in seconds that Equalizer waits for a server to respond to a connection request.
- **client timeout** applies to L7 clusters and is the time in seconds that Equalizer waits before closing an idle client connection.
- **server timeout** applies to L7 clusters and is the time in seconds that Equalizer waits before closing an idle server connection.
- **probe interval** is the target time in seconds for server health check probes. This value is solely a target, the monitoring process adjusts itself based on load.
- **strikeout threshold** is the number of failures to respond to a probe (strikes) before a server is declared down.
- **log hours** is the target number of hours of plot log data to retain. A zero in this field allots the numbers of hours based on the available memory.
- **plot clip** applies a threshold to limit the effect of spikes in plot data.
- **cycle time** is time in seconds for the master daemon to make one pass through all of the clusters. This value should not be modified.
- **probe interval** is the time in second between successive probes of servers. You can override this value for each cluster.
- **idle timeout** applies to L4 clusters and is the time in seconds before reclaiming idle Layer 4 connection records.
- **command** is string Equalizer executes when an event occurs. See “Configuring Custom Event Handling” on page 44
- **pedantic agent** applies only when clusters use server agents. When you check this box, Equalizer will treat a server as down when it can probe a server but receives no response from the server’s agent.
- **ICMP probe** makes Equalizer probe servers using a mix of L4, L7 and ICMP echo probes.
- **no plot** causes Equalizer to not record plotting data.
- **ignore case** applies to L7 and is the global setting to ignore case in match expressions. You can override this value per cluster and per match rule.
- **no outbound RST** applies to L4 and causes the Equalizer to not forward untranslated TCP RESET packets.

To modify a system parameter, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select Equalizer in the left frame. Equalizer displays the Equalizer status screen in the right frame.
3. Select **Change Equalizer Parameters** from the local menu. Equalizer displays the **modify system parameters** screen in the right frame
4. Change the appropriate field.
5. Click the **commit** button.

Changing the Administration Passwords

An administrator logged in under Edit mode can change both the View password and Edit password. If you are logged in under Edit mode, you can change the View password without specifying the current password.

To change the view or edit a password (see Figure 26), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Change Passwords** from the Equalizer menu in the main menu bar. The **change passwords** screen appears in the right frame.

Figure 26 Change passwords screen

3. Select the password to be changed: **View Password** or **Edit Password**.
4. For the Edit password only, enter the current password in the **current password** field.
5. Enter the new password in the **new password** field and then confirm it by entering it again in the **confirm password** field.
6. Click the **commit** button.

Note – If you have lost or forgotten the edit mode password, you can set it through the Equalizer Configuration Utility. For more information, refer to “Changing the Administration Interface Password” on page 25.

Saving or Restoring Your Configuration

Equalizer enables you to save or back up a configuration or restore a saved configuration.

Note – Equalizer passwords are not saved or restored, but IP configuration, clusters, and failover information are saved.

Saving Your Configuration

Use the Backup/Restore Configuration command to save your Equalizer configuration to a file or to load a saved configuration.

When you save your configuration, Equalizer wraps up the following four files in a bin file:

- `/etc/eq/conf`, which are cluster/server configurations that appear in the left pane of the interface.
- `/etc/eq.static`, which is the failover configuration file, which appears on the configure failover page of the interface.
- `/etc/rc.conf`, which is the IP information for the unit.
- `/etc/geo.cf`, which is the Envoy configuration, geographic cluster, and site information from the left pane of the interface.

Backing Up Your Configuration

To back up your current configuration to a file, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Backup/Restore Configuration** from the Equalizer menu in the main menu bar. The **backup/restore** screen (see Figure 27) appears in the right frame.




Figure 27 Backing up your Equalizer configuration

3. Click the **backup** button.
4. When prompted, specify the location where you want to save the configuration file; then click **OK**.

Restoring a Saved Configuration

To restore a saved configuration, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Configure > Backup/Restore Configuration** from the Equalizer menu in the main menu bar. The **backup/restore** screen appears in the right frame.

3. Click the **restore** button.

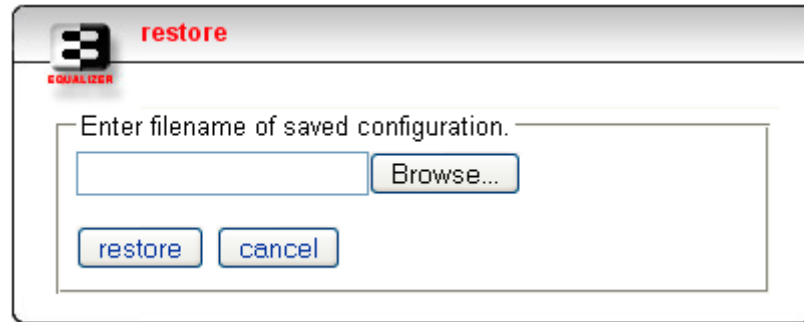


Figure 28 Restoring a saved configuration

4. Click **Browse...** to locate and select the configuration file that you want to use to restore the Equalizer configuration.
5. Click **restore** to upload the configuration file. Equalizer automatically reboots to update the configuration.

Note – Be very careful when restoring configurations. The saved IP information could cause conflicts on the network if the restored file comes from another Equalizer (for example, its backup).

Shutting Down Equalizer

Before turning off Equalizer or disconnecting the power, you should perform a clean shutdown.

To shut down Equalizer cleanly, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Shut Down Equalizer** from the Equalizer menu in the main menu bar. A confirmation dialog box appears (see Figure 29).




Figure 29 The Shutdown confirmation dialog box


3. In the confirmation dialog box, click **OK** to confirm that you really want to shut down Equalizer (or click **Cancel** to abort the shutdown request). If you click **OK**, Equalizer immediately initiates the shutdown cycle. After waiting 30 seconds, you can safely power down the Equalizer.

Rebooting Down Equalizer

You will only need to reboot the Equalizer after you have configured its failover.

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Reboot Equalizer** from the Equalizer menu in the main menu bar. A confirmation dialog box appears.
3. In the confirmation dialog box, click **OK** to confirm that you really want to reboot Equalizer.

6 Monitoring Equalizer Operation



The Equalizer Administration Interface provides several monitoring mechanisms that allow you to view the following:

- Global configuration information and connection statistics for Equalizer
- A status summary of currently configured clusters and servers
- The Equalizer system event log
- Cluster configuration parameters
- Server configuration parameters
- Graphical displays of the connection history for individual clusters and servers
- Real Path Server information and Real Path log

Displaying Equalizer Information

The Equalizer status screen displays information about Equalizer's operation modes and overall connection statistics:

- **Equalizer version** shows the current, running version of the Equalizer software.
- **system ID** shows the MAC address of the Equalizer unit.
- **platform** shows the type of Equalizer
- **external interface** is the name of this interface.
- **internal interface** is the name of this interface.
- **external address** is Equalizer's external IP address.
- **internal address** is Equalizer's internal IP address.
- **stale connection timeout** indicates the number of seconds before a stale connection is dropped.
- **passive FTP Translation** indicates whether PASV FTP mode is enabled or disabled.
- **failover mode** signifies whether this Equalizer is a primary or backup unit.
- **Envoy geographic load balancing** denotes whether geographic load balancing is currently enabled. This information appears only on the E350 and E450 platforms.
- **SSL acceleration** shows whether the optional XCEL™ card is installed, which enables SSL acceleration. This information appears only on the E350 and E450 platforms.
- **L4 total connections processed** is the number of Layer 4 (L4) connections processed in the last second.
- **L4 peak connections processed** shows the peak number of L4 connections processed per second since the beginning of this Equalizer session.

- **L4 connections timed-out** displays the number of L4 connections that have timed out. If Envoy is enabled, the following DNS status information appears at the bottom of the Current Status section:

- **DNS requests received** displays the total number of DNS requests received.
- **corrupt DNS requests received** shows the number of invalid DNS requests received.
- **DNS requests received for unknown domains** displays the number of unrecognized DNS requests received.

Users of the Equalizer E350/450 will also see this information:

- **L7 current active connections** is the number of active Layer 7 (L7) connections.
- **L7 total connections processed** shows the number of L7 connections processed in the last second.
- **L7 peak connections processed** is the peak number of L7 connections processed per second since Equalizer was started.

The screenshot shows the 'Equalizer status' page with a 'menu' link in the top right. The status information is as follows:

```

Equalizer version 7.2.0c
system ID 00:0e:a6:b6:cd:df:9a:9c:56:38
platform e450

external interface em1
internal interface em0
external address
internal address 10.0.0.66

stale connection timeout 50 sec
passive FTP translation disabled
failover mode primary

Envoy geographic load balancing enabled
SSL acceleration disabled
L4 total connections processed 0
L4 peak connections processed 0 per second
L4 connections timed-out 0
DNS requests received 0
corrupt DNS requests received 0
DNS requests received for unknown domains 0
L7 current active connections 0
L7 total connections processed 0
L7 peak connections processed 0
    
```

Figure 30 Equalizer status information

To display the global parameters, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. At the top of the column in the left frame, click the **Equalizer** entry. The **Equalizer status** screen appears in the right frame. You can also display this information by selecting **Equalizer Status** from the View menu in the main menu bar.

Displaying the Virtual Cluster Summary

The Virtual Cluster Summary (see Figure 31) lists the currently configured virtual clusters and their associated servers as well as the weight and status of each server.

To view the Virtual Cluster Summary, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. Select **Cluster Summary** from the View menu in the main menu bar. The **cluster summary** screen appears in the right frame.

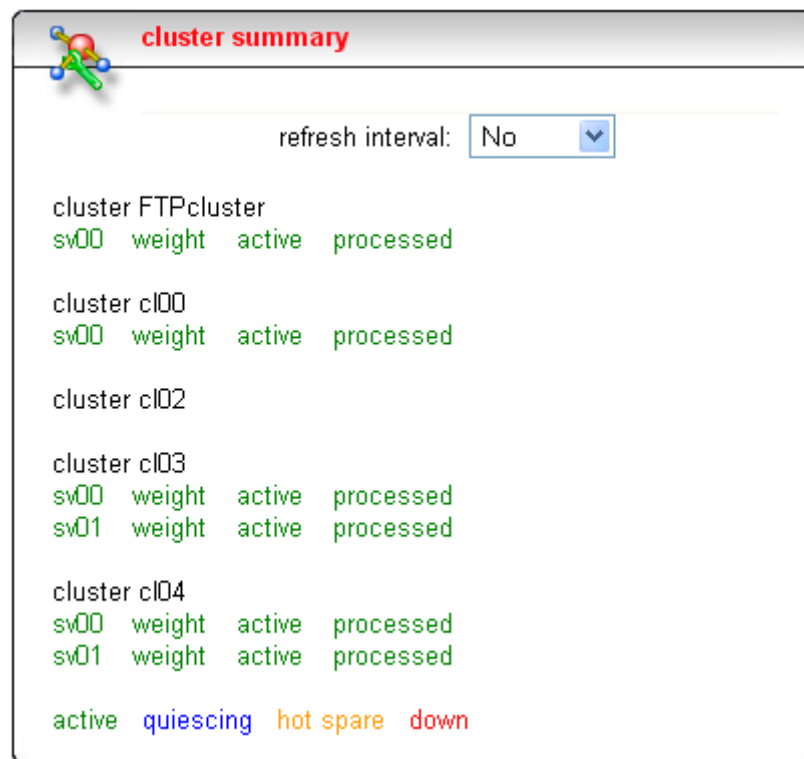


Figure 31 Viewing cluster summary information

This summary displays the status at the time the page was loaded. To set this information to automatically refresh, select a refresh interval.

The cluster summary indicates the following server states:

- Servers shown in green are currently active.
- Servers shown in blue are quiescing, that is, handling current connections but not accepting new ones.

- Servers shown in yellow are configured as hot spares.
- Servers shown in red are down. Equalizer monitors the status of active servers by periodically probing the IP address and Port specified by the server endpoint. If these probes fail two consecutive times, it marks the server *down*, gives the server a weight of zero, and stops routing new requests to that server. A server probe might fail even if the server machine is up and running. For instance, if the HTTP server daemon fails on a server machine, Equalizer will refuse connections to that endpoint.

Equalizer periodically queries servers that have gone down to determine if they have become available again. When a server comes back online, Equalizer begins to route requests to the server, slowly increasing the server's weight to its full capability.

For each server, the summary displays the following information:

weight: The server weights determine the relative proportion of connection requests that Equalizer routes to each server. If you have enabled automatic load balancing, these weights are the current, dynamically-adjusted values, not the static weights initially assigned by the administrator.

- **active:** The number of connections currently being processed by the server.
- **processed:** The total number of connections that have been processed by the server since the system was rebooted.

Displaying the System Event Log

The System Event Log (see Figure 32) displays start-up and server status messages. You can view the last 20, 50, 100, 200, 500, or 1000 entries.

To view the system event log and optionally change the number of entries on display, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. Select **Event Log** from the View menu in the main menu bar. The **log viewer** screen appears in the right frame.

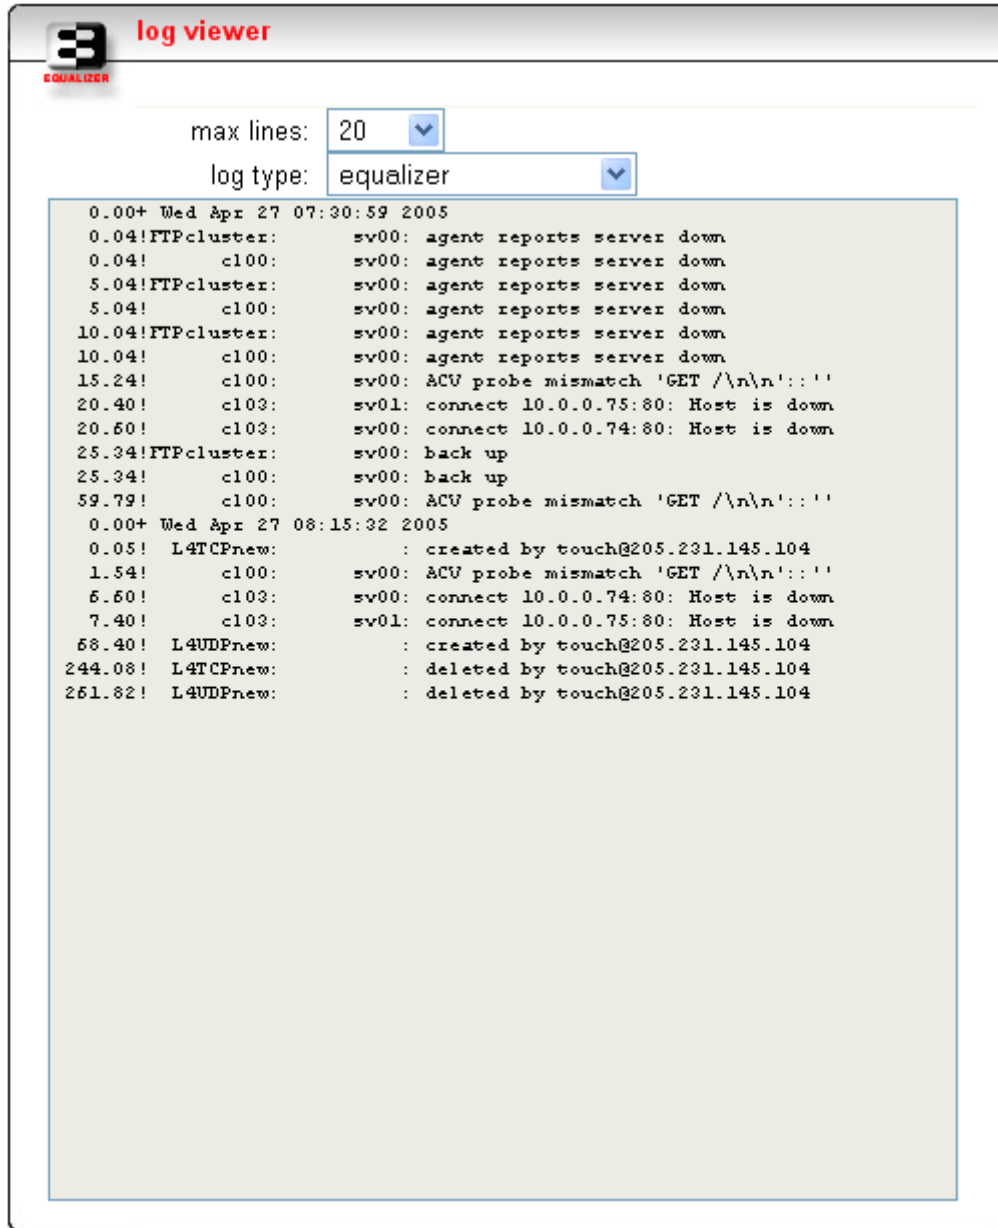


Figure 32 Viewing the system event log

3. To change the number of items displayed, select a value from the drop-down list; then click **Set**.
4. To look at the logs for the Equalizer, a virtual cluster, or the operating system, use the **log type** drop-down list.

To export the log, you can cut and paste its contents.

Displaying Cluster Information

The **cluster** screen (see Figure 33) displays information about a cluster's configuration. To display the parameters for a cluster, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the cluster whose parameters you want to view. The **cluster** screen appears in the right frame.

The screenshot shows the configuration page for cluster 'c100'. The page has a title bar with a cluster icon and the text 'cluster c100' and 'menu'. Below the title bar is a table of parameters and their values. At the bottom, there are two sections for flags and cookie_flags.

protocol	http
ip	10.0.0.67
port	80
policy	round_robin
responsiveness	5
cookie age	0
cookie domain	
cookie path	
ACV probe	GET /n\n
ACV response	<HTML>
server agent port	0

flags

advanced	<input type="checkbox"/>
disable	<input type="checkbox"/>
ignore case	<input checked="" type="checkbox"/>
server_agent	<input checked="" type="checkbox"/>
spooF	<input checked="" type="checkbox"/>
persist	<input checked="" type="checkbox"/>
once only	<input checked="" type="checkbox"/>

cookie_flags

always	<input checked="" type="checkbox"/>
--------	-------------------------------------

Figure 33 Viewing cluster information

The **cluster** screen shows the selected load balancing policy, the load-balancing responsiveness setting, the persistence parameters, and the server agent parameters. For more information about how Equalizer uses these parameters, see “Adding a Virtual Cluster” on page 65.

Plotting Cluster Performance History

The Plot Cluster History feature (see Figure 34) enables you to view a graphical representation of the performance history for any cluster. To plot the performance history for a cluster, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the cluster whose history you want to view.
3. Select **Plot Cluster History** from the local menu in the **cluster** screen. The graphical history for the selected cluster appears.

By default, the service time and active connections are plotted for the previous five minutes. To change the information plotted, select the categories and duration you want to plot and click the **Plot** button.

To zoom in on a portion of the graph, click the target area.

You can plot five values for a cluster:

- **Servers** is the average computed load of all the servers in the cluster. Because server computed loads are normalized by the cluster-wide average, the cluster-wide average should be 100. Certain events (for example, rapid fluctuations in the load, rebooting servers, and restarting application daemons such as httpd) can cause spikes in the computed load for the cluster.
- **Service Time** is the average service time of all of the servers in the cluster. The service time is the time it takes a server to start sending reply packets once it receives a client request. The average service time is a reasonable indication of the overall performance of the cluster.
- **Active Connections** is the total number of active connections on the servers in the cluster.
- **Hit Rate** is the number of connections served by the cluster each second. This is a good indication of how many “hits” the site is getting.
- **Server Agent** is the average of the dynamic server agent values for all servers in the cluster. If you have not configured server agents, this value defaults to 50 internally (that is, the agent sends 50 to the load balancing algorithm) but displays a value of 0.

For more information about these values, see the descriptions in “Plotting Server Performance History” on page 59.

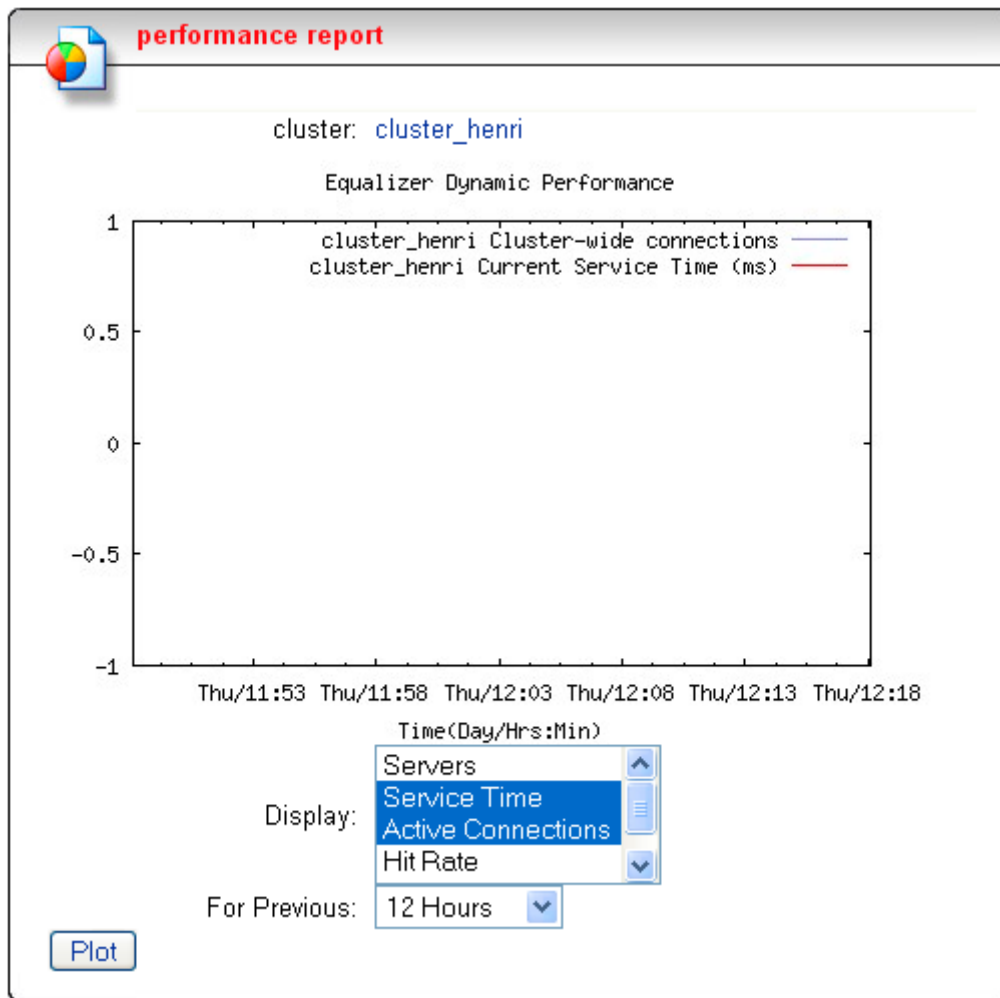


Figure 34 Viewing a cluster’s graphical history

Displaying Server Information

The server screen (see Figure 35) provide information about a particular server, including the following:

- The server’s name and the name of the cluster to which the server belongs.
- The server’s IP address and port.
- The static weight the administrator assigned to the server.
- Other configuration information such as being a hot spare or being quiesced.

To display a server’s parameters, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the server whose parameters you want to view. The server's parameters appear in the right frame.

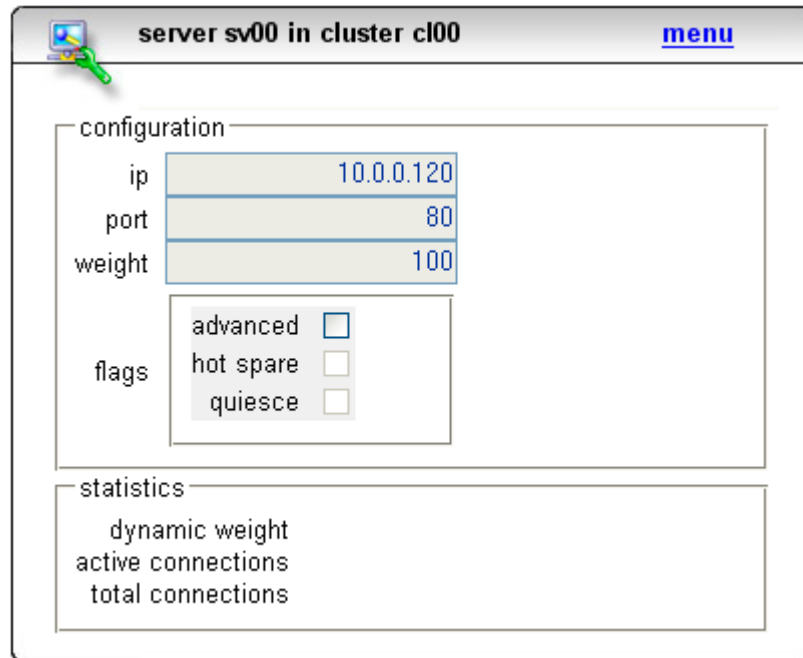


Figure 35 Viewing server information

Plotting Server Performance History

The Plot Server feature (see Figure 36) enables you to view a graphical representation of the performance history for any server.

To plot the a server's performance history, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the server whose history you want to view.
3. Select **Plot Server History** from the local menu in the server screen. The graphical history for the selected server appears.

By default, the active connections, service time, computed load, and dynamic weight are plotted for the previous 30 minutes. To change the information plotted, select the categories and duration you want to plot and click the **Plot** button.

To zoom in on a portion of the graph, click the area in which you are interested.

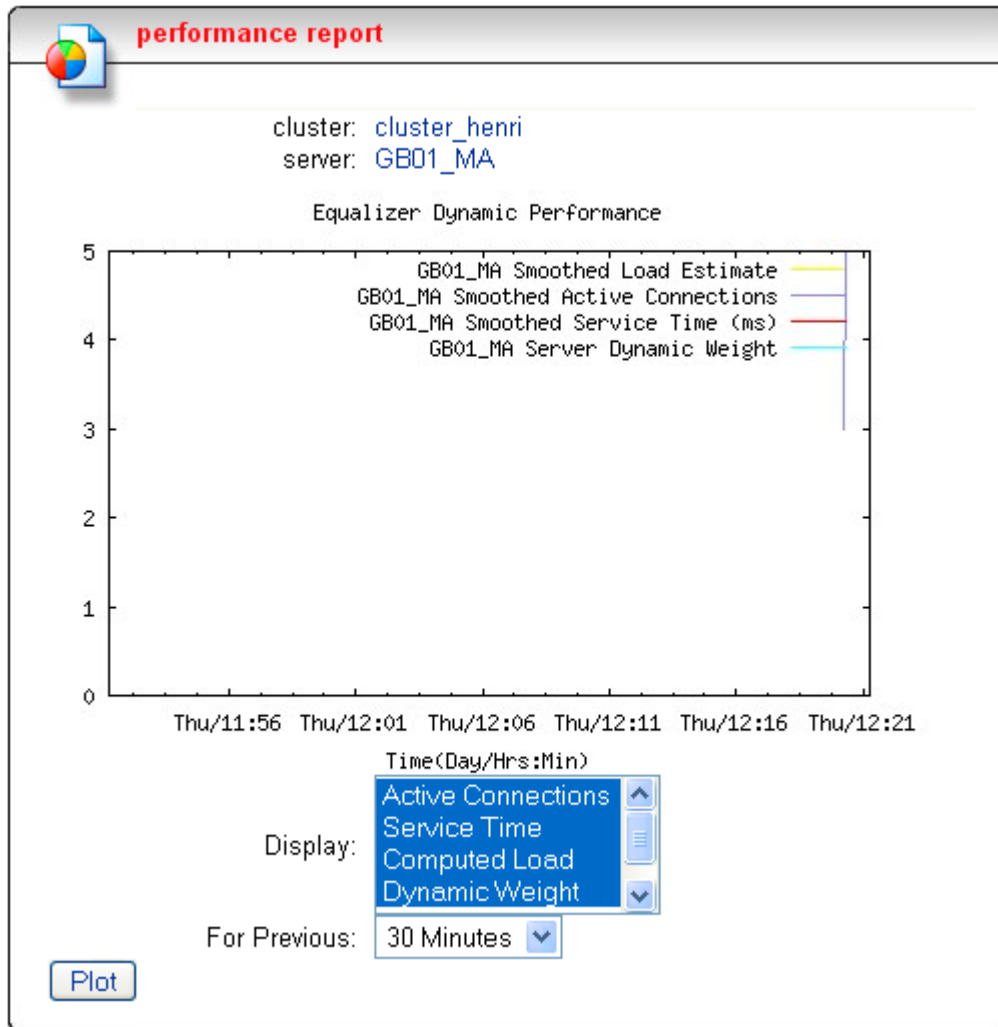


Figure 36 Viewing a server's graphical history

You can plot five values for a server:

- **Active Connections** shows the number of active connections on the server. Equalizer "smooths" the connection count using a sliding-window smoothing algorithm before being plotted. If you have enabled the sticky timer, note that the number of active connections on a server will be higher.
- **Service Time** indicates the time it takes a server to start sending reply packets once it has received a client request. This value is very small for servers that are primarily serving static HTML pages—typically 100-200 milliseconds. If the server is serving many active pages and cgi-bins, this value will be much higher. The service time increases when the server is under heavy load because client requests are queued until the server can handle them.
- **Computed Load** is a measure of the performance of the server relative to the overall performance of the cluster. Equalizer tries to normalize the cluster-wide computed

load value to 100. If the server's computed load value is above 100, it is performing below the overall cluster performance.

Equalizer derives a server's computed load value from its service time, number of active connections, and server agent value (if configured). It also takes into account the load balancing policy used by the cluster.

Ideally, a server's computed load should be around 100, though values in the range 85 to 115 are reasonable. If the server's computed load is higher than 115, the server is not performing well and you may need to add servers or upgrade to better servers. If you are using adaptive load balancing, Equalizer lowers the server's dynamic weight to reduce the number of connections sent to that server. If the server's computed load value is less than 85, the server is performing very well and Equalizer will attempt to improve cluster-wide performance by increasing the server's dynamic weight to direct more traffic to it. Such adjustments to the server's weight will in turn affect its computed load value.

- **Dynamic Weight** is the percentage of incoming traffic that Equalizer dispatches to this server. For example, if the cluster has three servers with dynamic weights of 100, 80, and 120, the first server will get $100/(100+80+120)$ or 33.3% of the incoming traffic.

If a server is down, its dynamic weight is zero. If a server crashes and reboots, the period that the server was down shows up as a gap in the dynamic weight plot.

If you are not using adaptive load balancing (for example, the load balancing policy is set to *round robin* or *static weight*), Equalizer does not use dynamic weights. For more information about setting the load balancing policy and adaptive load balancing, refer to "Configuring a Cluster's Load-Balancing Options" on page 70.

- **Server Agent** is the value that the server agent daemon returns. When queried, the server agent returns a value in the range 0-100. If you have not configured the cluster to use the server agent or the server agent daemon is not running on this server, the server agent value defaults to 50 internally (that is, the agent sends 50 to the load balancing algorithm) but displays a value of 0.

Server agent values above 60 to 70 indicate that the server is overloaded. If this persists and you have enabled adaptive load balancing, Equalizer responds by reducing the server's dynamic weight so that fewer requests are routed to the server.

Note – If all your servers have server agent values above 70, you probably have more traffic than your servers can handle efficiently. In this case, Equalizer can help by intelligently managing the overload, but the long-term solution is to upgrade the servers or add new ones.

Displaying Geographic Cluster Parameters

If you have installed Envoy for your Equalizer, you can view information about each of the geographic clusters that you have configured. For more information about Envoy, refer to Chapter 8, "Administering Geographic Clusters" on page 99.

To view the cluster-wide parameters, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the geographic cluster whose parameters you want to view. The Geographic Cluster Parameters screen appears in the right frame.

This page contains the following information:

- **Geographic Cluster**, which is the name of the cluster.
- **DNS TTL**, which is the amount of time, in seconds, that a name server is allowed to cache the domain information.
- **MX Exchanger**, which is the fully-qualified domain name that Equalizer will return if Equalizer receives a “mail exchanger” request for this geographic cluster. The mail exchanger is the host responsible for handling email sent to users in the domain.
- **Load Balancing Method** indicates the load-balancing method: round trip, adaptive, site load, or site weight. (For descriptions of these methods, refer to “Configuring a Geographic Cluster’s Load-Balancing Options” on page 105.)
- **Load Balancing Response** shows the type of response: slowest, slow, medium, fast, or fastest. This value controls how aggressively Equalizer adjusts the site’s dynamic weights. (For more information about the response settings, refer to “Adding a Geographic Cluster” on page 104.)
- **ICMP Triangulation** shows whether you have enabled ICMP triangulation, which routes client requests to the closest site geographically.

Plotting Geographic Cluster Performance History

If you have installed Envoy for your Equalizer, you can use the Plot Geographic Cluster feature to view a graphical representation of the performance history for the selected geographic cluster.

You can plot four values for a geographic cluster:

- **Request Rate** shows the number of requests received for the cluster per minute.
- **Active Requests** displays the number of requests that Equalizer is in the process of routing.
- **Network Latency** displays the average triangulation time when at least one site was able to respond. (This value does not include clients for which the default site was selected.)
- **Site Summary** shows the number of requests directed to all sites in the cluster for the specified duration. This plot appears by default when the plot site page is opened.

Note – You can only display the site summary separately; you cannot plot the site summary on the same graph as the other values.

To plot the performance history for a geographic cluster, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the geographic cluster whose history you want to view. The Geographic Cluster Parameters appear in the right frame.

3. Select **Plot GeoCluster History** from the local menu in the Geographic Cluster Parameters frame. The graphical history for the selected cluster appears in the right frame. By default, the site summary for the previous 30 minutes appears.
4. To change the information being plotted, select the categories and duration to be plotted; then click the **Plot** button. (To zoom in on a portion of the graph, click the area in which you are interested.)

Displaying Site Information

If you have installed Envoy, you can view configuration and status information for particular sites in a geographic cluster.

To view the information for a particular site, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the site whose information you want to view. The Site Parameters page appears in the right frame.

The Site Parameters page displays the following site parameters:

- **Geographic Cluster** is the name of the geographic cluster to which this site belongs.
- **Site** is the name of the target site.
- **Site IP Address** is the site's IP address.
- **Static Weight** shows the static weight assigned to the site.
- **Default Site** indicates whether the target site is the default site.
- **Resource** shows the IP address and port of the resource being monitored for this site.
- **Agent's Address** is the IP address of the Equalizer agent running on the site.
- **Resource Keepalive** shows the number of seconds between resource availability checks. If a resource fails its availability check, its site will not be returned to clients. Even after a resource is declared dead, Equalizer performs availability checks to determine when the resource is restored.

In addition to the site parameters, the site's current status appears as follows:

- **Resource Load** shows the load on the above resource that the Equalizer agent calculates. The load incorporates data on resource response time, number of active requests, and load-balancing variables.
- **Agent Retries** shows the number of probes Equalizer re-sent to its agent.
- **Agent Misses** shows the number of Equalizer-to-agent probes that received no response. Interruptions in network connectivity between the Equalizer server and site agents and site failures can result in missed probes.
- **Triangulation Time-outs** indicates the number of agent-to-client triangulation probes that timed out before Equalizer received a response.
- **Resource Errors** indicates the number of Equalizer-to-agent probes that returned a resource-unavailable error. If the Envoy on the remote site determines that the requested resource is unavailable, it returns a resource unavailable error.

- **Site Returned** shows the number of clients directed to this site. You can compare this number with the values for other sites to determine the relative number of users sent to each site. If a value for one site is zero and the others are non-zero, consider why the zero site has no traffic.
- **Returned as Default** indicates the number of clients directed to the default site.
- **Average Ping Time** shows the average triangulation time for all clients successfully contacted from this site. This represents all of the triangulation probes—whether or not this site was selected to process the request. This value gives you an idea of the network latency from this site to the user population. You can compare this value with the same value for other sites.

Plotting Site Performance History


If you have installed Envoy, the Plot Site feature enables you to view a graphical representation of the performance history for the selected site. To plot the performance history for a site, follow these steps:

1. Log into the Equalizer Administration Interface in either view or edit mode.
2. In the left frame, click the name of the site whose history you want to view. The Site Parameters appear in the right frame.
3. Select **Plot Site History** from the local menu in the Site Parameters frame. The graphical history for the selected site appears in the right frame.

By default, Equalizer plots the Request Rate and Resource Down values for the previous 30 minutes. To change the information plotted, select the categories and duration to be plotted; then click the **Plot** button. (To zoom in on a portion of the graph, click the area in which you are interested.)

You can plot the following six values for a site:

- **Probes Missed** is the number of requests in which an agent failed to reply to Equalizer's probes.
- **Triangulation Errors** shows the number of ICMP ECHO requests that the agent at this site sent to clients and for which the agent received no response.
- **Resource Down** indicates that the target resource failed to respond during the period plotted.
- **Site Chosen** shows the number of times that Equalizer returned this site in response to a client query.
- **Network Latency** shows the average network distance, in milliseconds, between the agent at this site and the clients that made DNS requests.
- **Resource Load** is the relative workload of this site during the plotted period.



Working with Virtual Clusters

A virtual cluster acts as the network-visible front-end for a group of servers. Use the Equalizer Administration Interface to add, configure, or remove virtual clusters.

Adding a Virtual Cluster

To add a new virtual cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Virtual Cluster** from the Add menu in the main menu bar. The **add cluster** screen appears in the right frame (see Figure 37). Another way to display this screen is to view the Equalizer status and select **Add Virtual Cluster** from the local menu.

The screenshot shows the 'add cluster' configuration window. The fields and their values are as follows:

cluster name	cl01
protocol	HTTP
ip	
port	80
policy	round robin
responsiveness	medium
cookie age	0
cookie domain	
cookie path	
ACV probe	
ACV response	
server agent port	0

The 'flags' section contains the following checked options:

- ignore case
- spoof
- persist
- once only

The 'cookie_flags' section has the following checked option:

- always

Figure 37 Adding a virtual cluster

3. Enter the **cluster name**, which is the logical name for the cluster, or accept Equalizer's default. Each cluster must have a unique name that begins with an alphabetical character (for example, *CPImages*).
4. Select one of the following **protocol** types for the cluster:
 - **HTTP**, Equalizer passes web server requests and route requests to particular servers based on the content of the request and various load-balancing criteria. (This protocol supports Layer 7 load balancing.)

- **HTTPS**, Equalizer passes secure web server requests and route requests to particular servers based on the content of the request and various load-balancing criteria. (This protocol supports Layer 7 load balancing.)
 - **L4 TCP**, Equalizer passes TCP-based requests and route requests based on configured load balancing criteria, the IP address, and TCP port number. Load balancing based on generic connection protocols can be quite efficient; however, routing decisions cannot take into account the content of the request. (This protocol supports Layer 4 load balancing.)
 - **L4 UDP**, Equalizer passes TCP-based requests and route requests based on configured load balancing criteria, the IP address, and UDP port number. Load balancing based on the generic connection protocols can be quite efficient, but routing decisions cannot take into account the content of the request. (This protocol supports Layer 4 load balancing.)
5. Enter the **ip** address, which is the dotted decimal IP address of the cluster. The IP address of the cluster is the external address (for example, 199 . 146 . 85 . 0) with which clients connect to the cluster.
 6. Enter the **port**, which is the numeric port number. For HTTP clusters, the port defaults to 80. For HTTPS clusters, the port automatically defaults to 443.

Note – In the typical Equalizer setup, configure the servers in an HTTPS cluster to listen and respond using HTTP. Equalizer rewrites (munges) responses from the server so that they are HTTPS. You can direct Equalizer pass these responses without rewriting them by enabling the cluster flag **dont munge**.

7. For all cluster protocols, choose the appropriate load-balancing **policy** to be used by this cluster. Choose from round robin (default), static weight, adaptive, fastest response, least connections, or server agent. For more information, refer to “Configuring a Cluster’s Load-Balancing Options” on page 70.
8. Enter values for:
 - **responsiveness** sets the load-balancing response setting for this cluster. For more information, refer to “Configuring a Cluster’s Load-Balancing Options” on page 70.
 - **ACV probe** is the active content verification probe string. For more information, refer to “Using Active Content Verification (ACV)” on page 76.
 - **ACV response** is the active content verification response string. For more information, refer to “Using Active Content Verification (ACV)” on page 76.
 - **server agent port** is the port used to contact server agents.
9. Set the flags:
 - **disable** causes the cluster to be unavailable. Use this flag before you modify a cluster’s parameters
 - **server agent** has Equalizer use server agents gather performance statistics from the servers in the cluster. If you enable this option, you must run Server Agent daemons on each server in the cluster and must specify a value in **server agent port**. See the appendix, “Using Server Agents” on page 109, for more information about configuring server agents.
 - **ignore case** causes all of the cluster’s match rules to use case insensitive comparisons when this box is checked. You can override this setting by changing **ignore case** for a specific match rule.

10. For HTTP and HTTPS clusters, choose from the following options:
 - **spoof** causes Equalizer to spoof the client IP address when Equalizer routes a request to a server in a virtual cluster. This option is checked by default. If you disable this option, the server receiving the request will see the Equalizer's address as the client address because the TCP connection to the client is terminated when the request is routed. When this is enabled, Equalizer must be the default route.
 - **persist** instructs Equalizer to use cookies to maintain a persistent session between a client and a particular server. This option is on by default. Equalizer "stuffs" a cookie into the server's response header on its way back to the client. This cookie uniquely identifies the server to which the client was just connected. With **persist** enabled, Equalizer routes only the first request from a client using load balancing criteria.
 - **once only** limits Equalizer to match only the first request of any client making multiple requests across a TCP session.
11. For HTTP and HTTPS clusters, if you enable **persist**, you may need to adjust the following:
 - **always** includes a cookie in the response whether or not the server actually set a cookie. If this is not selected, Equalizer only sends a persistence cookie when the server sends a cookie of its own.
 - **cookie age** sets the time, in seconds, over which the client browser maintains the cookie. After the specified number of seconds have elapsed, the browser can delete the cookie and any subsequent client requests will be handled by Equalizer's load-balancing algorithms.
 - **cookie domain** limits the presented cookie only to servers whose host name is within the specified domain. For example, if the cookie domain is `coyotepoint.com`, the browser will only present the cookie to servers in the `coyotepoint.com` domain (for example, `www.coyotepoint.com` or `my.coyotepoint.com`).
 - **cookie path** presents the cookie only when the path component of the request URI has the same prefix as that of the specified path. For example, if the cookie path is `/store/`, the browser presents the cookie only if the request URI includes a path such as `/store/mypage.html`.
12. For HTTPS clusters, choose from the following options:
 - **x509 verify** has Equalizer check that the certificate meets the X.509 standard when you upload a certificate. Certain self-signed or chained certificates will not pass this verification and in that instance, you will want to disable the test. To see this flag, check the **advanced** flag.
 - **dont munge** forces Equalizer to pass responses from the cluster's servers without rewriting them. The servers in a typical HTTPS cluster use the HTTP protocol—Equalizer normally rewrites their responses so that they are HTTPS.
13. For L4 TCP and L4 UDP clusters, choose from the following options:
 - **sticky time** is the number of seconds that Equalizer should "remember" connections from clients. If you don't need sticky connections, set this option to 0. For more information, refer to "Enabling Sticky Connections" on page 75.
 - **intercluster sticky** is an option that when enabled ensures that Equalizer directs requests from a particular user to the same server, even if the connection is to a different virtual cluster. For more information, refer to "Enabling Sticky Connections" on page 75.
 - **probe ssl** (L4 TCP only) causes Equalizer to use SSL when it uses the **ACV probe** string.

14. Click the **commit** button to add the virtual cluster.

Equalizer can refuse an Add Cluster command for several reasons, including:

- Attempting to add a cluster address that is already configured or is configured as a server address
- Specifying an invalid cluster names
- Specifying an invalid IP address or port number
- Attempting to add more clusters than are supported by Equalizer

When you check the advanced check box, you will see additional fields. For most operations the default values are acceptable. The modifiable fields are described below:

- **send buffer** applies to L7 clusters and is the amount of memory in kilobytes reserved by each L7 proxy process to store outgoing data before it is placed on the network interface.
- **receive buffer** applies to L7 clusters and is the amount of memory in kilobytes reserved by each L7 proxy process to store data that has been received on an interface before it is processed by an L7 proxy process.
- **request max** applies to L7 clusters and is the maximum number of kilobytes allotted for HTTP request headers.
- **response max** applies to L7 clusters and is the maximum number of kilobytes allotted for HTTP response headers.
- **cookie generation** applies to L7 clusters and is a value added to cookies when the cookie scheme is 2 or greater. In order for cookies to be valid, **cookie generation** must match the equivalent number embedded in the cookie. Conversely if you need to invalidate old cookies, increment this number.
- **probe delay** is the number of seconds between successive probes of the cluster's servers.
- **connect timeout** applies to L7 clusters and is the time in seconds that Equalizer waits for a server to respond to a connection request.
- **client timeout** applies to L7 clusters and is the time in seconds that Equalizer waits before closing an idle client connection.
- **server timeout** applies to L7 clusters and is the time in seconds that Equalizer waits before closing an idle server connection.
- **cipher suite** applies to HTTPS clusters and restricts the cipher suite offered by the server. When XCEL is detected, the Equalizer restricts the default cipher-suite to those accelerated by XCEL. If this is too restrictive, you can clear out the contents of this field.
- **sub-daemon max** applies to HTTPS clusters and is the maximum number of sub-daemons servicing the cluster.
- **session cache timeout** applies to HTTPS clusters and is number of seconds that Equalizer waits before disposing of an SSL session cache entry.
- **session cache kbytes** applies to HTTPS clusters and maximum number of kilobytes allotted to an SSL session cache.

- **certify_client** applies to HTTPS clusters and indicates whether a client has to present a valid certificate when making a request.
- **ssl_unclean_shutdown** applies to HTTPS clusters and should be checked if you see errors (cannot see pages) while trying to maintain HTTPS persistent connections over HTTP/1.1. This problem especially applies to connections between Internet Explorer and Apache Servers and usually occurs intermittently.

Deleting a Virtual Cluster

You cannot delete a cluster with servers assigned to it. So, before attempting to delete the cluster, delete all servers from the cluster. For information about removing servers from a cluster, refer to “Deleting a Server” on page 83.

To delete a cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster to be deleted. The cluster’s parameters appear in the right frame.
3. Select **Delete Cluster** from the local menu.
4. When prompted, click **OK** to confirm that you want to remove the cluster permanently.

Configuring a Cluster’s Load-Balancing Options

Configure load balancing policy and response settings for each cluster independently. Multiple clusters do not need to use the same load balancing configuration even if the same physical server machines host them. For example, if one cluster on port 80 handles HTML traffic and one on port 8000 serves images, you can configure different load balancing policies for each cluster.

When you use adaptive load balancing (that is, you have *not* set the cluster’s load balancing policy to round robin or static weight), you can adjust Equalizer to optimize cluster performance. For more information, see “Adjusting a Server’s Static Weight” on page 83.

Equalizer’s Load Balancing Policies

Equalizer supports the following load balancing policies, each of which is associated with a particular algorithm that Equalizer uses to determine how to distribute requests among the servers in the cluster:

- **round robin** load balancing distributes requests equally among all the servers in the cluster. Equalizer dispatches the first incoming request to the first server, the second to the second server, and so on. When Equalizer reaches the last server, it repeats the cycle. If a server in the cluster is down, Equalizer does not send requests to that server. This is the default method.

The round robin method does not support Equalizer’s adaptive load balancing feature; so, Equalizer ignores the servers’ static weights and does not attempt to dynamically adjust server weights based on server performance.
- **static weight** load balancing distributes requests among the servers depending on their static weights. A server with a higher static weight gets a higher percentage of the incoming requests. Think of this method as a *weighted round robin* implementation. Static weight load balancing does not support Equalizer’s adaptive load balancing feature; Equalizer does not dynamically adjust server weights based on server performance.

- **adaptive** load balancing distributes the load according to the following performance indicators for each server.

Server response time is the length of time for the server to begin sending reply packets after Equalizer sends a request.

Active connection count shows the number of connections currently active on the server.

Server agent value is the value returned by the server agent daemon running on the server.

- **fastest response** load balancing dispatches the highest percentage of requests to the server with the shortest response time. Equalizer does this carefully: if Equalizer sends too many requests to a server, the result can be an overloaded server with slower response time. The Fastest Response policy optimizes the cluster-wide response time.

Under Fastest Response, Equalizer checks the number of active connections and server agent values (if configured); but both of these have less of an influence than they do under adaptive load balancing. Even if a server's response time is the fastest in the cluster but its active connection count and server agent values are high, Equalizer might not dispatch new requests to that server.

- **least connections** load balancing dispatches the highest percentage of requests to the server with the least number of active connections. In the same way as Fastest Response, 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.
- **server agent** load balancing dispatches the highest percentage of requests to the server with the lowest server agent value. In a similar way to Fastest Response, Equalizer tries to avoid overloading the server by checking the number of connections and response time. This method only works if server agents are enabled. For more information about server agents, see "Configuring a Cluster to Use Server Agents" on page 73.

Equalizer's Load Balancing Response Settings

The **responsiveness** setting controls how aggressively Equalizer adjusts the servers' dynamic weights. Equalizer provides five response settings: Slowest, Slow, Medium, Fast, and Fastest. The response setting affects the dynamic weight spread, weight spread coefficient, and optimization threshold that Equalizer uses when it performs adaptive load balancing:

- **Dynamic Weight Spread** indicates how far a server's dynamic weight can vary (or *spread*) from its static weight.
- **Weight Spread Coefficient** regulates the speed of change to a server's dynamic weight. The weight spread coefficient causes dynamic weight changes to happen more slowly as the difference between the dynamic weight and the static weight increases.
- **Optimization Threshold** controls how frequently Equalizer adjusts dynamic weights. If Equalizer adjusts server weights too aggressively, oscillations in server weights can occur and cluster-wide performance can suffer. On the other hand, if Equalizer does not adjust weights often enough, server overloads might not be compensated for quickly enough and cluster-wide performance can suffer.

Modifying Equalizer's Load Balancing Options

To change a cluster's load-balancing options (see Figure 38), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster whose parameters to be changed. Equalizer displays the cluster's parameters in the right frame.
3. Select **Change Cluster Parameters** from the local menu. Equalizer opens the modify cluster screen in the right frame.

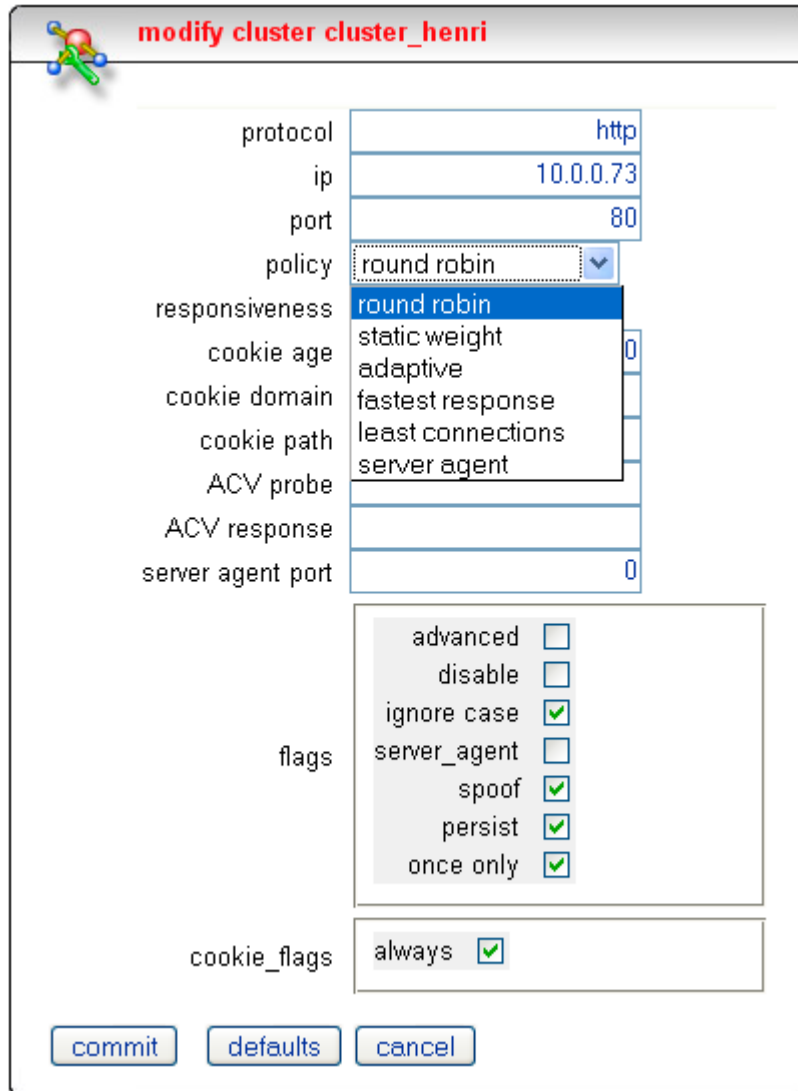


Figure 38 Changing load balancing options

4. Select a **policy**.
5. Choose a **responsiveness**.
6. Click the **commit** button.

Aggressive Load Balancing

After you fine-tune the static weights of each server in the cluster, you might discover that Equalizer is not adjusting the dynamic weights of the servers at all: the dynamic weights are very stable, even under a heavy load. In this case, you might want to set the cluster's load balancing response parameter to *fast*. Then Equalizer tries to optimize the performance of your servers more aggressively; this should improve the overall cluster performance. For more information about setting server weights, see "Adjusting a Server's Static Weight" on page 83.

Dynamic Weight Oscillations

If you notice a particular server's dynamic weight oscillates (for example, the dynamic weight varies from far below 100 to far above 100 and back again), you might benefit by choosing *slow* response for the cluster. You should also investigate the reason for this behavior; it is possible that the server application is behaving erratically.

Providing FTP Services on a Virtual Cluster

Virtual clusters that provide service on the FTP control port (port 21) must be layer 4 and have special requirements:

- FTP clusters occupy two virtual cluster slots, even though only one appears. This permits Equalizers NAT subsystem to rewrite server-originated FTP data connections as they "gateway" to the external network.
- FTP data connections always have a sticky time of one second. This is necessary to support the passive mode FTP data connection that most web browsers use.
- FTP virtual clusters do not support port redirection.

For more information about supporting passive mode FTP data connections, refer to "Enabling Passive FTP Connections" on page 41.

Configuring a Cluster to Use Server Agents

A *server agent* collects performance statistics from a server. If you configure a cluster to use server agents, Equalizer periodically contacts the server agent daemon running on each server and downloads the server performance statistics. You can also customize server agents to report on server resource availability; then Equalizer can stop sending requests to a server if a database or other vital resource is unavailable.

Note – While server agents are not required in most case, when you configure a cluster to use them, each server in the cluster **must** run the appropriate server agent daemon.

To configure a cluster to use server agents (see Figure 39), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster to be configured. The cluster's parameters appear in the right frame.
3. Select **Change Cluster Parameters** from the local menu. The modify cluster screen opens in the right frame.
4. Check the **server agent** checkbox.

- In the **server agent port** field, specify the port used to contact the server agents.

The screenshot shows a configuration window titled "modify cluster cluster_henri". The window contains the following fields and options:

- protocol: http
- ip: 10.0.0.73
- port: 80
- policy: round robin
- responsiveness: medium
- cookie age: 0
- cookie domain: (empty)
- cookie path: (empty)
- ACV probe: (empty)
- ACV response: (empty)
- server agent port: 1510
- flags:
 - advanced:
 - disable:
 - ignore case:
 - server_agent:
 - spoof:
 - persist:
 - once only:
- cookie_flags: always

At the bottom of the window are three buttons: "commit", "defaults", and "cancel".

Figure 39 Configuring a cluster to use server agents

- Click the **commit** button.

For information about writing your own server agents and using agents to monitor server resource availability, see “Using Server Agents” on page 109.

Enabling Persistent Sessions

For HTTP and HTTPS clusters that support Layer 7 (L7) load balancing, you can use cookies to maintain a persistent session between a client and a particular server for the duration of the session. For L4 TCP and L4 UDP clusters, which only support L4 load balancing, you can use IP-address based sticky connections to maintain persistent sessions.

When you use cookie-based persistence (**persist** checkbox) for HTTP and HTTPS clusters, Equalizer “stuffs” a cookie into the server’s response header on its way back to the client. This

cookie uniquely identifies the server to which the client was connected and is included automatically in subsequent requests from the client to the same cluster. Equalizer can use the information in the cookie to route the requests to the same server. If the server is unavailable, Equalizer automatically selects a different server.

Enabling Sticky Connections

Equalizer uses sticky connections to maintain persistent sessions for L4 TCP and L4 UDP clusters.

The *sticky time period* is the length of time over which Equalizer ensures that it directs new connections from a particular client to the same server. The timer for the sticky time period begins to expire as soon as there are no active connections between the client and the cluster. If Equalizer establishes a new connection to the cluster, Equalizer resets the timer for the sticky time period.

When you enable sticky connections, the memory and CPU overhead for a connection increase. This overhead increases as the sticky period increases. You should use the shortest reasonable period for your application and avoid enabling sticky connections for applications unless they need it. For most clusters, a reasonable value for the sticky time period is 600 seconds (that is, 10 minutes). If your site is extremely busy, consider using a shorter sticky time period.

When you enable *inter-cluster stickiness*, you can ensure that Equalizer directs requests from a particular client to the same server even if the connection is to a different virtual cluster. Inter-cluster stickiness only works for L4 clusters. Although L7 clusters automatically provide inter-cluster stickiness, inter-cluster stickiness will not work between L4 and L7 clusters.

You must enable inter-cluster stickiness for all the clusters to be bound together. The clusters with enabled inter-cluster stickiness should contain identical sets of server IP addresses. For example:

```
Cluster www.coyotepoint.com:http
  Server srv1@192.168.0.5
  Server srv2

Cluster www.coyotepoint.com:https
  Server srv1@192.168.0.5
  Server srv2
```

To enable sticky connections (see Figure 40), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster to be configured. The cluster's parameters appear in the right frame.

3. Select **Change Cluster Parameters** from the local menu. The modify cluster screen opens in the right frame.

The screenshot shows a dialog box titled "modify cluster c102". It contains the following fields and options:

- protocol: tcp_14
- ip: 10.0.0.68
- port: 80
- policy: round robin
- responsiveness: medium
- ACV probe: (empty)
- ACV response: (empty)
- server agent port: 0
- sticky time: 10
- flags:
 - advanced:
 - disable:
 - server_agent:
 - inter-cluster sticky:
 - probe ssl:

Buttons at the bottom: commit, defaults, cancel.

Figure 40 Setting the sticky time period

4. In the **sticky time** field, specify the sticky time period in seconds greater than zero.
5. To direct all requests from a particular client to the same server even if the connection is to a different virtual cluster, check the **inter-cluster sticky** checkbox.

Note – You can turn on inter-cluster stickiness only if you have enabled sticky connections by specifying a sticky time greater than zero.

6. Click the **commit** button.

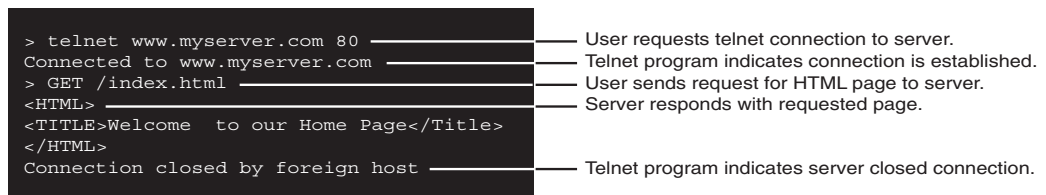
Using Active Content Verification (ACV)

Active Content Verification (ACV) is a mechanism for checking the validity of a server. When you enable ACV for a cluster, Equalizer requests data from each server in the cluster and verifies that the returned data contains a character string that indicates that the data is valid. You can use ACV with most network services that support a text-based request/response protocol, such as HTTP. However, you cannot use ACV with UDP-based services.

Controlling Server Verification Information

Specify an *ACV probe string* and an *ACV response string* to control the information that Equalizer uses to verify the servers. Equalizer uses the probe string to request data from each server. To verify the server's content, Equalizer searches the returned data for the response string. By default, Equalizer expects to receive a response within 10 seconds when performing active content verification. If there is no response or the response string does not appear in the first 1024 characters of the response, the verification fails and Equalizer stops routing new requests to that server. However, if Equalizer uses cookie-based persistence for a HTTP or HTTPS cluster, Equalizer continues to route requests from cookie holders to the server until its weight goes to zero.

For example, the HTTP protocol enables you to establish a connection to a server, request a file, and read the result. Figure 41 illustrates the connection process when a user requests a telnet connection to an HTTP server and requests an HTML page.



```

> telnet www.myserver.com 80
Connected to www.myserver.com
> GET /index.html
<HTML>
<TITLE>Welcome to our Home Page</Title>
</HTML>
Connection closed by foreign host

```

— User requests telnet connection to server.
 — Telnet program indicates connection is established.
 — User sends request for HTML page to server.
 — Server responds with requested page.
 — Telnet program indicates server closed connection.

Figure 41 Retrieving content from a server via telnet.

Equalizer can perform the same exchange automatically and verify the server's response by checking the returned data against an expected result.

Specify an *ACV probe string* and an *ACV response string* to control the information that Equalizer uses to perform the verification. Equalizer uses the probe string to request data from each server. To verify the server's content, Equalizer searches the returned data for the response string.

For example, you can use `GET /index.html` as the *ACV probe string* and you can set the response string to some text, such as `Welcome` in the example in Figure 41, which appears on the home page.

The response string should be text that appears only in a valid response. This string is case-sensitive. For example, most web servers automatically generate error pages that contain valid HTML, so using a response string of `HTML` would not be a good verification of the content.

Enabling ACV

To enable ACV (see Figure 42), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster to be configured. The cluster's parameters appear in the right frame.

3. Select **Change Cluster Parameters** from the local menu. The modify cluster screen opens in the right frame.

modify cluster cluster_henri

protocol

ip

port

policy

responsiveness

cookie age

cookie domain

cookie path

ACV probe

ACV response

server agent port

flags

- advanced
- disable
- ignore case
- server_agent
- spooof
- persist
- once only

cookie_flags

Figure 42 Enabling active content verification.

4. In the **ACV probe** field, specify a non-empty string. Equalizer sends this string to each server in the cluster to request verifiable data.

Note – When you set up a L7 cluster and add a probe string, `\n\n` is automatically added to the end of the string, which is sometimes a confusing change. On the other hand, when you set up a L4 cluster and add a probe string, `\n\n` is not automatically added to the end of the string. The reason for this different behavior is that L7 “knows” the protocol is HTTP/HTTPS but L4 does not know the protocol.

5. In the **ACV response** field, specify a case-sensitive string that is not empty. Equalizer uses this string to verify the data with which the server responds to the ACV probe. For content

verification to succeed, the specified string must appear in the first 1024 characters of the server's response.

6. Click the **commit** button.

Using Secure Server Certificates for HTTPS Clusters

For HTTPS clusters, Equalizer supports the use of secure server certificates. When you install the certificates on Equalizer, Equalizer handles the necessary authentication with clients and communicates in clear text with the servers in the HTTPS cluster. For even faster encryption and decryption, equip your Equalizer with an XCEL card.

Equalizer supports server certificates from Trusted Root Certificate Authorities and from certificate authorities (CAs) without their own Trusted Root CA certificates. If a CA without its own Trusted Root CA certificate issues your certificate, you may need to install two certificates: a server certificate and a chained root certificate for the CA. The chained root certificate associates the server certificate with a Trusted Root CA certificate.

Coyote Point's web site has some information about generating certificates.

Installing a Certificate for an HTTPS or SSL Cluster

To support secure connections to an HTTPS cluster, you must install a secure server certificate issued by a certificate authority (CA) such as VeriSign or Thawte. Until you install the certificate, the cluster is disabled and its name appears red in Equalizer. After installing the certificate, the name for a cluster turns green. Certificates are associated with host names and not IP addresses, therefore you do not need a separate certificate per server in a cluster. You will need a separate certificate per cluster.

You can install certificates in a PEM or PKCS12 format.

To install a certificate, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the HTTPS cluster for which you want install a certificate. The cluster's parameters appear in the right frame. If no certificate is installed for the cluster, you will see a warning message above the parameters stating that you must install a certificate to activate the cluster.
3. Look at the cluster's parameters. If **x509 verify** is checked, Equalizer will verify that the certificate is compliant with the X.509 standard. Certain self-signed or chained certificates will not pass this verification. If you have trouble uploading your certificate, you may need to disable this field.

4. Select **Manage SSL Certificates** from the local menu. The install SSL certificate screen appears in the right frame.



Figure 43 The install certificate screen

5. Enter the path of the certificate file, or click **Browse** to select the file through the Choose File dialog box. The certificate file should be a PEM-encoded or PKCS12-encoded composite Certificate and Private Key.
6. If applicable, enter the password for the certificate. When you enter a password for a password-protected certificate, the certificate is protected only if you have an XCEL card installed on the Equalizer.

Note – When you upload a composite certificate, your private key is stored on the Equalizer. Keep in mind that users with access to the Equalizer will potentially have access to your private key. An optional XCEL SSL accelerator card is available for the Equalizer that provides secure key storage as well as hardware-based SSL encryption and decryption. When you upload your private key to an Equalizer with the XCEL SSL accelerator installed, the key is stored in write-only memory that can only be accessed by the accelerator hardware. This prevents unauthorized access to your private key.

7. Click **upload** to upload and install the specified certificate.

Installing a Chained Root or Intermediate Certificate

If your certificate authority issued you a chained root certificate, you must install this to complete the installation process for HTTPS clusters. The chained root certificate must be in a PEM format.

To install a certificate, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the HTTPS cluster for which you want install a certificate. The Cluster Parameters frame appears in the right frame. If no certificate is installed for the cluster, you will see a warning message above the cluster parameters. The message states that you must install a certificate to activate the cluster.
3. Select **Manage SSL Certificates** from the local menu in the Cluster Parameters frame. The Install SSL Certificate screen appears in the right frame.
4. Enter the fully-qualified name of the certificate file, or click **Browse** to select the file through the Choose File dialog box. The certificate file should be a PEM-encoded or PKCS12-encoded Certificate Authority bundle.
5. Click **Upload** to upload and install the specified certificate.

Using Certificates with the XCEL SSL Accelerator Card

The Equalizer XCEL SSL accelerator card is an add-on for Equalizer that provides secure key storage as well as hardware-based SSL encryption and decryption. All private keys uploaded to an Equalizer with an installed XCEL card get placed in write-only memory that can only be accessed by the accelerator hardware. This prevents unauthorized access to your private keys.

The XCEL card provides 128 kilobits of memory for private keys. This will hold up to 32 four-kilobit keys or 64 two-kilobit keys.

If you install the XCEL card in a production Equalizer, you must delete any HTTPS clusters and add them in order to store the private keys on the card.

Using Certificates in Failover Configurations

In failover configurations, you must install the certificates on the primary and backup Equalizers.

Managing Servers

In this section, you will learn how to work with servers: adding them, adjusting their static weight, shutting them down, and deleting them.

Adding a Server to a Cluster



In general, you should configure your network topology so that Equalizer is the gateway for *all* traffic for its virtual clusters. Each server in a cluster uses Equalizer as the gateway for any response packets to clients that contacted the server through a virtual cluster address. However, you do not need to configure Equalizer as the gateway for the servers in L7 clusters if you have *disabled* IP spoofing for the cluster.

To add a server (see Figure 44) to a virtual cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the cluster to which you want to add the server. The cluster's parameters appear in the right frame.

3. Select **Add Server** from the local menu. The add server screen opens in the right frame.

Figure 44 Adding a server

4. Enter **server name**, which is the server's logical name, or accept Equalizer's default. Each server in a cluster must have a unique name that begins with an alphabetical character, not a numeral (for example, *Phoenix*).
5. Enter **ip** which is the IP address of the server endpoint you are adding to the cluster.
6. Enter **port**, which is the port number of the service on the server machine. Unless you want to set up port redirection, you can usually accept the default value, which is the same as the port of the virtual cluster.

Note – Equalizer performs all the encryption and decryption for HTTPS clusters, so traffic between the Equalizer and the servers in an HTTPS cluster uses the HTTP protocol. When you add servers to an HTTPS cluster, you should configure them on port 80.

7. Enter **weight**, which determines a starting point (static weight) for the percentage of requests to route to each server. For information about selecting an appropriate static weight, refer to "Adjusting a Server's Static Weight" on page 83.
8. Check the **hot spare** checkbox if you plan to use this server as a backup server, in case the other servers in the cluster fail. Checking **hot spare** forces Equalizer to direct incoming connections to this server only if *all* the other servers in the cluster are down. You will not configure most servers as hot spares.

For example, you might configure a server as a hot spare if you are using licensed software on your servers and the license allows you to run the software only on one node at a time. In this situation, you could configure the software on two servers in the cluster and then configure one of those servers as a hot spare. Equalizer will use the second server only if the first goes down, enabling you to make your application available without violating the licensing terms or having to buy two software licenses.

9. Click the **commit** button.

Equalizer can refuse an Add Server command for several reasons, including:

- Attempting to add a server address that is already configured or is configured as a cluster address
- Specifying an invalid IP address or port number
- Attempting to add more servers than are supported by Equalizer

Deleting a Server

To delete a server from a virtual cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the server to be removed.
3. Select **Delete Server** from the local menu.
4. When prompted, click **OK** to confirm that you want to remove the server from the cluster.

If you attempt to delete a server with active connections, a confirmation dialog box appears. Click **Force** to remove the server anyway. This action removes the server and deletes the active connections and the user sessions they represent. To cancel the deletion, click **Cancel**.

Adjusting a Server's Static Weight

Equalizer uses a server's static weight as the starting point for determining the percentage of requests to route to that server. Equalizer assigns servers with a higher static weight a higher percentage of the load. The *relative* values of server static weights are more important than the actual values. For example, if a cluster contains two servers and one server has roughly twice the "horsepower" of the other, setting the static weights to 50 and 100 is equivalent to setting the static weights to 100 and 200.

If Equalizer is performing adaptive load balancing (ALB), you should generally use higher static weights. When you have enabled Equalizer's ALB feature (and the load balancing policy is *not* set to round robin or static weight), using higher static weights will produce finer-grained load balancing. Higher weights enable Equalizer to adjust server weights more gradually; increasing the weight by 1 produces a smaller change if the starting weight is 100 than it does if the starting weight is 50.

Dynamic server weights might vary from 50-150% of the statically assigned values. To optimize cluster performance, you might need to adjust the static weights of the servers in the cluster based on their performance.

Reasonable values for server weights are generally in the range 20-200. When you install servers, set each server's static weight value in proportion to its "horsepower." For example, you might assign an P3/900Mhz-based server a value of 100 and an P3/500Mhz-based server a value of 90. All

the static weights in a cluster do not need to add up to any particular number, but a total that is close to the value 100 is preferable.

Note – Equalizer stops dynamically adjusting server weights if the load on the cluster drops below a certain threshold. For example, if web traffic slows significantly at 4:00 AM PST, Equalizer will not modify server weights until traffic increases again. Because a server's performance characteristics can be very different under low and high loads, Equalizer optimizes only for the high-load case. Keep this in mind when you configure new Equalizer installations; to test Equalizer's ALB performance, you'll need to simulate expected loads.

To change a server's static weight (see Figure 45), follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the server to be modified. The server's parameters appear in the right frame.
3. Select **Change Server Parameters** from the local menu. The modify server screen opens in the right frame.

The screenshot shows a dialog box titled "modify server GB01_MA in cluster cluster_henri". It contains the following fields and options:

- ip**: 10.0.0.74
- port**: 80
- weight**: 100
- flags**:
 - advanced
 - hot spare
 - quiesce

At the bottom of the dialog are three buttons: **commit**, **defaults**, and **cancel**.

Figure 45 Changing a server's static weight

4. Enter the new weight in the **weight** field.
5. Click the **commit** button.

Setting Static Weights for Homogenous Clusters

If all the servers in a cluster have the same hardware and software configurations, you should set their static weights to the same value initially. We recommend that you use a static weight of 100 and set the load-balancing response parameter to *medium*.

As with any new configuration, you will need to monitor the performance of the servers under load for two to three hours. If you observe that the servers differ in the load they can handle, adjust their static weights accordingly and again monitor their performance. You should adjust server weights by small increments; for example, you might set the static weight of one server to 110 and the other

to 90. Fine-tuning server weights to match each server's actual capability can easily improve your cluster's response time by 5 to 10%.

Note – Equalizer's ALB algorithm can take 10-15 minutes to fine-tune cluster performance when you change static weights. After you change static weights, wait 30 minutes before you judge the cluster's ALB performance.

Setting Static Weights for Mixed Clusters

Equalizer enables you to build heterogeneous clusters using servers of widely varying capabilities. Adjust for the differences by assigning static weights that correspond to the relative capabilities of the available servers. This enables you to get the most out of your existing hardware, so you can use an older server side-by-side with a new one.

After you assign relative static weights, monitor cluster performance for two to three hours under load. You will probably fine-tune the weights and optimize performance of your cluster two or three times.

Continue monitoring the performance of your cluster and servers and watch for any trends. For example, if you notice that Equalizer *always* adjusts the dynamic weights so that the weight of one server is far below 100 and the weight of another is far above 100, the server whose dynamic weight is consistently being reduced might have a problem.

Shutting Down a Server Gracefully

To avoid interrupting user sessions, make sure that a server to be shut down or deleted from a cluster no longer has any active connections. When a server's static weight is zero, Equalizer will not send new requests to that server. Connections that are already established continue to exist until the client and server application end them or they time out because they are idle.

To shut down servers in a generic TCP or UDP (L4) cluster, you can set the server's weight to zero and wait for the existing connections to terminate. However, you need to quiesce servers in HTTP and HTTPS (L7) clusters to enable servers to finish processing requests for clients that have a persistent session with the server. When you quiesce a server, Equalizer does not route new connections from new clients to the server, but will still send requests from clients with persistent session with the server to the server. Once all the persistent sessions on the server have expired, you can set the server's static weight to zero so; then Equalizer will not send additional requests to the server.

Removing a Layer 7 Server from Service

To remove a Layer 7 server from service, follow these steps:

1. In the left frame, click the name of the server to be quiesced. The server's parameters appear in the right frame.
2. Select **Change Server Parameters** from the local menu. The modify server screen opens in the right frame.
3. Check the **quiesce** checkbox; then click **commit** to save your changes.
4. Once all the server's persistent sessions have expired, use Change Server Parameters to set the server's static weight to zero.

Removing a Layer 4 Server from Service

To remove a Layer 4 server from service, follow these steps:

1. In the left frame, click the name of the server to be removed. The server's parameters appear in the right frame.
2. Select **Change Server Parameters** from the local menu. The Change Server Parameters dialog box opens in the right frame.
3. Set the server's weight to 0; then click **commit** to save your changes. This action prevents Equalizer from routing new connections to the server.
4. Wait until there are no active connections and the server's idle time is greater than the your application's session lifetime before taking the server offline. To check these values, click on the server in the left frame to show the server's statistics.



Overview of Match Rules

Layer 7 clusters require match rules in order to control the processing of the data stream from the client. Match rules specify the actions to take based on Layer 7 protocol-specific attributes. The most useful action, from a load balancing perspective, is the selection of the set of servers to load balance the requests over. Equalizer's support for Layer 7 content-sensitive load balancing enables you to define match rules for routing HTTP and HTTPS requests. For each virtual cluster, you can specify any number of match rules. Then for each match rule, you specify the subset of servers that can handle requests that meet the rule criteria.

A match rule provides for custom processing of requests within connections. Equalizer provides common and protocol-specific match functions that enable dynamic matching based on the request's contents. Protocol-specific match functions typically test for the presence of particular attributes in the current request. For example, a Layer 7 HTTP virtual cluster can specify matching on specific pathname attributes to direct requests to subsets of servers so that all requests for images are sent to the image servers.

A match rule specifies match expressions that are combinations of match functions with logical operators. This allows for matching requests that have, for example, `attribute A AND NOT attribute B`. You can construct arbitrarily complex logical expressions in this manner.

If the match expression evaluates to *true*, then the data in the request has selected the match rule, and the match body applies, and no further attempts are made to match to subsequent expressions. The *match body* contains statements that affect the subsequent handling of the request. Once the data in the request selects one match expression, no further matching is performed for that request and Equalizer makes a load balancing decision.

Note – Certain Layer 7 protocols can have multiple requests on the same TCP/IP connection, in which case, the default mode of operation is to match each individual request on the stream, not just the initial one. A flag, `once_only`, can be set to match only on the initial request.

If the match expression evaluates to *false*, then Equalizer processes each subsequent match rule in the list of match rules for the virtual cluster until a match occurs. If no match occurs, the connection from the client is dropped. However, virtual clusters created using the Equalizer Administration Interface are always provided with a default match rule, which will always match any request and which will use the entire set of servers for load balancing unless it is modified.

Each virtual cluster can have any number of match rules, and each match rule can have arbitrarily complex match expressions. Keep in mind that Equalizer interprets these expressions for every Layer 7 request processed, so it is a good idea to keep the expressions simple.

General Match Expressions and Match Bodies

A match rule consists of a *match expression* and a *match body*, which identifies the operations to perform if the expression is satisfied by the request. Match syntax is as follows:

```
match name { expression } then { body }
```

Each match has a name, which is simply a label. The name must follow the same restrictions as those for cluster names and server names. All match names within a cluster must be unique.

Match Expressions

Match expressions affect the subsequent processing of the request stream using URI, host, or other information. Match expressions are made up of match functions, most of which are protocol-specific, joined by logical operators, optionally preceded by the negation operator, with sets of beginning and end parentheses for grouping where required. This may sound complex, and it can be, but typical match expressions are simple; it is usually best from a performance perspective to keep them simple.

The most simple match expression is one made up solely of a match function. The truth value (*true* or *false*) of this expression is then returned by the match function. For example, a match function common to all Layer 7 protocols is the `any()` function, which always returns *true*, independent of the contents of the request data. So, the most simple match expression is:

```
any()
```

which will always result in the match rule being selected.

Use the logical NOT operator, (sometimes `!`), to invert the sense of the truth value of the expression. So, you can use the NOT operator to logically invert a match expression, as follows:

```
NOT expression
```

giving rise to the next simplest example:

```
NOT any()
```

which will always result in the match rule not being selected (which is not all that useful in this example).

With the addition of the logical OR (`||`) and logical AND (`&&`) operators, you can specify complex expressions, selecting precise attributes from the request:

```
NOT red() || (round() && happy())
```

Match expressions are read from left to right. Expressions contained within parentheses get evaluated before other parts of the expression. The previous expression would match anything that was not red or that was round and happy.

Unlike the previous example, match functions correspond to certain attributes in a request header.

For example, a request URI for a web page might look like this:

```
Get /somedir/somepage.html http/1.1
Accept: text/html, text/*, *.*
Accept-Encoding: gzip
Host: www.somesite.com
User-Agent: Mozilla/4.7 [en] (Win98; U)
```

Various functions return true when their arguments match certain components of the request URI. Using the above request URI, for example, you could use several match functions: `pathname` returns true if its argument matches `/somedir/somepage.html`, `dirname` returns true if its argument matches `/somedir/`, and `filename` returns true if its argument matches `somepage.html`. Some of the other functions can evaluate the contents of the host attribute in the request URI; `host` (`www.somesite.com`), `host-prefix` (`www`), and `host_suffix` (`somesite.com`).

Some function arguments can take the form of a regular expression¹. Note that you cannot put regular expressions into match expressions except as an argument to a function whose definition admits regular expressions.

Note – Matching regular expressions (*regex*) is many times more processing-intensive than matching other built-in request attributes. So whenever possible, use the other predefined request attributes.

Match Bodies

Match bodies specify the actions to take if the match expression selects the request. This is specified in the form of statements that provide values to variables used by the load balancer to process the request. The most common (and most useful) match body selects the set of servers over which to apply the load balancing:

```
servers = all;
```

The `servers` assignment statement takes a comma-separated list of server names, which specifies the set of servers to be used for load balancing all requests that match the expression in the match rule. The reserved server names `all` and `none` specify respectively the set of *all* servers in the virtual cluster and *none* of the servers in the virtual cluster. If you do not assign servers, none will be available for load balancing; as a result, the connection to the client will be dropped.

In general, you can override most cluster-specific variables in a match body. (You can override protocol-specific variables as well, but that does not always make sense.) One useful example of overriding variables is as follows:

```
servers = s0, s1, s2;
flags = ! once_only;
```

which would load-balance across the specified servers (which first must be defined in the virtual cluster) and also turn off the `once_only` flag for the duration of processing of that connection.

Match Rule Example

A full example of a match rule is:

```
match example {
    client_ip("199.98.84.1")
} then {
    servers = s2, s5;
```

1. Regular expressions are specified according to IEEE Std 1003.2 (“POSIX.2”).

```

        flags = once_only, ! spoof;
    }

```

In this example (the match rule is named “example”), the match function, `client_ip`, has an argument that matches all requests from IP address `199.98.84.1`. Servers `s2` and `s5` are the only ones used for load balancing of matching requests. Additionally, this rule sets the `once_only` flag (that is, we perform processing only on the initial request on this connection) and clears the `spoof` flag (that is, when the connection is made to the server, the server sees a connection to the Equalizer, not to the client).

Constructing Match Rules

The Equalizer Administration Interface allows you to create and modify match rules, without requiring a detailed knowledge of the configuration language syntax. It helps to understand the general concepts of match rules covered in “General Match Expressions and Match Bodies” on page 88.

Viewing the Default Match Rule

All Layer 7 clusters created via the Equalizer Administration Interface start with a single match rule (named Default) that matches all requests and selects all servers.

```

match Default {
    any()
} then {
    servers = all;
}

```

The default rule simply specifies that all servers defined in the cluster should be used for load balancing the request. This rule must remain the last match rule in the ordered list of match rules for a cluster. You cannot modify this match rule.

The default rule can be viewed by clicking in the left frame on **match Default** for any Layer 7 cluster. (If you have not created a Layer 7 cluster, see “Working with Virtual Clusters” on page 65). Figure 46 shows a default match rule.

The screenshot shows a dialog box titled "match Default in cluster TC_MNY" with a "menu" link in the top right. The dialog is divided into two main sections:

- Match Expression:** A text area containing the expression `any()`.
- Load Balancing Settings:** A section titled "load balance with these settings." containing:
 - A "servers" field with an empty text input box.
 - A "flags" field with a text input box containing the word "advanced" and an unchecked checkbox.

Figure 46 A Default match rule shown in the Match Rule dialog box

This screen shows how match rules appear in the administration tool. The first section contains the match rule expression. The second section shows the set of servers used for load-balancing when the expression matches.

Defining a Match Rule

To add a match rule to a virtual cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the Layer 7 cluster to which you want to add the rule.

- In the right frame, select **Add Match Rule** from the local menu. The create match rule screen appears in the right frame.

Figure 47 The Match Rule dialog box in which you create a match rule

- Enter a name for the new rule in the **match name** field.
All match names within a cluster must be unique.
- Select the placement of the rule by choosing a rule from the **immediately before** list box.
The ordering of match rules is important, as they are processed from first to last until one of them evaluates to *true*, at which time the match body is processed. The initial match expression of a new rule, NOT *any()* is one that will always evaluate to *false* meaning that this match rule will never be selected. It is good practice to be cautious when adding new match rules to ensure that all the traffic to a cluster does not get mishandled. A new match rule will not be committed until you click **commit**. You can cancel the entire process by clicking **cancel**.

Note – Note that the GUI displays the logical negation operator as NOT, rather than !.

- To place or modify a match function, click the appropriate part of the expression.
The part of the expression that editor will directly affect is red and the affiliated parts to the selection are green. Pay attention to the colors of various parts of the match expression, these colors show what will be affected.

7. From the drop-down list below the match expression, select the match function with which you want to build or edit the rule. (To learn more about match functions, refer to “Match Functions” on page 94.)

The drop-down list of edit actions are different depending on what you select in the expression and whether the cluster is HTTP or HTTPS. All lists have some common match functions and structural editing operators. In any list of edit actions, *selection* refers to the green and red parts of the match expression and *self* refers to the red portion.

Some of the structural editing operators include the function you are replacing (for example, replace with host AND any). When modifying the structure of an any function, it may be helpful to temporarily change the function to something more distinct (so that you will not have to interpret the expression, “replace with any AND any”).

8. Click the **continue** button, Equalizer shows the new version of the match expression.

Depending on the new function, you may have to fill in information in the **arg0** and **arg1** text boxes. These fields supply arguments, as required, to the selected match function.

If there are any syntax errors, an error screen appears when you click the **continue** button. This most likely occurs if there are missing arguments or syntax errors in the argument strings.

If you click a different part of the match expression without clicking the **continue** button first, you will lose any changes since you last clicked **continue**.

9. You construct complicated Boolean expressions using the structural editing operators.
10. To undo the latest changes, click the **undo** button.
11. To add to or change the match expression, repeat steps 6 through 10. Equalizer continues to show your additions and modifications.
12. Select the servers used to load balancing matching requests for this match rule.
13. Check **advanced** if you want to override the inheritance of **spoof**, **once only**, or **persist**. To override, clear the corresponding **inherit** checkbox and make the appropriate change to the flag.
14. When you have finished specifying expressions for your match rule, click the **commit** button.

Modifying a Match Rule

To edit a match rule, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the match rule to be changed.
3. In the right frame, click **Menu** and select **Edit Match Rule** from the local menu. The modify match rule screen opens in the right frame.
4. Make the desired changes to the match expression using steps similar to the prior section, “Defining a Match Rule” on page 91.
5. Make the desired changed to the list of servers.
6. To save your changes, click the **commit** button.

Removing a Match Rule

To delete a match rule, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the match rule to be deleted.
3. In the right frame, click **Menu** and select **Delete This Match** from the local menu.
4. Click **OK** to confirm that you want to permanently remove the match rule.

Match Functions

To build or edit a match expression, click part of the expression and select a request attribute from a dynamic drop-down list. The highlighted section of the expression determines the contents of the list. For instance, if the current selection is a match function, you will view a list of items that can replace the function. This section lists Equalizer's match functions.

Certain match functions have similar functions that solely apply to the *prefix*, *suffix*, *substr*, or *regex*. Prefix and suffix are self-explanatory. Substr will match an argument to a substring of the entity. Regex interprets an argument as a regular expression and then tries to match. Regular expressions can be very costly to compute, so it is best to use the prefix, suffix, or substr testing functions or Boolean combinations of prefix and suffix testing, rather than the regular expression function.

Common Match Functions

The common match functions are defined below.

any()

This function always evaluates to *true*.

client_ip(string)

This function evaluates to *true* only if the IP address of the client machine making the connection matches the string-valued argument. This function can be useful in restricting match expressions to a particular client, which can aid debugging a new match rule when a cluster is in production. Only the test client will match, leaving other clients to be handled by other match rules.

debug_msg(string)

This function always evaluates to *true*. It writes the single string valued argument str to the debug log. It can be useful for debugging match expressions.

ignore_case()

This function evaluates to true regardless of the capitalization.

observe_case()

This function evaluates to true for the correct capitalization.

HTTP Protocol and Request URI Match Functions

This section includes the functions that test for the attributes of the request protocol level (HTTP 0.9 or above) and the request URI.

Any string comparisons done by any of the functions is case insensitive, that is, the case of strings is ignored.

http_09()

This function takes no arguments and evaluates to *true* if the HTTP protocol used by the request appears to be HTTP 0.9. This is done by inference—an explicit protocol level is absent after the request URI.

host(string)

This function evaluates to *true* if the hostname portion of the request matches the string-valued argument. In the case of HTTP 0.9, the host is a portion of the request URI. All other HTTP protocol versions require a *Host* header to specify the host, which would be compared to the string.

filename(string)

This function evaluates to *true* if the string-valued argument matches the filename portion of the URI path. This portion does not include the trailing pathname component separator, as that is considered part of the directory.

pathname(string)

This function evaluates to *true* if the string-valued argument matches the path component of the request URI.

dirname(string)

This function evaluates to *true* if the string-valued argument matches the directory portion of the path component of the request URI.

HTTP Header Matching Functions

Certain HTTP request headers are searched for when the request is being processed. All match functions dealing with request headers take an initial string-valued argument, which selects the header of interest. If this header is not present in the request, the match function evaluates to *false*. Otherwise, the text associated with the header is examined depending on the particular function. Although HTTP permits a header to span multiple request lines, none of the functions matches text on more than one line. The list of supported headers follows:

Table 48: Supported HTTP Headers for Matching

accept	expect	proxy-authorization
accept-charset	from	range
accept-encoding	host	referer
accept-language	if-match	te
authorization	if-modified-since	trailer

Table 48: Supported HTTP Headers for Matching

cache-control	if-none-match	transfer-encoding
connection	if-range	upgrade
content-length	if-unmodified-since	user-agent
cookie	max-forwards	via
date	pragma	warning

header_prefix(header, str)

This function evaluates to *true* if the selected header is present and if the string-valued argument **str** is a prefix of the associated header text.

header_suffix(header, string)

This function evaluates to *true* if the selected header is present and if the string-valued argument **str** is a suffix of the associated header text.

header_substr(header, string)

This function evaluates to *true* if the selected header is present and if the string-valued argument **str** is a sub-string of the associated header text.

header_regex(header, string)

This function evaluates to *true* if the selected header is present and if the string-valued argument **str**, interpreted as a regular expression, matches the associated header text. Regular expressions can be very costly to compute, so it is best to use the prefix, suffix or sub-string testing functions, or Boolean combinations of prefix, suffix and sub-string testing, rather than the regular expression function.

HTTPS Specific Match Functions

Equalizer permits the construction of virtual clusters running the HTTPS protocol. HTTPS is HTTP running over an encrypted transport, typically SSL version 2.0 or 3.0 or TLS version 1.0. All of the functions available for load balancing HTTP clusters are available for HTTPS. In addition, there are some additional match functions.

Note – Given that HTTPS runs encrypted using SSL and TLS as the transport, in order for any Layer 7 processing, the Equalizer must terminate the SSL/TLS encrypted connection. This can have deleterious effects on performance, as the encryption and decryption process is extremely compute-intensive. A hardware accelerator is available which can be added to the Equalizer platform to ameliorate this problem.

ssl2()

This function evaluates to *true* if the client negotiated the encrypted connection using SSL version 2.0.

ssl3()

This function evaluates to *true* if the client negotiated the encrypted connection using SSL version 3.0.

tls1()

This function evaluates to *true* if the client negotiated the encrypted connection using TLS version 1.0.



Geographic Load Balancing with Envoy

Coyote Point's Envoy geographic load balancer, which is an optional software add-on for the Equalizer product line, supports geographic clustering and load balancing. Geographic clustering and load balancing enables requests to be automatically distributed across servers in different physical locations or on different networks.

Equalizer and its set of servers in a particular location forms a *site* (or Envoy site). A geographic cluster contains multiple sites, and Equalizer's geographic load balancing technology balances incoming requests across those sites.

Note – To perform geographic load balancing, you need to enable the Envoy add-on to your Equalizer system. Not all Equalizer systems allow you to do this. Check the Coyote Point website to see a list of the Equalizer systems that support Envoy.

Equalizer performs the following steps to determine the site in the geographic cluster that should handle the request:

1. The selected Equalizer identifies the geographic cluster that has been configured with the requested domain name—in this case, `www.coyotepoint.com` (see Figure 49).




Figure 49 Sending name resolution requests to an Equalizer in a geographic cluster

2. The selected Equalizer sends *geographic query protocol probes* to *agents* running at each site in the cluster. These probes contain information about the requesting client and the resource that is being resolved. The site handling the resolution request (Site A) also queries its local agent (see Figure 50).




Figure 50 The selected Equalizer queries other Equalizers and its own servers in the geographic cluster

3. The agent checks the availability of the requested resource (see Figure 51).
 - If the resource is unavailable at the agent's site, the agent sends an error message to Equalizer.
 - If the resource is available and ICMP triangulation is enabled, the agent sends an ICMP echo request (*ping*) to the requesting client. When the echo reply arrives, the agent forwards the latency information to the Equalizer that sent the geographic probe.




Figure 51 The selected Equalizer receives availability and triangulation (latency) information

- The Equalizer that sent the geographic probe returns the best Equalizer site to the client's local DNS (see Figure 52).




Figure 52 The client's local DNS receives the best Equalizer site

- The selected Equalizer uses the information gathered from each site to determine the site that is best able to process the request for the client and then forwards the request to that site (see Figure 53).




Figure 53 Site B handles the client's connection

Installing and Configuring Envoy

Each site in an Envoy configuration has an Envoy-enabled Equalizer and any number of servers. Once you've completed the normal Equalizer installation and configuration at each location, you can install Envoy and configure your authoritative name server to work with Envoy.

After you have licensed your Equalizer and Envoy and completed the Envoy installation and DNS configuration described in this section, you can set up the geographic clusters and define the available sites for each cluster through the Equalizer Administration Interface.

Installing Envoy

You must license the Envoy software on each of the Equalizers that will be part of the geographic cluster. Envoy software is pre-installed on each Equalizer and is enabled through the registration and licensing process. To enable Envoy, follow these steps:

1. Follow the registration procedure and make sure that you enter the serial number for your Envoy software.
2. Follow the instructions provided to license your Equalizer.
3. Shut down Equalizer and reboot the machine. (For information about how to safely shut down Equalizer, see "Shutting Down Equalizer" on page 49.)

Configuring the Authoritative Name Server to Query Envoy

You must configure the authoritative name server(s) for the domains that are to be geographically load balanced to delegate authority to the Envoy sites. You need to delegate each of the fully-qualified subdomains to be balanced.

For example (see Figure 54), assume you must balance `www.coyotepoint.com` across a geographical cluster containing two Envoy sites, `east.coyotepoint.com` (at `192.168.2.44`) and `west.coyotepoint.com` (at `10.0.0.5`). In this case, you must configure the name servers that will handle the `coyotepoint.com` domain to delegate authority for `www.coyotepoint.com` to both `east.coyotepoint.com` and `west.coyotepoint.com`. When queried to resolve `www.coyotepoint.com`, `coyotepoint.com`'s name servers should return name server (NS) and alias (A) records for both Envoy sites.

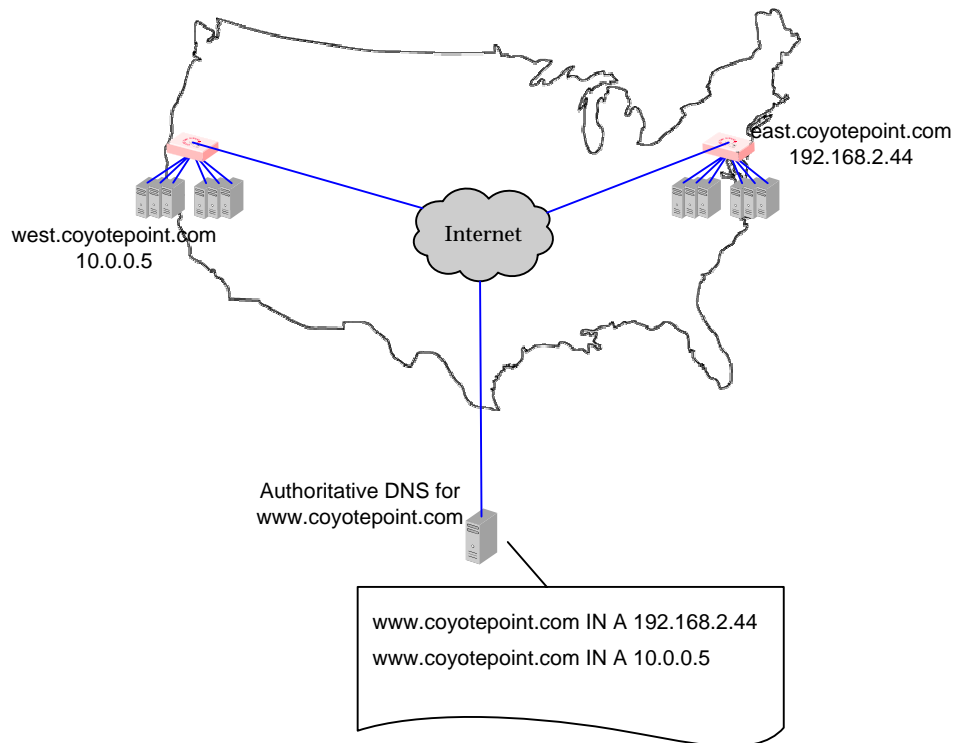


Figure 54 Two-site DNS example

Using Envoy with Firewalled Networks

Envoy sites communicate with each other using Coyote Point's UDP-based Geographic Query Protocol. Similarly, Envoy sites communicate with clients using the DNS protocol. If you protect one or more of your Envoy sites with a network firewall, you must configure the firewall to permit the Envoy packets to pass through.

To use Envoy with firewalled networks, you need to configure the firewalls so that the following actions occur:

- Envoy sites communicate with each other on UDP ports 5300 and 5301. The firewall must allow traffic on these ports to pass between Equalizer/Envoy sites.
- Envoy sites and clients can exchange packets on UDP port 53. The firewall must allow traffic on this port to flow freely between an Envoy server and any Internet clients so that clients trying to resolve hostnames via the Envoy DNS server can exchange packets with the Envoy sites.
- Envoy sites can send ICMP echo request packets out through the firewall and receive ICMP echo response packets from clients outside the firewall. (When a client attempts a DNS resolution, Envoy sites send an ICMP echo request (ping) packet to the client and the client might respond with an ICMP echo response packet.)

Working with Geographic Clusters

This section shows you how to add or delete a geographic cluster and how to configure a geographic cluster’s load-balancing options.

Configuring a geographic cluster and its sites is similar to configuring a virtual cluster and its servers. If you need information about how to access the administration interface, refer to “Introducing the Equalizer Administration Interface” on page 29.

Adding a Geographic Cluster

To add a new geographic cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. Select **Add GeoCluster** from the Add menu in the main menu bar. The Add Geographic Cluster dialog box appears in the right frame. You also can display this dialog box by selecting **Add GeoCluster** from the local menu when you view Equalizer’s global parameters.
3. Enter the **Geographic Cluster Name**, which is the fully-qualified domain name (FQDN) of the geographic cluster (for example, `www.coyotepoint.com`). The FQDN must include all name components up to the top level (com, net, org, etc). Do not include the trailing period.
4. Enter the **CTTL** (cache-time-to-live), which is the length of time, in seconds, that the client’s DNS server should cache the resolved IP address. Longer times will result in increased failover times in the event of a site failure, but are more efficient in terms of network resources. A reasonable value would be 120 (that is, 2 minutes).
5. Enter the **MX Exchanger**, which is the fully qualified domain name to be returned if Equalizer receives a “mail exchanger” request for this geographic cluster. The mail exchanger is the host responsible for handling email sent to users in the domain.
6. Specify the **Load Balancing Method**:
 - **Round Trip**: This method weights the client’s network proximity more heavily than other criteria. This option only works if you enable Ping Triangulation.
 - **Adaptive**: This method takes all available information into account when selecting a site. This setting is a reasonable default.
 - **Site Load**: This method weights the current load at each site more heavily than other criteria.
 - **Site Weight**: This method weights the user-defined static weight for each site more heavily than other criteria.
7. Specify the **Load Balancing Response**. This value controls how aggressively Equalizer adjusts the site’s dynamic weights. Equalizer provides five response settings: Slowest, Slow, Medium, Fast, and Fastest. Faster settings enable Equalizer to adjust its load balancing criteria more frequently and permit a greater variance in the relative weights assigned to sites. Slower settings cause site measurements to be averaged over a longer period of time before Equalizer applies them to the cluster-wide load balancing; slower settings also tend to ignore spikes in cluster measurements caused by intermittent network glitches. We recommend that you select the *Medium* setting as a starting point.
8. Check or clear the **ICMP triangulation** checkbox. When you check ICMP triangulation, Equalizer pings the client and collects latency information, which provides more accurate client

location information. If you do not want to allow Equalizer to ping clients when they make a request, clear the ICMP triangulation checkbox.

9. Click the **Add** button to add the geographic cluster. An entry for the new geographic cluster appears in the left frame.

Equalizer can refuse an Add GeoCluster command for several reasons, including:

- Attempting to add a cluster for a FQDN that is already configured
- Attempting to add more clusters than are supported by Equalizer

Configuring a Geographic Cluster's Load-Balancing Options

You can change the load balancing policy and response settings for a geographic cluster from the Geographic Cluster Parameters frame. Configure these parameters independently for each geographic cluster. (For more information about the load balancing policy and response settings, see “Adding a Geographic Cluster” on page 104.)

You might want to fine-tune the static weights of the geographic cluster's sites to optimize cluster performance. For more information, see “Adjusting a Site's Static Weight” on page 107.

To change a geographic cluster's load-balancing options, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the link formed by the domain name associated with the geographic cluster. The Geographic Cluster Parameters page opens in the right frame.
3. Select Change GeoCluster Parameters from the local menu. The Change Geographic Cluster dialog box appears in the right frame.
4. Select a Load Balancing Method. The load balancing method determines the algorithm that Equalizer will use to distribute requests among the sites in the cluster:
 - **Round Trip**, which weights the client's network proximity more heavily than other criteria.
 - **Adaptive**, which takes all available information into account when selecting a site. This setting is a reasonable default.
 - **Site Load**, which weights the current load at each site more heavily than other criteria.
 - **Site Weight**, which weights the user-defined static weight for each site more heavily than other criteria.
5. Specify the **Load Balancing Response**, which controls how aggressively Equalizer adjusts the site's dynamic weights: Slowest, Slow, Medium, Fast, and Fastest. The faster settings enable Equalizer to adjust its load balancing criteria more frequently and permit a greater variance in the relative weights assigned to sites. A slow setting causes site measurements to be averaged over a longer period of time before Equalizer applies them to the cluster-wide load balancing and tend to ignore spikes in cluster measurements caused by intermittent network glitches.
6. Click the **Set** button.

Deleting a Geographic Cluster

To delete a geographic cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the geographic cluster to be deleted. The Geographic Cluster Parameters frame appears in the right frame.
3. Select **Delete Cluster** from the local menu in the Geographic Cluster Parameters frame.
4. When prompted, click **OK** to verify that you really want to remove the cluster. Equalizer deletes the cluster and all its sites.

Working with Sites

This section describes how to use Equalizer to add or delete a site from a geographic cluster and how to adjust a site's static weight.

Adding a Site to a Geographic Cluster

To add a site to an existing geographic cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the geographic cluster to which you want to add the site.
3. Select **Add Site** from the local menu. The Add Site dialog box opens in the right frame.
4. Enter the **Site Name**, which is a symbolic name that represents this site. For example, the east-coast site for `www.coyotepoint.com` might be `eastCOAST`.
5. Enter the **Site Address**, which is the IP address of the site. This is the address of an Equalizer cluster that is returned if the site is chosen.
6. Click the **Peer** or **Default** radio button. You can designate only one site in a cluster as the default.
 - Equalizer returns a peer site's IP address based on the selected load balancing algorithms.
 - Choose the default site if the client's DNS server did not respond to ICMP echo requests from any site. This can happen if a firewall blocks ICMP packets between the client's DNS and the internet.
7. Enter the **Keepalive** value, which is how often the agent should probe the resource. The default value of 100 results in the resource's availability being tested every 100 seconds.
8. Enter the **Static Weight** value, which represents the site's capacity. (This value is similar to a server's static weight.) Valid values range between 10 and 200. Use the default of 100 if all sites are configured similarly; otherwise, adjust higher or lower for sites that have more or less capacity.
9. Enter the **Resource Address**, which is the IP address of the resource that is monitored for this site. This must be the same address as a configured Equalizer cluster and is generally the same value as the site address. For example, `east.coyotepoint.com` might have resource IP=`192.168.0.5` and Port=`80` if this cluster were configured on Equalizer.
10. Enter the **Port**, which is the TCP port number of the resource that is monitored for this site.
11. Enter the **Agent Address**, which is the IP address of the site monitoring agent. Usually, this is the external (or Envoy failover) address of the Equalizer at this site.
12. Click the **Add** button.

Equalizer can refuse an Add Site command for several reasons, including attempting to add:

- A site with a name or IP address that is already configured
- More sites than are supported by Equalizer
- A default site when you have already configured a default site

Adjusting a Site's Static Weight

Equalizer uses a site's static weight as the starting point for determining what percentage of requests to route to that site. Equalizer assigns sites with a higher static weight a higher percentage of the load. The *relative* values of site static weights are more important than the actual values. For example, if two sites are in a geographic cluster and one has roughly twice the capacity of the other, setting the static weights to 50 and 100 is equivalent to setting the static weights to 100 and 200.

Dynamic site weights can vary from 50% to 150% of the assigned static weights. To optimize geographic cluster performance, you might need to adjust the static weights of the sites in the cluster based on their performance.

Site weights can range from 10 to 200. When you set up sites in a geographic cluster, you should set each site's static weight value in proportion to its capacity for handling requests. It is not necessary for all of the static weights in a cluster to add up to any particular number.

To change a site's static weight, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the site to be modified. The Geographic Site Parameters page opens in the right frame.
3. Select **Change Site Parameters** from the local menu. The Change Site dialog box appears in the right frame.
4. Enter the new weight in the **Static Weight** field.
5. Click the **Set** button.

Deleting a Site from a Geographic Cluster

To delete a site from a geographic cluster, follow these steps:

1. Log into the Equalizer Administration Interface in edit mode.
2. In the left frame, click the name of the site to be deleted. The Site Parameters frame appears in the right frame.
3. Select **Delete Site** from the local menu in the Site Parameters frame.
4. When prompted, click **OK** to confirm that you really want to remove the site.



Introducing Server Agents

You can configure Equalizer's load balancing algorithms to accept direct feedback from servers that describe the current server load or availability of critical resources. To use server agents, you must install agent software on each of the servers in a cluster. These agents must be able to respond to TCP connections on a well-known port. This response is in the form of an ASCII string that represents the current load on the server or indicates that service is not available. If an agent indicates that service is unavailable, Equalizer will automatically stop sending requests to that server.

You configure server agents on a cluster-wide basis—all the servers in a virtual cluster must be running agents for server agents to be used for adaptive load balancing. When you have enabled server agents, Equalizer periodically probes the agent at each server's IP address through the configured agent port. Equalizer uses the collected server agent values when performing adaptive load balancing calculations. You configure Equalizer to use server agents through the Change Cluster dialog box. (For more information, see “Configuring a Cluster to Use Server Agents” on page 73.)

Custom Server Agents


You can write custom agents as shell scripts or in Java, C, or other languages. Upon request, Coyote Point can provide sample agents written in C that you can modify for your specific needs.

Equalizer's agent protocol is extremely simple: when Equalizer connects to the agent's port, the agent must respond with an ASCII string (a number that represents the current condition of the server) and then close the port.

Conditions on the server are as follows:

- 1** Service is unavailable. Might also be used to indicate that a required resource, such as a database, is unavailable.
- 0 to 100** 0 indicates that the server is very lightly loaded; 100 indicates that the server is very overloaded.

B Using Reserved IP Addresses



Equalizer supports placing servers on *reserved*, non-routable networks such as the class A network 10.0.0.0 and the class C network 192.168.2.0. In environments in which the conservation of IP addresses is important, using reserved IP addresses can minimize the number of “real” IP addresses needed.

For example, an ISP hosting several hundred unique web sites replicated on three servers might not want to assign real IP addresses for all of them because each virtual cluster would consume four addresses: three on the back-end servers and one for the virtual cluster. In this case, the ISP might use 10.0.0.0 (the now-defunct Arpanet) as the internal network and assign virtual server addresses out this network for the servers. Figure 67 illustrates a typical reserved internal network.

Note – Due to the additional overhead introduced by enabling outbound NAT, approach using reserved internal networks with caution.

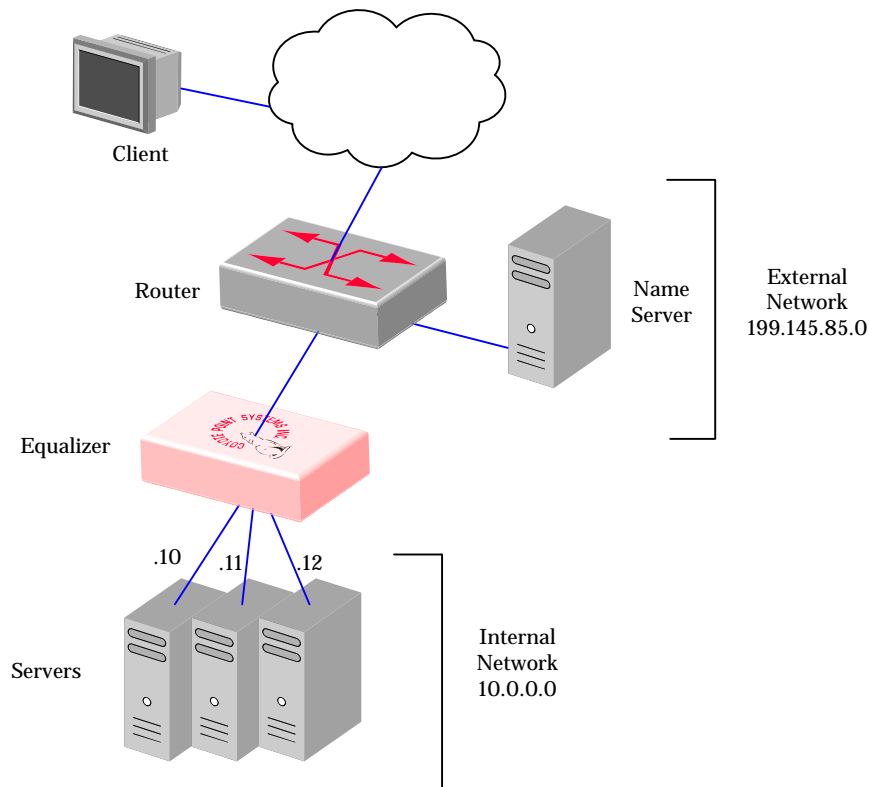


Figure 55 Reserved Internal Network

If servers placed on a non-routable network need to communicate with hosts on the Internet for any reason (such as performing DNS resolution or sending email), you must configure Equalizer to perform *outbound NAT* (network address translation). When you have enabled outbound NAT, Equalizer translates connections originating from the servers on the reserved network so that external hosts will not see packets originating from non-routable addresses.

To enable Equalizer to perform outbound NAT, follow these steps:

1. Open the Equalizer Administration Interface and log in under edit mode.
2. In the left frame, click the **Equalizer** entry at the top of the column.
3. Select **Change Equalizer Parameters** from the local menu. The modify system parameters screen appears in the right frame.
4. Check the **enable outbound NAT**.
5. Click the **commit** button.

Note – If you're using a Failover configuration, outbound NAT must be enabled on both Equalizers.

You will find a worksheet for configuring and using reserved IP addresses in “Equalizer Configuration Worksheets” on page 16.



This appendix describes Equalizer's regular expressions and their format. You can use IEEE Std 1003.2 ("POSIX.2") regular expressions to specify the Equalizer Match Rules. Note that regular expressions are case-insensitive.

Terms

The terms in this section describe the components of regular expressions.

- A *regular expression* (RE) is one or more non-empty branches, separated by pipe symbols (`|`). An expression matches anything that matches one of the branches.
- A *branch* consists of one or more concatenated pieces. A branch matches a match for the first piece, followed by a match for the second, and so on.
- A *piece* is an atom optionally followed by a single `*`, `+`, or `?`, or by a bound.
 - An atom followed by an asterisk matches a sequence of 0 or more matches of the atom.
 - An atom followed by a plus sign matches a sequence of 1 or more matches of the atom.
 - An atom followed by a question mark matches a sequence of 0 or 1 matches of the atom.
- A *bound* consists of an open brace (`{`) followed by an unsigned decimal integer, between 0 and 255 inclusive. You can follow the first unsigned decimal integer with a comma, or a comma and a second unsigned decimal integer. Close the bound with a close brace (`}`). If there are two integers, the value of the first may not exceed the value of the second.

Learning About Atoms

An *atom* followed by a bound that contains one integer i and no comma matches a sequence of exactly i matches of the atom. An atom followed by a bound that contains one integer i and a comma matches a sequence of i or more matches of the atom. An atom followed by a bound containing two integers i and j matches a sequence of i through j (inclusive) matches of the atom. An atom can consist of any of the following:

- A regular expression enclosed in parentheses, which matches a match for the regular expression.
- An empty set of parentheses, which matches the null string.
- A bracket expression.
- A period (`.`), which matches any single character.
- A carat (`^`), which matches the null string at the beginning of a line.
- A dollar sign (`$`), which matches the null string at the end of a line.

- A backslash (\) followed by one of the following characters: ^,[\$()|*+?{\, which matches that character taken as an ordinary character.
- A backslash (\) followed by any other character, which matches that character taken as an ordinary character (as if the \ had not been present).
- A single character with no other significance, which simply matches that character.
- An open brace ({} followed by a character other than a digit is an ordinary character, not the beginning of a bound. It is illegal to end a real expression with a backslash (\).

Creating a Bracket Expression

A *bracket expression* is a list of characters enclosed in brackets ([...]). It normally matches any single character from the list. If the list begins with ^, it matches any single character not from the rest of the list. Two characters in a list that are separated by '-' indicates the full range of characters between those two (inclusive) in the collating sequence; for example, '[0-9]' in ASCII matches any decimal digit. It is illegal for two ranges to share an endpoint; for example, 'a-c-e'. Ranges are very collating-sequence-dependent, and portable programs should avoid relying on them.

- To include a literal ']' in the list, make it the first character (following an optional '^').
- To include a literal '-', make it the first or last character, or the second endpoint of a range.
- To use a literal '-' as the first endpoint of a range, enclose it in '[' and ']' to make it a collating element (see below).

With the exception of these and some combinations using '[' (see next paragraphs), all other special characters, including '\', lose their special significance within a bracket expression.

Within a bracket expression, a collating element (a character, a multi-character sequence that collates as if it were a single character, or a collating-sequence name for either) enclosed in '[' and ']' stands for the sequence of characters of that collating element. The sequence is a single element of the bracket expression's list. A bracket expression containing a multi-character collating element can thus match more than one character; e.g., if the collating sequence includes a 'ch' collating element, then the real expression '[[.ch.]]*c' matches the first five characters of 'chchcc'.

Within a bracket expression, a collating element enclosed in '[' and `]' is an equivalence class, representing the sequences of characters of all collating elements equivalent to that one, including itself. (If there are no other equivalent collating elements, the treatment is as if the enclosing delimiters were '[' and '].') For example, if 'x' and 'y' are the members of an equivalence class, then '[[x]]', '[[y]]', and '[xy]' are all synonymous. An equivalence class may not be an end-point of a range.

Within a bracket expression, the name of a character class enclosed in '[' and ':' stands for the list of all characters belonging to that class.

There are two special cases of bracket expressions: the bracket expressions '[:<:]' and '[:>:]' match the null string at the beginning and end of a word respectively. A word is defined as a sequence of word characters that is neither preceded nor followed by word characters. A word character is an alnum character (as defined by ctype(3)) or an underscore. This is an extension, compatible with but not specified by IEEE Std 1003.2 ("POSIX.2"), and should be used with caution in software intended to be portable to other systems.

Matching Expressions

If a real expression could match more than one substring of a given string, the real expression matches the one starting earliest in the string. If the real expression could match more than one substring starting at that point, it matches the longest. Subexpressions also match the longest possible substrings, subject to the constraint that the whole match be as long as possible, with subexpressions starting earlier in the real expression taking priority over ones starting later. Note that higher-level subexpressions thus take priority over their lower-level component subexpressions.

Match lengths are measured in characters, not collating elements. A null string is considered longer than no match at all. For example, 'bb*' matches the three middle characters of 'abbbc', '(weelweek)(knights|nights)' matches all ten characters of 'weeknights', when '(.*).*' is matched against 'abc' the parenthesized subexpression matches all three characters, and when '(a*)*' is matched against 'bc' both the whole real expression and the parenthesized subexpression match the null string.



Instead of supporting sendmail, Equalizer supports a program called `mini_sendmail`, which is stored in the `/usr/local/sbin` folder and which accepts e-mail on behalf of the standard sendmail program. The `mini_sendmail` program is authored and copyrighted (C) 1999 by Jef Poskanzer (jef@acme.com). All rights are reserved.

Syntax

The syntax of `mini_sendmail` is as follows:

```
mini_sendmail [-f<name>] [-t] [-s<server>] [-T<timeout>] [-v]
[address ...]
```

When you issue the `mini_sendmail` command without flags, `mini_sendmail` reads its standard input up to an end-of-file and sends a copy of the message found there to all the addresses listed. The `mini_sendmail` program sends the message by connecting to a local SMTP server, which means you can use `mini_sendmail` to send e-mail from inside a `chroot(2)` area.

Flags

The `mini_sendmail` program supports the following flags:

- `-f`, which sets the name of the sender of the mail
- `-t`, which reads the message for recipients. The `mini_sendmail` program scans the `To:`, `Cc:`, and `Bcc:` lines for recipient addresses, and deletes the `Bcc:` line before transmission.
- `-s`, which specifies the SMTP server to use. Without this flag, `mini_sendmail` uses `local-host`.
- `-T`, which specifies the timeout. This value defaults to one minute.
- `-v`, which indicates verbose mode (that is, the conversation with the SMTP server appears).

Example

The following shows a sample of `mini_sendmail` usage:

```
/usr/local/sbin/mini_sendmail -f equalizer -s
your_local_SMTP_server_IP
```




You usually can diagnose Equalizer installation and configuration problems using standard network troubleshooting techniques. This section identifies some common problems, the most likely causes, and the best solutions.

Equalizer Doesn't Boot

Serial cable not connected to Equalizer

Check the cable connection.

Clients Time Out While Trying to Contact a Virtual Cluster

Equalizer is not gatewaying reply packets from the server

Log on to the server(s) and check the routing tables. Perform a `traceroute` from the server to the client. Adjust the routing until Equalizer's address shows up in the `traceroute` output.



All packets sent from the server back to clients must pass through Equalizer.

Test client is on the same network as the servers

If the test client is on the same network as the servers, the servers will probably try to send data packets directly to the client, bypassing Equalizer. You can correct this by adding *host routes* on the servers so that the servers send their reply packets via Equalizer.

No active servers in the virtual cluster

Possible solutions:

- Check the Equalizer Summary page. Are there any servers in that virtual cluster? Are all the servers marked DOWN?
- Log onto the server and run the `netstat` command (Unix servers). If the `netstat` output shows connections in the SYN-RCVD state, the server is not forwarding its reply packets to Equalizer.

Equalizer is not active

Is Equalizer functioning? Try to `ping` the administration address. If you do not get a response, “Equalizer Doesn't Respond to Pings to the Admin Address” provides additional troubleshooting information.

Primary and Backup Equalizer Are in a Conflict Over Primary

Certain switches (often those from Cisco and Dell) have Spanning Tree enabled by default. This can cause a delay in the times that the network is accessible and cause the backup Equalizer to enter into failover mode. If you cannot disable Spanning Tree, enable FastPort for all ports connected to the Equalizers.

Backup Equalizer Continues to Boot

Primary and Backup Equalizer Are in a Conflict over Primary

Certain Dell and Cisco switches have Spanning Tree enabled by default. This can cause a delay in the times that the network is accessible and cause the backup Equalizer to enter into failover mode. If you cannot disable Spanning Tree, enable PortFast for all ports connected to the Equalizers.

Can't View Equalizer Administration Pages

Equalizer is not active

Is Equalizer functioning? Try to ping the administration address. If you do not get a response, see “Equalizer doesn't respond to pings to the admin address” below, which provides additional troubleshooting information.

Equalizer Administration Page Takes a Long Time to Display

DNS server configured on Equalizer is not responding

Possible solutions:

- Check that Equalizer has IP connectivity to the name server configured using the serial configuration utility.
- If you want to disable DNS lookups on Equalizer, specify a name server IP address of 0.0.0.0 in Equalizer's serial configuration utility.

Equalizer Doesn't Respond to Pings to the Admin Address

Equalizer is not powered on

Check that power switch is on and the front panel LED is lit. Connect the keyboard and monitor, cycle the power, and watch the startup diagnostic messages.

Equalizer isn't connected to your network

Check the network wiring.

Administration address not configured on the external interface

This applies to dual network configurations. Use the Equalizer Configuration Utility to set the IP address and netmask for external interface. Be sure to commit your changes.

Browser Hangs When Trying to Connect Via FTP to an FTP Cluster

FTP server returns its private IP address in response to a “PASV” command

Enable PASV mode FTP translation on the Advanced options page of the Equalizer Administration utility. (For more information, see “Enabling Passive FTP Connections” on page 41.) This behavior is likely to cause problems if you’re using reserved internal addresses for the server

Return Packets from the Server Aren’t Routing Correctly

IP spoofing is enabled

This problem normally occurs in a single network setup. When you enable IP spoofing, clustered servers see the client’s IP address. If the server tries to reply directly to the client, the client will reject the reply (it had sent its request to a different address).

Run a `traceroute` to ensure that routes from a server to a client go through Equalizer and not directly back to the client. If Equalizer does not appear, modify the route to include Equalizer. Alternatively, you can disable IP spoofing.

Web Server Cannot Tell Whether Incoming Requests Originate Externally or Internally

IP Spoofing is not enabled

Check the cluster’s configuration and enable IP spoofing. This causes Equalizer to pass the client’s IP address. Make sure that responses from the server go through the Equalizer.



SOFTWARE LICENSE

PLEASE READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THE SOFTWARE. BY USING THIS SOFTWARE YOU AGREE TO BE BOUND BY THE TERMS AND CONDITIONS OF THIS LICENSE. IF YOU DO NOT AGREE WITH THE TERMS OF THIS LICENSE, PROMPTLY RETURN THE UNUSED SOFTWARE, MANUAL, AND RELATED EQUIPMENT AND HARDWARE (WITH PROOF OF PAYMENT) TO THE PLACE OF PURCHASE FOR A FULL REFUND.

Coyote Point Systems, Inc. (“Coyote Point Systems”) and its suppliers grant to Customer (“Customer”) a nonexclusive and nontransferable license to use the Coyote Point Systems software (“Software”) in object code form solely on a single central processing unit owned or leased by Customer or otherwise embedded in equipment provided by Coyote Point Systems. Customer may make one (1) archival copy of the software provided Customer affixes to such copy all copyright, confidentiality, and proprietary notices that appear on the original. EXCEPT AS EXPRESSLY AUTHORIZED ABOVE, CUSTOMER SHALL NOT COPY, IN WHOLE OR IN PART, SOFTWARE OR DOCUMENTATION; MODIFY THE SOFTWARE; REVERSE COMPILE OR REVERSE ASSEMBLE ALL OR ANY PORTION OF THE SOFTWARE; OR RENT, LEASE, DISTRIBUTE, SELL, OR CREATE DERIVATIVE WORKS OF THE SOFTWARE.

Customer agrees that aspects of the licensed materials, including the specific design and structure of individual programs, constitute trade secrets and/or copyrighted material of Coyote Point Systems. Customer agrees not to disclose, provide, or otherwise make available such trade secrets or copyrighted material in any form to any third party without the prior written consent of Coyote Point Systems. Customer agrees to implement reasonable security measures to protect such trade secrets and copyrighted material. Title to Software and documentation shall remain solely with Coyote Point Systems.

This License is effective until terminated. Customer may terminate this License at any time by destroying all copies of Software including any documentation. This License will terminate immediately without notice from Coyote Point Systems if Customer fails to comply with any provision of this License. Upon termination, Customer must destroy all copies of Software. Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licenses to export, reexport, or import Software.

This License shall be governed by and construed in accordance with the laws of the State of New York, United States of America, as if performed wholly within the state and without giving effect to the principles of conflict of law. If any portion hereof is found to be void or unenforceable, the remaining provisions of this License shall remain in full force and effect. This License constitutes the entire License between the parties with respect to the use of the Software.

Restricted Rights - Coyote Point Systems' software is provided to non-DOD agencies with RESTRICTED RIGHTS and its supporting documentation is provided with LIMITED RIGHTS. Use, duplication, or disclosure by the U.S. Government is subject to the restrictions as set forth in

subparagraph “C” of the Commercial Computer Software - Restricted Rights clause at FAR 52.227-19. In the event the sale is to a DOD agency, the U.S. Government's rights in software, supporting documentation, and technical data are governed by the restrictions in the Technical Data Commercial Items clause at

DFARS

252.227-7015 and DFARS 227.7202.

LIMITED WARRANTY

This document includes Limited Warranty information for Coyote Point Systems products. For products purchased in the European Union, please refer to the European Union Amendment.

General Terms.

EXCEPT AS EXPRESSLY SET FORTH IN THIS LIMITED WARRANTY, COYOTE POINT SYSTEMS MAKES NO OTHER WARRANTIES OR CONDITIONS, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. COYOTE POINT SYSTEMS EXPRESSLY DISCLAIMS ALL WARRANTIES AND CONDITIONS NOT STATED IN THIS LIMITED WARRANTY. ANY IMPLIED WARRANTIES THAT MAY BE IMPOSED BY LAW ARE LIMITED IN DURATION TO THE LIMITED WARRANTY PERIOD. SOME STATES OR COUNTRIES DO NOT ALLOW A LIMITATION ON HOW LONG AN IMPLIED WARRANTY LASTS OR THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES FOR CONSUMER PRODUCTS. IN SUCH STATES OR COUNTRIES, SOME EXCLUSIONS OR LIMITATIONS OF THIS LIMITED WARRANTY MAY NOT APPLY TO YOU.

This Limited Warranty applies to the Coyote Point Systems software and hardware products sold by Coyote Point Systems, Inc., its subsidiaries, affiliates, authorized resellers, or country distributors (collectively referred to in this Limited Warranty as "Coyote Point Systems") with this Limited Warranty.

Software. Coyote Point Systems warrants that: (i) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (ii) the Software substantially conforms to its published specifications during the Limited Warranty Period. Except for the foregoing, the Software is provided AS IS.

Hardware. Coyote Point Systems warrants that the Hardware product will be free from defects in material and workmanship under normal use during the Limited Warranty Period.

The Limited Warranty Period is for one year from the date of shipment. Your dated delivery receipt, showing the date of shipment of the product, is your proof of the shipment date. You may be required to provide proof of purchase as a condition of receiving warranty service. You are entitled to warranty service according to the terms and conditions of this document if a repair to your Coyote Point Systems software or hardware is required within the Limited Warranty Period. This Limited Warranty extends only to the original purchaser of this Coyote Point Systems product and is not transferable to anyone who obtains ownership of the Coyote Point Systems product from the original purchaser.

Coyote Point Systems products are manufactured using new materials or new and used materials equivalent to new in performance and reliability. Replacement products are guaranteed to have functionality at least equal to our published specifications. Replacement parts are warranted to be

free from defects in material or workmanship for the remainder of the Limited Warranty Period. Repair or replacement of a part will not extend the Limited Warranty.

During the Limited Warranty Period, Coyote Point Systems will repair or replace the defective component parts or the hardware product. All component parts or hardware products removed under this Limited Warranty become the property of Coyote Point Systems. Coyote Point Systems, at its discretion, may elect to provide you with a replacement unit of Coyote Point Systems' choosing that is at least equivalent to your Coyote Point Systems product in hardware performance. Coyote Point Systems reserves the right to elect, at its sole discretion, to give you a refund of your purchase price instead of a replacement. This is your exclusive remedy for defective products.

To request Limited Warranty service, you must contact Coyote Point Systems Technical Support, which can be reached at (888) 891-8150 or via E-mail at support@coyotepoint.com. Coyote Point Systems Technical Support will determine the nature of the problem, and if a return is necessary, issue a Return Materials Authorization (RMA). No returned product will be accepted without an RMA number obtained in advance and clearly marked on the outside of the shipping container. All products to be returned must be in the original manufacturer's undamaged packaging along with all accessories shipped with the original product including cables, handles and manuals. If you did not retain the original packaging materials, there may be a charge for replacement packaging.

If a defective product is returned, the cost of incoming freight and insurance is the responsibility of the customer. The cost of return freight is the responsibility of Coyote Point Systems, if shipped within the United States. Shipments to other locations will be freight collect. You are responsible for missing or physically damaged parts on the returned defective product, if they are not covered under the product Limited Warranty. You are responsible for all customs fees, taxes or VAT that may be due (excluding income taxes). A product returned for repair but found to be in good working order will be charged a \$75 "No Trouble Fee".

Coyote Point Systems does not warrant that the operation of this product will be uninterrupted or error-free. Coyote Point Systems is not responsible for damage that occurs as a result of your failure to follow the instructions that came with the Coyote Point Systems product.

Restrictions.

This Limited Warranty does not extend to software errors that can not be reproduced, or for any product from which the serial number has been removed, or that has been damaged or rendered defective (a) as a result of accident, misuse, abuse, or other external causes; (b) by operation outside the usage parameters stated in the user documentation that shipped with the product; (c) by the use of parts not manufactured or sold by Coyote Point Systems; or (d) by modification or service by anyone other than (i) Coyote Point Systems, (ii) a Coyote Point Systems authorized service provider, or (iii) your own installation of end-user replaceable Coyote Point Systems or Coyote Point Systems approved parts.

COYOTE POINT SYSTEMS WILL NOT HAVE ANY LIABILITY FOR ANY DAMAGES ARISING FROM THE USE OF THE PRODUCTS IN ANY HIGH-RISK ACTIVITY, INCLUDING, BUT NOT LIMITED TO, THE OPERATION OF NUCLEAR FACILITIES, AIRCRAFT NAVIGATION OR COMMUNICATION SYSTEMS, AIR TRAFFIC CONTROL, MEDICAL SYSTEMS, LIFE SUPPORT, OR WEAPONS SYSTEMS.

These terms and conditions constitute the complete and exclusive warranty agreement between you and Coyote Point Systems regarding the Coyote Point Systems product you have purchased. These terms and conditions supersede any prior agreements or representations including representations made in Coyote Point Systems sales literature or advice given to you by Coyote Point Systems or an agent or employee of Coyote Point Systems that may have been made in connection with your purchase of the Coyote Point Systems product. No change to the conditions of this Limited

Warranty is valid unless it is made in writing and signed by an authorized representative of Coyote Point Systems.

Limitation of Liability

IF YOUR COYOTE POINT SYSTEMS SOFTWARE OR HARDWARE PRODUCT FAILS TO WORK AS WARRANTED ABOVE, YOUR SOLE AND EXCLUSIVE REMEDY SHALL BE REPAIR OR REPLACEMENT (INCLUDING REFUND). COYOTE POINT SYSTEMS' MAXIMUM LIABILITY UNDER THIS LIMITED WARRANTY IS EXPRESSLY LIMITED TO THE LESSER OF THE PRICE YOU HAVE PAID FOR THE PRODUCT OR THE COST OF REPAIR OR REPLACEMENT OF ANY SOFTWARE OR HARDWARE COMPONENTS THAT MALFUNCTION IN CONDITIONS OF NORMAL USE. COYOTE POINT SYSTEMS IS NOT LIABLE FOR ANY DAMAGES CAUSED BY THE PRODUCT OR THE FAILURE OF THE PRODUCT TO PERFORM, INCLUDING ANY LOST PROFITS OR SAVINGS OR DATA, OR SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, OR PUNITIVE DAMAGES. COYOTE POINT SYSTEMS IS NOT LIABLE FOR ANY CLAIM MADE BY A THIRD PARTY OR MADE BY YOU FOR A THIRD PARTY.

THIS LIMITATION OF LIABILITY APPLIES WHETHER DAMAGES ARE SOUGHT, OR A CLAIM MADE, UNDER THIS LIMITED WARRANTY OR AS A TORT CLAIM (INCLUDING NEGLIGENCE AND STRICT PRODUCT LIABILITY), A CONTRACT CLAIM, OR ANY OTHER CLAIM. THIS LIMITATION OF LIABILITY CANNOT BE WAIVED OR AMENDED BY ANY PERSON. THIS LIMITATION OF LIABILITY WILL BE EFFECTIVE EVEN IF YOU HAVE ADVISED COYOTE POINT SYSTEMS OR AN AUTHORIZED REPRESENTATIVE OF COYOTE POINT SYSTEMS OF THE POSSIBILITY OF ANY SUCH DAMAGES. THIS LIMITATION OF LIABILITY, HOWEVER, WILL NOT APPLY TO CLAIMS FOR PERSONAL INJURY. THE FOREGOING LIMITATIONS SHALL APPLY EVEN IF THE ABOVE-STATED WARRANTY FAILS OF ITS ESSENTIAL PURPOSE.

THIS LIMITED WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS. YOU MAY ALSO HAVE OTHER RIGHTS THAT MAY VARY FROM STATE TO STATE OR FROM COUNTRY TO COUNTRY. YOU ARE ADVISED TO CONSULT APPLICABLE STATE OR COUNTRY LAWS FOR A FULL DETERMINATION OF YOUR RIGHTS.

IN THE EVENT OF INCONSISTENCY BETWEEN ANY TERMS OF THIS DISCLAIMER OF WARRANTIES AND LIMITED WARRANTY AND ANY TRANSLATION THEREOF INTO ANOTHER LANGUAGE, THE ENGLISH LANGUAGE VERSION SHALL PREVAIL.

THIS DISCLAIMER OF WARRANTIES AND LIMITED WARRANTY ARE GOVERNED BY THE LAWS OF THE STATE OF NEW YORK, UNITED STATES OF AMERICA, WITHOUT REGARD TO THE CONFLICT OF LAWS PROVISIONS THEREOF. THE UNITED NATIONS CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALE OF GOODS SHALL NOT APPLY TO THESE TERMS IN ANY RESPECT.

THIS DISCLAIMER OF WARRANTIES AND LIMITED WARRANTY ARE SUBJECT TO THE TERMS OF SALE OF THE COYOTE POINT SYSTEMS' PRODUCT.



This glossary defines some of the key terms used in this document. Some of the glossary definitions are based on RFC1208, “A Glossary of Networking Terms.”¹

ACV	Active Content Verification; an Equalizer mechanism for checking the validity of a server. ACV does not support UDP-based services.
administration address	The IP address assigned to Equalizer on the internal network. See internal network and IP address.
administration interface	The browser-based interface for setting up and managing the operation of Equalizer.
address translation	The modification of external addresses to standardized network addresses and of standardized network addresses to external addresses.
agent	An application that gathers or processes information for a larger application. See server agent.
aggregation	A summary of all the data that is computed from detailed information. See sticky network aggregation.
alias	A nickname that replaces a long name or one that is difficult to remember or spell. See IP alias.
aliased IP address	A nickname for an IP address. See IP alias.
algorithm	Instructions, procedures, or formulas used to solve a problem.
application layer	Layer 7 (L7); the highest layer of standards in the Open Systems Interconnection (OSI) model (according to The Microsoft Press <i>Computer Dictionary</i>), which helps a user perform work such as transferring files, formatting e-mail messages, and accessing remote computers.
atom	The smallest part of a regular expression in Equalizer. See branch, piece, and regular expression.
authoritative name server	A name server that maintains the complete information for a particular part of the domain name space. See name server.
back-end server	A physical server on the internal network that receives connection requests from Equalizer.
backup Equalizer	The backup unit, which replaces the primary Equalizer if that Equalizer fails. See hot backup and primary Equalizer.
bound	A character that represents the limit of part of a regular expression.
bracket expression	In a regular expression, a list of characters enclosed in brackets ([...]).
branch	In an Equalizer regular expression, a complete piece of a regular expression. You can concatenate and/or match branches. See atom, piece, and regular expression.
cache	An area in which information is temporarily stored.
Class A	An ISO/IEC 11801 standard for twisted pair cabling rated to 100 KHz; similar to Category 1 cabling. Use the Class A standard for voice and low frequency applications. According to the Microsoft Press <i>Computer Dictionary</i> , you can use Class A networks “for sites with few networks but numerous hosts.” See ISO/IEC.
Class B	An ISO/IEC 11801 standard for twisted pair cabling rated to 1 MHz; similar to Category 2 cabling. Use the Class B standard for medium bit rate applications. See ISO/IEC.

Class C	An ISO/IEC 11801 standard for twisted pair cabling rated to 16 MHz; similar to TIA/EIA Category 3 cabling. Use the Class C standard for high bit rate applications, in which the network allocates 24 bits for the IP address network-address field. A Class C network allocates 24 bits for the IP address network-address field and 8 bits for the host field. <i>See</i> ISO/IEC.
cluster	A set of networked computer systems that work together as one system. <i>See</i> server cluster and virtual cluster.
cluster address	The IP address assigned to a particular cluster configured on Equalizer.
computed load	A measure of the performance of a server relative to the overall performance of the cluster of which the server is a part.
cookie	Data that a Web server stores on a client on behalf of a Web site. When a user returns to the Web site, the server reads the cookie data on the client, providing the Web site all the saved information about the user.
daemon	An application that runs in the background and performs one or more actions when events trigger those actions.
DNS	Domain Name System or Domain Name Service; used to map domain names to Internet servers in order to link to IP addresses or map IP addresses to domain names. <i>See</i> IP address.
DNS TTL	The amount of time, in seconds, that a name server is allowed to cache the domain information. <i>See</i> DNS and TTL.
domain	The highest level in an IP address and the last part of the address in the URL. The domain identifies the category under which the Web site operates. For example, in <code>www.coyotepoint.com</code> , <code>com</code> is the domain, where <code>com</code> represents a <i>commercial</i> site. <i>See</i> domain name, IP address, and subdomain. <i>See also</i> DNS.
domain name	The owner of an IP address. The next highest level in an IP address and the next-to-last part of the address. For example, in <code>www.coyotepoint.com</code> , <code>coyotepoint</code> is the domain name. <i>See</i> domain, IP address, and subdomain. <i>See also</i> DNS.
dynamic weight	The weight that Equalizer assigns to a particular server during operation. <i>See</i> server weight, static weight, and weight.
echo	The transmittal of data that has been sent successfully back to the originating computer. <i>See</i> ping. <i>See also</i> CMP echo request.
edit mode	One of two modes in which Equalizer can be administered. In edit mode, you can view and modify parameters. <i>See</i> view mode.
EIA	Electronic Industries Association; a trade association that sets standards for electrical and electronic components.
endpoint	An IP address-port pair that identifies the start or end of an address; a value that ends a process.
Envoy	Equalizer add-in; software that supports geographic clustering and load balancing. <i>See</i> geographic cluster, geographic load balancing, and load balancing. <i>See also</i> intelligent load balancing.
Equalizer Administration Interface	An Equalizer window with which you can monitor Equalizer's operation; view statistics; add, modify, or clusters; add, modify, and delete servers; and shut down a server or Equalizer through a Javascript-enabled browser.
Equalizer Configuration Utility	An Equalizer feature that enables you to configure Equalizer, set parameters, and shut down and upgrade Equalizer.
external address	The IP address assigned to Equalizer on the external network.
external interface	A network interface used to connect Equalizer to the external network. <i>See</i> interface, internal interface, and network interface.

1. O. Jacobsen and D. Lynch, Interop, Inc. March 1991.

external network	The subnet to which the client machines and possibly the Internet or an intranet are connected.
failover	The act of transferring operations from a failing component to a backup component without interrupting processing.
firewall	A set of security programs, which is located at a network gateway server and which protect the network from any user on an external network. <i>See gateway.</i>
FQDN	<i>See Fully Qualified Domain Name (FQDN).</i>
FTP	File Transfer Protocol; rules for transferring files from one computer to another.
FTP cluster	A virtual cluster providing service on the FTP control port (port 21). <i>See cluster and virtual cluster.</i>
Fully Qualified Domain Name (FQDN)	The complete, registered domain name of an Internet host, which is written relative to the root domain and unambiguously specifies a host's location in the DNS hierarchy. For example, <code>east</code> is a hostname and <code>east.coyotepoint.com</code> is its fully qualified domain name. <i>See also domain name.</i>
gateway	A network route that typically translates information between two different protocols.
geographic cluster	A collection of servers (such as Web sites) that provide a common service over different physical locations. <i>See cluster.</i>
geographic load balancing	Distributing requests as equally as possible across servers in different physical locations. <i>See load balancing. See also intelligent load balancing.</i>
geographic probe	A query sent to a site in a geographic cluster to gather information so Equalizer can determine the site that is best able to process a pending request. <i>See geographic cluster.</i>
header	One or more lines of data that identify the beginning of a block of information or a file.
hot backup	Configuring a second Equalizer as a backup unit that will take over in case of failure. Also known as a hot spare. <i>See backup Equalizer. See also primary Equalizer.</i>
HTTP	HyperText Transfer Protocol; the protocol with which a computer or user access information on the World Wide Web.
HTTPS	HyperText Transfer Protocol (Secure); a server application programmed to run under the Windows NT operating system.
hub	A device that joins all the components attached to a network.
ICMP	<i>See Internet Control Message Protocol.</i>
ICMP echo request	The act of repeating a stream of characters (for example, echoing on the computer screen characters as a user types those characters). <i>See ping. See also echo.</i>
ICMP triangulation	Routing client requests to the closest site geographically based on triangulation, a method of calculating the location of a site using the known locations of two or more other sites.
intelligent load balancing	A request for load balancing using Equalizer-based algorithms that assess the configuration options set for cluster and servers, real-time server status information, and information in the request itself. <i>See algorithm and load balancing. See also geographic load balancing.</i>
interface	The place at which two or more systems connect and communicate with each other. <i>See external interface, internal interface, and network interface.</i>
internal address	The IP address assigned to Equalizer on the external network.
internal network	The subnet to which the back-end server machines are connected.
Internet Control Message Protocol (ICMP)	The ISO/OSI Layer 3, Network, protocol that controls transport routes, message handling, and message transfers during IP packet processing. <i>See ICMP triangulation and ISO/OSI model.</i>
IP	Internet protocol; the TCP/IP protocol that controls breaking up data messages into packets, sending the packets, and reforming the packets into their original data messages. <i>See Internet protocol stack, IP address, packet, and TCP/IP.</i>

IP address	A 32-bit address assigned to a host using TCP/IP. IP addresses are written in dotted decimal format, for example, 192 . 22 . 33 . 1.
ISO/IEC	International Organization for Standardization/International Electrotechnical Commission; international standards organizations.
ISO/OSI model	<p>International Organization for Standardization/Open Systems Interconnection model, a standard that consists of seven layers that control how computers communicate with other computers over a network.</p> <ul style="list-style-type: none">▪ Layer 1, Physical, which sets the rules for physical connections via hardware, is the lowest layer.▪ Layer 2, Data-link, uses Layer 1 and its own rules to control coding, addressing, and transmitting information.▪ Layer 3, Network, uses the prior two layers rules as well as its own rules to control transport routes, message handling, and message transfers.▪ Layer 4, Transport, uses its rules and those of the previous layers to control accuracy of message delivery and service.▪ Layer 5, Session, uses its rules and those of the previous layers to establish, maintain, and coordinate communication.▪ Layer 6, Presentation, uses its rules and those of the previous layers to control text formatting and appearance as well as conversion of code.▪ Layer 7, Application, uses its rules and those of the other layers to control transmission of information from one application to another. Layer 7 is the highest layer. <p>See Layer 4, Layer 7, and transport layer.</p>
L4	See Layer 4.
L7	See Layer 7.
latency	The time over which a signal travels over a network, from the starting point to the endpoint. See ping. See also CMP echo request and echo.
Layer 4 (L4)	The transport layer; Layer 4 uses its rules and those of the previous three layers to control accuracy of message delivery and service.which controls accuracy of message delivery and service. See ISO/OSI model and Layer 7.
Layer 7 (L7)	The application layer; Layer 7 uses its rules and those of the other layers to control transmission of information from one application to another. Layer 7 is the highest layer in the ISO/OSI model. See ISO/OSI model and Layer 4.
load	A job that can be processed or transported once. See load balancing. See also geographic load balancing and intelligent load balancing.
load balancing	Moving a load from a highly-used resource to a resource that is used less often so that operations are efficient. Equalizer balances loads over a wide physical area or by using algorithms that assess options and real-time information. See geographic load balancing and intelligent load balancing.
MX exchanger	Mail exchanger; a fully qualified domain name to be returned if a server receives a mail exchanger request.
name server	A server that stores information about the domain name space.
NAT	Network Address Translation; an Internet standard that defines the process of converting IP addresses on a local-area network to Internet IP addresses. See NAT subsystem.
NAT subsystem	The Equalizer subsystem responsible for transferring connections to and from the back-end servers.

netmask	Address mask; a bit mask used to select bits from an Internet address for subnet addressing. The mask is 32 bits long and selects the network portion of the Internet address and one or more bits of the local portion.
Network Address Translation (NAT)	See NAT.
network interface	The place at which two or more networks connect and communicate with each other. See interface. See external interface, interface, and internal interface.
network route	See gateway.
OSI network	A network that uses the International Organization for Standardization/Open Systems Interconnection model. See ISO/OSI model, Layer 4, Layer 7, and transport layer.
packet	A group of data that is transmitted as a single entity.
passive FTP connection	An Equalizer option that rewrites outgoing FTP PASV control messages from the servers so that they contain the IP address of the virtual cluster rather than that of the server. See FTP and PASV.
PASV	Passive mode FTP; a mode with which you can establish FTP connections for clients that are behind firewalls. See firewall, FTP, and passive FTP connections.
pattern match	A pattern of ASCII or hexadecimal data that filters data.
payload	The set of data to be transmitted. A payload contains user information, user overhead information, and other information that a user requests. A payload <i>does not</i> include system overhead information. Also known as the mission bit stream.
persistence	The act of storing or retaining data for use at a later time, especially data that shows the state of the network before processing resumes. See cookie and IP-address-based persistence.
physical server	A machine located on the internal network that provides services on specific IP addresses and ports. See server and virtual web server. See also authoritative name server, back-end server, name server, and proxy server.
piece	An atom followed by a single *, +, or ?, or by a bound. See atom, branch, and regular expression.
ping	A program used to test reachability of destinations by sending them an ICMP echo request and waiting for a reply. See echo and probe. See also CMP echo request
port	The abstraction used by Internet transport protocols to distinguish among multiple simultaneous connections to a single destination host.
port number	The number used to identify a service contact port, such as HTTP port 80.
primary Equalizer	The primary unit that handles requests. If the primary Equalizer fails, the backup unit replaces it. See also backup Equalizer and hot backup.
probe	An action that obtains status information about a computer, network, or device. See geographic probe and ping.
protocol	A set of rules that govern adherence to a set of standards. See protocol stack.
protocol stack	A layer of protocols that process network actions cooperatively and in tandem. See protocol.
proxy server	A utility, which is part of a firewall, that helps the regular tasks of managing data transmittal from a network to the Internet and from the Internet to the network. See also firewall.
quiesce	To become quiet or more quiet than previously.
RADIUS	Remote Authentication Dial-In User Service; a protocol that authorizes and authenticates a user trying to link to a network or the Internet.
redirection	The process of receiving input from or sending output to a different resource than usual.
regular expression (RE)	One or more non-empty branches, separated by pipe symbols (). An expression matches anything that matches one of the branches. See atom, branch and piece.

request packet	A packet that contains information that requests a response. <i>See</i> packet and response packet.
reserved network	A network consisting of “phony” IP addresses, which are not registered and cannot be made visible outside of the internal network.
resolution	The process of interpreting all the messages between an IP address and a domain name address.
response packet	A packet that contains information that responds to a request. <i>See</i> packet and request packet.
router	A network device that facilitates the transmission (that is, <i>routing</i>) of messages.
routing table	A database, which is static or dynamic, that contains a set of route addresses and routing information familiar to the router. A human being enters and updates the information in a static routing table; routers operate and constantly update a dynamic routing table.
RST	The reset command, which instructs a device to end a connection.
Secure Sockets Layer (SSL)	A protocol, which uses public-key encryption, that enables secure communications between a client and Web server, typically for guarding financial transactions.
server	A computer or application that controls access to a network and its associated devices and applications. A server communicates with one or more clients as well as other servers. <i>See</i> authoritative name server, back-end server, name server, physical server, proxy server, and virtual web server.
server address	The IP address of a server on the internal interface. Multiple IP addresses can be aliased to a single physical server. <i>See</i> server.
server agent	An agent that provides Equalizer with real-time performance statistics for a specified server. <i>See</i> server.
server cluster	A group of servers that are components in a network and joined through hardware or software. <i>See</i> cluster. <i>See also</i> FTP cluster, geographic cluster, and virtual cluster. <i>See</i> server.
server endpoint	An IP address-port pair that identifies a physical or virtual server on the internal network to which Equalizer can route connection requests. <i>See</i> server.
server weight	A value that indicates the relative proportion of connection requests that a particular server will receive. <i>See</i> dynamic weight, server, static weight, and weight.
site	A cluster of servers under Equalizer control that is part of a geographic cluster.
spoofing	Fooling a system into thinking that a transmission comes from an authorized user when that may not be the case.
SSL	<i>See</i> Secure Sockets Layer (SSL).
stack	An area of reserved memory in which applications place status data and other data. <i>See</i> protocol stack.
stale connection	A partially open or closed connection.
state	Status; the current condition of a network, computer, or peripherals.
stateless	A condition in which a server processes each request from a site independently and cannot store information about prior requests from that site. Each request stands on its own. <i>See also</i> DNS and RADIUS.
static weight	The weight that an administrator assigns to a particular server. During operation, Equalizer dynamically adjusts the server weights (that is, dynamic weight), so a server's weight at a particular time might be different from the static weight originally set by the administrator. <i>See</i> dynamic weight, server weight, and weight.
sticky connection	A connection in which a particular client remains connected to same server to handle subsequent requests within a set period of time.
sticky timer	A countdown timer that tracks periods of inactivity between a particular client and server.

subdomain	A section, which is formally named, that is under a domain name; analogous to the relationship between a subfolder and folder. For example, in <code>www.coyotepoint.com</code> , <code>www</code> is the subdomain. <i>See</i> domain, domain name, and IP address. <i>See also</i> DNS.
subnet	Part of a network that has the same address as the network plus a unique subnet mask.
switch	A device, which is attached to a network and which controls the route over which data is sent.
SYN/ACK	Synchronize and acknowledge; a message that synchronizes a sequence of data information and acknowledges the reception of that information.
syslog	A system log file, in which information, warning, and error messages are stored in a file, sent to a system, or printed.
TCP	Transmission Control Protocol; the rules for the conversion of data messages into packets. <i>See</i> ISO/OSI model, Layer 4, packet, transport layer, and TCP/IP.
TCP/IP	Transmission Control Protocol/Internet Protocol; the rules for transmitting data over networks and the Internet. <i>See</i> TCP.
Telnet	Part of TCP/IP, a protocol that enables a user to log onto a remote computer connected to the Internet. <i>See</i> TCP/IP.
traceroute	A utility that shows the route over which a packet travels to reach its destination.
Transmission Control Protocol (TCP)	<i>See</i> TCP.
Transmission Control Protocol/Internet Protocol (TCP/IP)	<i>See</i> TCP/IP.
transport layer	<i>See</i> Layer 4. <i>See also</i> ISO/OSI model.
TTL	Time-to-live, the length of time, in seconds, that a client's DNS server should cache a resolved IP address.
User Datagram Protocol (UDP)	Within TCP/IP, a protocol that is similar to Layer 4 (the transport layer). UDP converts data into packets to be sent from one server to another but does not verify the validity of the data. <i>See</i> ISO/OSI, TCP/IP, and transport layer.
view mode	One of two modes in which Equalizer can be administered: edit and view. In view mode, you can view—but not edit—parameters. <i>See</i> edit mode.
virtual cluster	An endpoint that acts as the network-visible port for a set of hidden back-end servers. <i>See</i> cluster, endpoint, FTP cluster, geographic cluster, and server cluster.
virtual server address	An IP address that is aliased to a physical server that has its own, separate IP address. <i>See</i> virtual web server.
virtual web server	Software that imitates HTTP server hardware. A virtual web server has its own domain name and IP address. <i>See</i> domain name, HTTP, IP address, server, and virtual server address. <i>See also</i> authoritative name server, back-end server, name server, physical server, and proxy server.
WAP	<i>See</i> Wireless Application Protocol.
weight	The relative proportion of a single item in a population of similar items. <i>See</i> dynamic weight, server weight, and static weight.
Wireless Application Protocol (WAP)	A set of rules that govern access to the Internet through wireless devices such as cellular telephones, pagers, and two-way communication devices.



"A Glossary of Networking Terms" 127

&& 88

|| 88

A

accessing

Equalizer Administration interface 29

active

connections 71

Active Connections cluster value 57

Active Connections server value 60

Active Content Verification 2

Active Content Verification. See ACV.

Active Requests geographic-cluster value 62

actual value, server static weight 83

ACV 2, 76, 127

enabling 77

probe string 77

response string 77

ACV probe string 2

ACV Probe String field 78

ACV Response string 2

ACV Response String field 78

adaptive load balancing 70, 71, 83, 104, 105

Add Geographic Cluster dialog box 104

Add Site dialog box 106

adding

geographic cluster 104

match rule to virtual cluster 91

server to cluster 81

server to virtual cluster 81

site to geographic cluster 106

virtual cluster 65

address

administration 9, 127

aliased IP 127

cluster 128

external 27, 128

failover gateway 30

internal 27, 129

IP 24

resource 106

server 132

translation 127

virtual server 133

adjusting

server's static weight 83

static weight of site 107

administration

address 9, 127

interface 29, 127

interface, changing password 25

administration password 47

agent 127

custom 109

custom server 109

Equalizer 100

IP address 106

retries 63

server 73, 132

site 63

Agent Address field 106

Agent Misses status 63

Agent Retries status 63

Agent's Address site parameter 63

agent-to-client triangulation probe 63

aggregation 127

sticky network 3, 42

aggressive load balancing 73

ALB algorithm 85

algorithm 127

algorithms

load balancing 106

load-balancing 5

alias 127

failover gateway 14

server 28

alias, failover 12

aliased IP address 127

all 89

any() 88

application layer. See Layer 7 (L7).

atom 113, 127

authoritative name server 5, 7, 26, 127

configuring 102

auto-sensing power supply 19

average network distance 64

Average Ping Time status 64

B

- back-end server 127
- backing up configuration 48
- backup 12
 - default 37
 - Equalizer 14, 36, 127
 - hot 12, 27, 36, 129
 - mode 14, 37, 39
 - server 82
 - unit 12, 27, 36
- backup Equalizer 12
- backup unit 12, 51
- beginning configuration 22
- boot process 21
- bound 113, 127
- BPDU (bridge protocol data unit) 37
- bracket expression 114, 127
- branch 113, 127
- bridge protocol data unit (BPDU) 37
- browser
 - Javascript-enabled 29, 30

C

- cache 127
- cache-time-to-live field 104
- card, XCEL 51
- certify_client 70
- Change Server Parameters dialog box 86
- Change Site dialog box 107
- changing
 - administration password 25, 47
 - configuration 38, 44, 45
 - console password 25
 - server's static weight 83, 84
 - static weight of site 107
- character-based interface 22
- checkboxes
 - Enable Outbound Network Address Translation 40
 - ICMP Triangulation 104
- checking
 - validity of server 76
- cipher suite 69
- Class A 127
- Class A network 42
- Class B 127
- Class B network 42
- Class C 128
- Class C network 42
- client request packet 8
- client timeout 46, 69
- cluster 128
 - adding 65
 - adding geographic 104

- adding server to 81
- adding site to geographic 106
- address 128
- configuration 56
- configuring to use server agents 73
- deleting 70
- deleting geographic 105
- deleting site from geographic 107
- displaying information about 56
- FTP 129
- geographic 5, 63, 129
- geographic load balancing 105
- HTTP 74
- HTTPS 74, 82
- Layer 4 (L4) 75, 78, 85
- Layer 7 (L7) 78, 85
- NFS server 2
- optimizing performance of geographic 105
- server 132
- statistics, plotting 57
- virtual 65, 133

- cluster performance, optimizing 83
- cluster value
 - Active Connections 57
 - Hit Rate 57
 - Server Agent 57
 - Servers 57
 - Service Time 57
- cluster, virtual 28
- clusters
 - heterogeneous 85
 - setting static weight for homogenous 84
 - setting static weights for mixed 85
- cluster-wide parameter
 - DNS TTL 62
 - Geographic Cluster 62
 - ICMP Triangulation 62
 - Load Balancing Method 62
 - Load Balancing Response 62
 - MX Exchanger 62
- collating element 114
- command 46
- commands
 - RST (reset) 41
- Commit option 24
- committing
 - changes to configuration parameters 24
- computed load 128
- Computed Load server value 60
 - server value
 - Computed Load 61
- conditions, server 109
- configuration

- backing up 48
- backup 12
- beginning 22
- cluster 56
- examples 16
- failover 12, 15, 27, 36, 37
- initial 21
- network 81
- network parameters 23
- parameters, committing changes to 24
- restoring saved 48
- saving 48
- server 27
- single network 11
- single-network 27
- testing 28
- two network 10
- two-network 27, 28
- understanding 8
- configuration utility, Equalizer 22
- configuration worksheet 16
- Configure Network Interfaces window 23
- configuring
 - authoritative name server to query Envoy 102
 - cluster to use server agents 73
 - cluster's load balancing options 70
 - Equalizer 21
 - events 44, 45
 - geographic cluster load-balancing options 105
 - network parameters 23
 - redundancy 38
 - second Equalizer as hot backup 36
 - servers 27
- connect timeout 46, 69
- connection
 - passive FTP 131
 - stale 51, 132
 - statistics 51
 - sticky 132
 - terminal 39
- connection timeout, stale 41
- connections
 - FTP data 73
 - sticky 3, 42, 75
- connector, RJ-45 network 19
- conserving IP addresses 14
- console 39
 - changing password 25
 - logging into 22
- console messages 39
- Console option 25
- cookie 128
 - holders 77

- lifetime 68
 - stuffing 74
- cookie generation 69
- Cookie Lifetime option 68
- cookie-based persistence 3, 77
- cord, power 19
- creating
 - custom server agent 109
- CTTL field 104
- custom event handling 44
- custom server agent 109
- cycle time 46
- cycle, diagnostic 14

D

- daemon 54, 128
 - server agent 73
- data connections, FTP 73
- date, setting 25
- decrypting HTTPS clusters 82
- default
 - backup 37
 - primary unit 27
 - route 27
- default match rule 87
- Default Router field 22
- Default Site site parameter 63
- defining
 - match rule 91
- delegating authority to Envoy site 102
- deleting 94
 - cluster 70
 - geographic cluster 105
 - match rule 94
 - server 83
 - site from geographic cluster 107
- device probe message 21
- diagnosing Equalizer installation and configuration problems 119
- diagnostic cycle 14
- diagnostic messages 21
- dialog boxes
 - Add Geographic Cluster 104
 - Add Site 106
 - Change Server Parameters 86
 - Change Site 107
- displaying
 - cluster information 56
 - Equalizer information 51
 - geographic cluster parameters 61
 - server information 58
 - site information 63
 - system log 54

- virtual cluster summary 53
- DNS 2, 5, 15, 21, 27, 52, 104, 128
- DNS Requests Received for Unknown Domains parameter 52
- DNS Server field 22
- DNS TTL 128
- DNS TTL cluster-wide parameter 62
- domain 5, 128
- domain name 5, 128
 - fully-qualified 5
- domain name server 24
- domain name service 5
- domain name, fully-qualified 104
- down 2, 14, 54
- dynamic
 - server agent 57
- dynamic weight 71, 83, 105, 128
 - oscillations 73
 - spread 71
- Dynamic Weight server value 61
- Dynamic Weight Spread option 71

E

- echo 128
- echo request, ICMP 100
- edit mode 25, 29, 32, 128
- edit mode password 30, 47
- Edit Password 47
- editing
 - match rule 93
- EIA 128
- element, collating 114
- emulation, VT100 22
- emulator, terminal 19
- enable outbound NAT 112
- Enable Outbound Network Address Translation checkbox 40
- enabling
 - ACV 77
 - inter-cluster stickiness 75
 - outbound NAT 112
 - passive FTP connections 41
 - persistent sessions 74
 - sticky connections 75
- encrypting HTTPS clusters 82
- endpoint 54, 128
- endpoint, server 132
- Envoy 1, 5, 26, 52, 63, 64, 99, 128
 - installing 102
- Envoy Geographic Load Balancing parameter 51
- Envoy site, delegating authority to 102
- Equalizer
 - agent 100

- ALB algorithm 85
- backup 14, 36, 127
- configuration utility 22
- configurations 8
- displaying information 51
- entry 53, 112
- kernel 14
- primary 30, 131
- second 12, 27
- shutting down 26, 49
- updating software 25
- upgrading software 25
- Equalizer Administration interface 29, 30, 38, 53, 87, 128
- Equalizer Configuration Menu window 22, 23, 25
- Equalizer Configuration Utility 128
- Equalizer front panel 19
- Equalizer Version parameter 51
- Equalizer's regular expressions (RE) 113
- Equalizers
 - Equalizers
 - two 36
- event handling, custom 44
- expression
 - bracket 127
 - regular (RE) 131
- expressions
 - bracket 114
 - regular (RE) 113
- external
 - address 27, 128
 - interface 8, 128
 - network 8, 10, 11, 129
 - test machine 28

F

- failed unit 27
- failover 12, 129
 - configuration 12, 15, 36, 37
 - gateway address 30
 - process 14, 36
- failover alias 12
- failover configuration 27
- failover configuration screen 38
- failover gateway
 - alias 14
- false 87, 88
- fine-tuning
 - site weight 105
- firewall 26, 129
 - network 103
- firewalled networks, using Envoy with 103
- forwarding log information 44, 45
- FQDN 104, 129

- front panel 19
- FTP 129
 - control port 73
 - data connections 73
 - passive connections 41
 - passive mode 73
 - passive translation 41
 - services, providing 73
- FTP cluster 129
- FTP connection, passive 131
- FTP PASV 41
- Fully Qualified Domain Name (FQDN) 129
- fully-qualified domain name 5, 104

G

- gateway 12, 24, 30, 81, 129
 - default route 27
 - using Equalizer between networks 10
- Gateway field 24
- geographic
 - cluster 5, 129
 - load balancing 1, 5, 26, 129
 - probe 100, 129
- geographic cluster 63
 - adding 104
 - adding site to 106
 - deleting 105
 - deleting site from 107
 - load balancing options 105
 - optimizing performance of 105
 - removing site from 107
- Geographic Cluster cluster-wide parameter 62
- Geographic Cluster Name field 104
- Geographic Cluster Parameters frame 105, 106
- Geographic Cluster Parameters page 62
- geographic cluster parameters, displaying 61
- Geographic Cluster site parameter 63
- geographic load balancing 26
- Geographic Query Protocol 26
- Geographic Site Parameters page 107
- geographic-cluster value
 - Active Requests 62
 - Network Latency 62
 - Request Rate 62
 - Site Summary 62
- global statistics 52
- graph, zooming in on 57, 60, 63, 64
- graphical history 57

H

- header 129
- headers
 - IP 2

- TCP/UDP 2
- heterogeneous clusters 85
- history, plotting geographic cluster 62
- Hit Rate cluster value 57
- holders, cookie 77
- homogenous clusters, setting static weight for 84
- host 5, 9
- Host field 23
- host, syslog 44
- Hostname field 21
- hot
 - backup 27, 36
 - spare 27
- hot backup 12, 14, 129
- HTTP 3, 4, 76, 77, 78, 85, 129
 - clusters 74
 - protocol 82
 - request 87
- HTTP server daemon 54
- HTTPS 3, 4, 77, 78, 85, 129
 - clusters 74, 82
 - request 87
- hub 37, 129
- HyperText Transfer Protocol (Secure). See HTTPS.
- HyperText Transfer Protocol. See HTTP.

I

- ICMP ECHO request 64
- ICMP echo request 100, 103, 106, 129
- ICMP echo request packet 27
- ICMP echo response packet 27
- ICMP probe 46
- ICMP triangulation 100, 104, 129
- ICMP Triangulation checkbox 104
- ICMP Triangulation cluster-wide parameter 62
- idle timeout 46
- ignore case 46
- initial configuration 21
- installation and configuration problems 119
- installing
 - Envoy 102
 - latest Equalizer software 25
- intelligent load balancing 129
- inter-cluster stickiness 3, 75
- interface 129
 - administration 29, 127
 - Equalizer Administration 29, 30, 38, 53, 128
 - external 8, 128
 - network 131
 - single-network 11
- interfaces
 - character-based 22
- Interfaces option 23

- internal
 - address 27, 129
 - network 10, 11, 129
- internal interface parameters 24
- internal-network test machine 28
- Internet 10
- Internet Control Message Protocol (ICMP) 129
- intranet 10
- IP 129
- IP address 9, 24, 28, 130
 - reserved 14
 - site agent 106
- IP Address field 24
- IP address, aliased 127
- IP address, site 63
- IP headers 2
- IP spoofing 81
- IP-address based persistence 3
- ISO/IEC 130
- ISO/IEC 11801 standard 127
- ISO/OSI model 130

J

- Javascript-enabled web browser 29, 30

K

- Keepalive value 106
- kernel, Equalizer 14

L

- L4 Connections Timed Out parameter 52
- L4 Peak Connections Processed parameter 51
- L4 Total Connections Processed parameter 51
- L4. See Layer 4 (L4).
- L7 Current Active Connections parameter 52
- L7 Peak Connections Processed parameter 52
- L7 Total Connections Processed parameter 52
- L7. See Layer 7 (L7).
- latency 5, 100, 130
- layer
 - Secure Sockets 132
- Layer 4 (L4) 41, 51, 52, 86, 130
 - cluster 75, 78, 85
 - load balancing 74
- Layer 4 load balancing 42
- Layer 7 (L7) 1, 3, 4, 52, 85, 130
 - cluster 78, 85
 - load balancing 74, 87
 - rules 87
- Layers 1, 2, 3, 5, and 6 130
- license 35, 123
- licensing 35
- load 130

- computed 128
- load balancing 87, 130
 - adaptive 70, 71, 83, 104, 105
 - aggressive 73
 - algorithms 106
 - geographic 1, 5, 26, 129
 - geographic cluster 105
 - intelligent 129
 - Layer 4 42
 - Layer 4 (L4) 74
 - Layer 7 (L7) 74, 87
 - methods 70, 105
 - options 70, 71
 - policy 57, 70, 83
 - response 104
 - responsiveness 57
 - round robin 70, 83
 - round trip 104, 105
 - site load 104, 105
 - site weight 104, 105
 - static weight 70, 83
 - UDP 2
 - WAP gateways 2
- Load Balancing Method cluster-wide parameter 62
- Load Balancing Response cluster-wide parameter 62
- Load Balancing Response option 104, 105
- load distribution, geographic 5
- load-balancing algorithms 5
- local menus 33
- local name server 6
- log hours 46
- logging into
 - Equalizer console 22
- logical AND 88
- logical name 82
- logical NOT operator 88
- logical OR 88
- login prompt 22

M

- machine
 - external test 28
 - internal-network test 28
 - test 28
- managing
 - servers 81
- match body 87, 89
- match expressions 88
- match rule 87, 94
 - adding to virtual cluster 91
 - defining 91
 - editing 93
- match rule, default 87

- matching expressions 115
- menu
 - configuration utility 22
- menus
 - local 33
- messages
 - console 39
 - device probe 21
 - diagnostic 21
 - server status 54
 - start-up 54
- mini_sendmail 117
- minimizing number of IP addresses needed 14
- mode
 - backup 14, 37, 39
 - edit 25, 29, 32, 128
 - operation 51
 - primary 14, 39
 - view 29, 32, 133
- model
 - ISO/OSI 130
- monitoring
 - cluster performance 84
- MX exchanger 130
- MX Exchanger cluster-wide parameter 62
- MX Exchanger field 104

N

- name
 - logical 82
- name resolution request 5
- name server 130
- Name Server field 24
- name server, authoritative 26, 127
- NAT 39, 130
 - outbound 15
 - subsystem 2
- NAT subsystem 73, 130
- NAT, outbound 112
- netmask 131
- network
 - address translation 2
 - average distance 64
 - configuration 81
 - connectivity 24
 - external 8, 10, 11, 129
 - interface 131
 - internal 10, 11, 129
 - latency 5
 - OSI 131
 - parameters, configuring 23
 - reserved 132
 - RJ-45 connector 19

- route 12
 - sticky aggregation 42
 - troubleshooting techniques 119
- Network Address Translation. See NAT.
- Network Configuration window 23
- network environment, using Equalizer in single 11
- network firewall 103
- Network Interfaces field 21
- Network Latency geographic-cluster value 62
- Network Latency site value 64
- networks
 - Class A 42
 - Class B 42
 - Class C 42
 - non-routable 111, 112
 - placing servers on 111
 - reserved 111, 112
- NFS server cluster 2
- no outbound RST 46
- no plot 46
- none 89
- non-routable networks 111, 112
- NOT operator 88

O

- operation modes 51
- Optimization Threshold 71
- optimization threshold 71
- optimizing
 - cluster performance 83, 107
 - geographic cluster performance 105
- optimizing cluster performance 83
- options
 - load balancing 70
- oscillations, dynamic weight 73
- OSI network 131
- outbound
 - NAT 15
- outbound NAT 39, 112

P

- P3/500Mhz-based server 83
- P3/900Mhz-based server 83
- packet 27, 131
 - ARP 14
 - ICMP echo request 27
 - ICMP echo response 27
 - request 8, 132
 - response 8, 132
 - SYN 41
 - TCP/UDP 4
- pages
 - parameters 33

- panel, front 19
- parameters
 - displaying geographic cluster 61
 - internal interface 24
 - server agent 57
- parameters pages 33
- partition 36
- passive
 - FTP connections 41
 - FTP translation 41
- passive FTP connection 131
- passive FTP mode 73
- password 22, 30
 - administration interface 25, 30, 47
 - changing 47
 - console 25
 - edit mode 25, 30, 47
- Password option 25
- PASV 41, 73, 131
- PASV FTP 51
- pattern match 131
- payload 131
- pedantic agent 46
- peer site 106
- performance
 - improving 5
 - monitoring 84
 - optimizing 83
 - optimizing cluster 83
 - statistics 73
- performing
 - outbound NAT 112
- persistence 131
 - cookie-based 3, 77
 - IP-address based 3
- persistent sessions 3
 - enabling 74
- physical network layout 16
- physical server 28, 131
- piece 113, 131
- ping 24, 27, 28, 100, 103, 131
- placing servers on networks 111
- plot clip 46
- Plot GeoCluster History 63
- Plot Site 64
- plotting
 - cluster statistics 57
 - geographic cluster history 62
 - geographic cluster statistics 62
 - site statistics 64
- port 9, 131
 - redirection 73, 82
- port 21 73

- Port field 106
- port number 131
- PortFast 37
- power cord 19
- power supply, auto-sensing 19
- primary
 - Equalizer 30
 - mode 14, 39
 - role 30
 - unit 12, 36, 37
- primary Equalizer 131
- primary unit 27, 51
- probe 131
 - agent-to-client triangulation 63
 - device 21
 - geographic 100, 129
 - server 54
 - site 63
- probe delay 69
- probe interval 46
- probe string
 - ACV 77
- Probes Missed site value 64
- problems, solving 119
- protocol 131
 - SSL 4
 - stateless 2
- protocol stack 131
- protocols
 - HTTP 82
 - UDP-based Geographic Query 26
- providing
 - FTP services on virtual cluster 73
- proxy server 131

Q

- quiesce 131
- quiescing servers 85

R

- RADIUS 2, 131
- receive buffer 46, 69
- record
 - sticky 3
- redirection 131
- redirection, port 73, 82
- register (see license) 35
- regular expression 89
- regular expression (RE) 131
- regular expressions (RE) 113
- relative value, server static weight 83
- relative workload 64
- Remote Authentication Dial-In User Service. See RADIUS.

- removing
 - cluster 70
 - geographic cluster 105
 - Layer 4 server from service 86
 - Layer 7 server from service 85
 - server 83
 - site from geographic cluster 107
- request
 - HTTP 87
 - HTTPS 87
- request max 69
- request packet 2, 8, 132
- Request Rate geographic-cluster value 62
- reserved network 14, 132
- reserved networks 111, 112
- resolution 132
- resolution request 5
- resource
 - address 106
- Resource Address field 106
- Resource Down site value 64
- Resource Errors status 63
- Resource Keepalive site parameter 63
- Resource Load site value 64
- Resource Load status 63
- Resource site parameter 63
- response
 - settings 70, 71
- response max 69
- response packet 8, 132
- response string
 - ACV 77
- response time, server 71
- restoring
 - saved configuration 48
- retries, agent 63
- Returned as Default status 64
- RFC1208 127
- RJ-45 network connector 19
- round robin load balancing 70, 83
 - weighted 70
- round trip load balancing 104, 105
- route, network 12
- router 9, 22, 132
- routing table 132
- routing tables 9
- RST 132
- RST (reset) command 41
- rules
 - Layer 7 (L7) 87
 - match 87

S

- saving
 - configuration 48
- second Equalizer 12, 27
- Secure Sockets Layer (SSL) 132
- send buffer 45, 69
- serial
 - terminal 19
- server 54, 132
 - adding 81
 - address 132
 - agent 132
 - alias 28
 - authoritative name 5, 7, 26, 127
 - back-end 127
 - checking validity 76
 - cluster 132
 - conditions 109
 - configuration 27
 - displaying information about 58
 - domain name 24
 - endpoint 132
 - IP address 28
 - local name 6
 - name 130
 - physical 131
 - proxy 131
 - resource availability 73
 - response time 71
 - shutting down 85
 - virtual web 133
 - weight 83, 132
 - weights 83
- server address, virtual 133
- server agent 73
 - configuring cluster to use 73
 - daemon 73
 - parameters 57
 - using 73
 - value 71
- Server Agent cluster value 57
- Server Agent server value
 - server value
 - Server Agent 61
- server agent, custom 109
- server agents 2
- server status messages 54
- server timeout 46, 69
- server value
 - Active Connections 60
 - Computed Load 60
 - Dynamic Weight 61
- servers

- backup 82
- changing static weight of 84
- deleting 83
- Layer 4 (L4) 86
- Layer 7 (L7) 85
- managing 81
- P3/500Mhz-based 83
- P3/900Mhz-based 83
- placing on networks 111
- quiescing 85
- Servers cluster value 57
- Service Time cluster value 57
- Service Time server value
 - server value
 - Service Time 60
- session
 - telnet 27
- session cache kbytes 69
- session cache timeout 69
- sessions
 - enabling persistent 74
 - persistent 3
- setting
 - date and time 25
 - static weights for homogenous clusters 84
 - static weights for mixed clusters 85
 - time zone 25
- settings
 - response 70
- Shutdown option 26
- shutting down
 - server 85
- shutting down Equalizer 26, 49
- sibling 12
- sibling Equalizers 36
- single network environment, using Equalizer in 11
- single-network
 - configuration 11
 - interface 11
- single-network configuration 27
- site 5, 132
 - adding to geographic cluster 106
 - deleting 107
 - displaying information about 63
 - IP address 63
 - peer 106
- Site Address field 106
- Site Chosen site value 64
- Site IP Address site parameter 63
- site load balancing 104, 105
- Site Name field 106
- site parameter
 - Agent's Address 63
 - Default Site 63
 - Geographic Cluster 63
 - Resource 63
 - Resource Keepalive 63
 - Site 63
 - Site IP Address 63
 - Static Weight 63
- Site Parameters frame 107
- Site Returned status 64
- Site site parameter 63
- site summary 63
- Site Summary geographic-cluster value 62
- site value
 - Network Latency 64
 - Probes Missed 64
 - Resource Down 64
 - Resource Load 64
 - Site Chosen 64
 - Triangulation Errors 64
- site weight
 - fine-tuning 105
 - load balancing 104, 105
- site-wide failure 5
- software license 123
- software, updating Equalizer 25
- solving installation and configuration problems 119
- Spanning Tree 37
- spoofing 132
 - IP 81
- SSL Acceleration parameter 51
- SSL protocol 4
- SSL. See Secure Sockets Layer.
- ssl_unclean_shutdown 70
- stack 132
- stack, protocol 131
- stale connection 51, 132
- stale connection timeout 41
- standards
 - ISO/IEC 11801 127
- start-up messages 54
- state 132
 - valid 36
- stateless 132
- stateless protocol 2
- static weight 1, 58, 132
 - changing 83, 84
 - load balancing 70, 83
- static weight of site, adjusting 107
- Static Weight option 107
- Static Weight site parameter 63
- Static Weight value 106
- statistics
 - connection 51

- performance 73
- plotting 57
- plotting geographic cluster history 62
- plotting site 64
- status
 - Agent Misses 63
 - Agent Retries 63
 - Average Ping Time 64
 - Resource Errors 63
 - Resource Load 63
 - Returned as Default 64
 - Site Returned 64
 - Triangulation Time-outs 63
- stickiness, inter-cluster 3
- sticky
 - connection 132
 - connections 3, 42
 - network aggregation 3, 42
 - record 3
 - time period 3, 73
 - timer 3, 132
- sticky connections, enabling 75
- sticky time period 75
- strikeout threshold 46
- stuffing cookie 74
- sub-daemon max 69
- subdomain 133
- subnet 133
- summary
 - virtual cluster 53
- summary, site 63
- switch 133
- SYN flood attack 41
- SYN packet 41
- SYN/ACK 41, 133
- syslog 133
- syslog host 44
- System Event Log 54
- system log, displaying 54

T

- table, routing 132
- TCP 42, 85, 133
- TCP/IP 133
- TCP/UDP
 - headers 2
 - packet 4
- Telnet 133
- telnet 27, 28
- telnet session 27
- terminal 21
 - emulator 19
 - serial 19

- terminal connection 39
- test machine 28
- test machine, external 28
- testing configuration 28
- threshold, optimization 71
- time
 - server response 71
 - setting 25
 - sticky 3
- Time option 25
- time period, sticky 73, 75
- Time Zone option 25
- time zone, setting 25
- timeout, stale connection 41
- timer, sticky 3, 132
- traceroute 27, 133
- translation, address 127
- Transmission Control Protocol. See TCP.
- Transmission Control Protocol/Internet Protocol. See TCP/IP.
- transport layer. See Layer 4 (L4).
- triangulation
 - ICMP 104
 - Triangulation Errors site value 64
 - Triangulation Time-outs status 63
- triangulation, ICMP 100, 129
- troubleshooting installation and configuration problems 119
- troubleshooting techniques, network 119
- true 87, 88
- truth value 88
- TTL 133
- two-network configuration 10, 27, 28

U

- UDP 2, 27, 42, 76, 85, 133
- UDP load balancing 2
- UDP-based Geographic Query Protocol 26
- unit
 - backup 27, 36
 - default primary 27
 - failed 27
 - primary 27, 36, 37
- Upgrade option 25
- upgrading Equalizer 25
- URL 30
- User Datagram Protocol. See UDP
- user mode 32
- user name 30
- using
 - ACV 76
 - Envoy with firewalled networks 103
 - Equalizer as gateway between networks 10

- Equalizer in single network environment 11
- reserved IP addresses 14
- second Equalizer as backup 12
- server agents 73
- utilities
 - Equalizer Configuration 128
- utility
 - Equalizer configuration 22

V

- valid state 36
- view mode 29, 32, 133
- View Password 47
- viewing
 - a cluster's graphical history 58
 - a server's graphical history 60
 - a site's graphical history 64
 - cluster information 56
 - Equalizer information 52
 - geographic cluster's graphical history 62
- virtual
 - cluster 133
 - server address 133
 - web server 133
- virtual cluster 1, 28, 65
 - adding 65
 - adding match rule to 91
 - adding server to 81
 - deleting 70
 - FTP services, providing 73
 - geographic 5
- virtual cluster summary 53
- VT100 emulation 22

W

- WAP gateway 2
- WAP. See Wireless Application Protocol
- warranty 19
- web browser
 - Javascript-enabled 29, 30
- web server, virtual 133
- weight 133
 - adjusting server 83
 - changing static 84
 - dynamic 71, 83, 105
 - fine-tuning site 105
 - oscillations 73
 - server 83, 132
 - site 104
 - spread coefficient 71
 - static 1, 58, 132
- Weight Spread Coefficient option 71
- weighted round robin load balancing 70

- window
 - Configure Network Interfaces 23
 - Equalizer Configuration Menu 22, 23, 25
 - Network Configuration 23
- Wireless Application Protocol (WAP) 133
- wireless application protocol (WAP) 2
- workload
 - relative 64
- worksheets
 - configuration 16
- writing
 - custom agents 109

X

- XCEL card 51

Z

- zooming in on graph 57, 60, 63, 64