



Load Balancer Administration Guide

Contents

Chapter 3. Product overview 9

What is new in this release	10
Components of Load Balancer that are available to use	11
Dispatcher component	12
Content Based Routing (CBR) component	16
Site Selector component	19
High availability with Load Balancer	21
Managing servers	24
Types of cluster, port, and server configurations for Dispatcher	26

Chapter 4. Installing Load Balancer 29

Setting up IBM Installation Manager	29
Installing the product on distributed operating systems	33
Installing the product with the GUI	36
Installing the product silently	38
Uninstalling the product	44
Uninstalling Load Balancer from distributed operating systems using the GUI	45
Uninstalling Load Balancer from distributed operating systems using response files	45
Uninstalling Load Balancer if it was installed without Installation Manager	48
Installing fix packs	52
Installing fix packs on distributed operating systems using the GUI	53
Installing fix packs on distributed operating systems using the command line	54
Uninstalling fix packs from distributed operating systems using the GUI	58
Directory conventions	59

Chapter 5. Configuring Load Balancer 61

Configuring Dispatcher	61
Quick start configuration	61
Plan for Dispatcher	64
Methods of configuration for Dispatcher	65
Types of cluster, port, and server configurations for Dispatcher	67
Configuring the Dispatcher machine	69
Configuring the server machines	73
Configuring collocated servers for the Dispatcher component	80
Load balancing a private network	82
Configuring Site Selector	83
Quick start configuration for Site Selector	83
Plan for Site Selector	85
Methods of configuration for Site Selector	88
Setting up Site Selector	91
Configuring Content Based Routing	92
Quick start configuration for Content Based Routing	92
Plan for Content Based Routing	95

Methods of configuration for Content Based Routing	98
Types of cluster, port, and server configurations for Content Based Routing	102
Setting up the Content Based Routing component	105

Chapter 6. Administering Load Balancer 111

Enabling advisors to manage load balancing	111
Advisors	113
List of advisors	114
Getting service-specific advice with the advisor request or response option	118
Configuring the LDAP URI advisor	119
Getting advice with Metric Server	121
The Workload Management Advisor	124
Creating a custom advisor	125
Configuring high availability	146
How high availability works	148
Detecting server failures with heartbeats and reach targets	149
High Availability recovery strategy for failed servers	150
Scripts to run with high availability	151
Using NAT forwarding to forward traffic	152
Use encapsulation forwarding to forward traffic across network segments	154
Quiesce servers or services for maintenance windows	156
Restricting incoming traffic with ipchains and iptables	157
Logging with Load Balancer	158
Logging server statistics with binary logging	160
Support for ICMP forwarding and messaging	162
Configure rules to manage traffic to busy or unavailable servers	162
Optimize connections with client-to-server affinity	166
Optimize connections with client-to-server affinity for the Dispatcher component	166
Optimize connections with client-to-server affinity for the Content Based Router component	169
Optimize connections with client-to-server affinity for the Site Selector component	171
Sample scripts to generate alerts and record server failure	174

Chapter 7. Tuning Load Balancer 175

The manager report	175
Optimizing the manager interval	177
Tuning the proportion of importance given to status information	178
Managing traffic with server weights	179
Reset connections to down servers	180

Update the network configuration automatically	181
Optimizing the sensitivity threshold	182
Optimizing the smoothing index	182
Controlling connection records with the staletimeout value	183

Chapter 8. Troubleshooting Load

Balancer components 185

Troubleshooting Content Based Router	185
Checking CBR port numbers	189
Troubleshooting Dispatcher	189
Checking Dispatcher port numbers	202
Troubleshooting Site Selector	203
Checking Site Selector port numbers	206

Chapter 9. Reference 207

Advanced configuration	207
Directory conventions	207
Types of cluster, port, and server configurations for Dispatcher	207
GUI: General instructions	209
Custom advisor methods and function calls	210
List of advisors	215
Sample scripts to generate alerts and record server failure	219
High Availability recovery strategy for failed servers	220
Scripts to run with high availability	220
Sample 2-tier high availability configuration with Dispatcher, CBR, and Caching Proxy	221
Content rule (pattern) syntax	224
Commands	226

dscontrol advisor	227
dscontrol binlog	231
dscontrol cluster	231
dscontrol executor	233
dscontrol file	234
dscontrol help	236
dscontrol highavailability	236
dscontrol logstatus	239
dscontrol manager	239
dscontrol metric	243
dscontrol port	244
dscontrol rule	247
dscontrol server	249
dscontrol set	253
dscontrol status	253
sscontrol advisor	254
sscontrol file	258
sscontrol help	259
sscontrol logstatus	260
sscontrol manager	261
sscontrol metric	264
sscontrol nameserver	265
sscontrol rule	265
sscontrol server	268
sscontrol set	269
sscontrol sitename	270
sscontrol status	272
Examples	273
Example: Sample advisor	273
Example: Implementing custom advisors	276
Glossary	287

Chapter 1. Edge Components, Version 8.5

View the latest Edge Components information. This information applies to the current version and to all subsequent releases and modifications until otherwise indicated in new editions. Edge is supported on Network Deployment.

- Latest updates to the information center - Read the latest on changes made to the information and the frequency of information center updates.
- Survey: Importance of High Quality Technical Information - Provide quantitative feedback on the technical information that you access through mechanisms such as information centers, embedded help in the user interface, the Support Portal, the developerWorks Web site, and IBM Redbooks publications.
- "Learning"
- "Tasks" on page 2
- "Community and Support" on page 2

Learning

The following topics in the information center will help you learn about the product, as well as the various technologies for supporting and enhancing your WebSphere® applications.

Product tips

Avoid trouble Search this site for 'gotcha' to find common troubleshooting tips.

Behavior changes Search this site for 'trns' to find behavior changes between versions or service releases.

Best practices Search this site for 'bprac' to find common best practices tips.

Configuration notes Search this site for 'sptcfg' to find supported configurations.

Deprecations Search this site for 'depfeat' to find deprecated items.

Mixed version considerations Search this site for 'mixv' to find issues in environments that use mixed versions of WebSphere Application Server.

New features Search this site for 'newfeat' to learn about new features in the product.

Requires restart Search this site for 'restart' to find tasks that require you to restart WebSphere Application Server.

What's new

What is new for Load Balancer for IPv4 and IPv6

Getting started with Product information

Detailed system requirements

Library page (documentation in PDF format and more)

developerWorks®

Education center

WebSphere Application Server zone

Education and certification

IBM® Education Assistant

WebSphere Application Server software education

WebSphere software certification
WebSphere training - IBM IT Education Services
WebSphere software platform training and certification

Information roadmaps

Roadmap for WebSphere Application Server

Learning resources

New to WebSphere
Resource reference list for WebSphere software products
Software product compatibility reports

Product information

WebSphere Application Server software
WebSphere software

Redbooks®

Redbooks query for the WebSphere Application Server
WebSphere Redbooks Domain

Third-party information

Java™ EE on the Oracle website

Tasks

The primary mission of the information center is to support your quick, easy, and successful performance of tasks.

Using Load Balancer for IPv4 and IPv6

Chapter 2, "Load Balancer for IPv4 and IPv6," on page 5

Glossary

View terms and definitions

Community and Support

Join a community to discover and share information with other product users.

Content feeds

Articles and Tutorials for WebSphere (developerWorks)
Articles and Tutorials about web services (developerWorks)
Downloads and Demos for WebSphere (developerWorks)
IBM Redbooks - WebSphere

Community links

developerWorks blogs
Follow WASInfo4U on Twitter
Global WebSphere Community

Discussion spaces

Newsgroup - [ibm.software.websphere.*](https://groups.google.com/group/ibm.software.websphere) on Google groups
WebSphere forums and communities

Feedback

Participate in product survey to help make IBM products easier to install and use Help make IBM products easier to install and use.

Send feedback to the documentation team.

Share and discover

All groups on developerWorks

All wikis on developerWorks

developerWorks articles

developerWorks forum

developerWorks "Meet the Experts"

Subscribe to Wiki updates

TechNotes

Techdocs - The Technical Sales Library

Wiki - About WebSphere product information

WebSphere events

WebSphere Technical Exchange Webcasts

Support

IBM Support Assistant

Steps to getting support

Support Technotes

WebSphere Application Server support

© Copyright IBM Corporation 2002, 2012.

This information center is Built on Eclipse™. (www.eclipse.org)

Chapter 2. Load Balancer for IPv4 and IPv6

View the latest information for Load Balancer for IPv4 and IPv6. This information applies to the current version and to all subsequent releases and modifications until otherwise indicated in new editions. Load Balancer is supported on Network Deployment.

- Latest updates to the information center - Read the latest on changes made to the information and the frequency of information center updates.
- Survey: Importance of High Quality Technical Information - Provide quantitative feedback on the technical information that you access through mechanisms such as information centers, embedded help in the user interface, the Support Portal, the developerWorks Web site, and IBM Redbooks publications.
- "Learning"
- "Tasks" on page 6
- "Community and Support" on page 7

Learning

The following topics in the information center will help you learn about the product, as well as the various technologies for supporting and enhancing your WebSphere applications.

Product tips

New features Search this site for 'newfeat' to find learn about new features in the product.

Best practices Search this site for 'bprac' to find common best practices tips.

Avoid trouble Search this site for 'gotcha' to find common troubleshooting tips.

Deprecations Search this site for 'depfeat' to find deprecated items.

What's new

Release Notes

Getting started with Product information

Detailed system requirements

Library page (documentation in PDF format and more)

Search keywords

Behavior changes: trns

New features: newfeat

Deprecations: depfeat

Best practices: bprac

Configuration notes: sptcfg

Mixed versions: mixv

Avoid trouble: gotcha

Learning resources

IBM Education Assistant

New to WebSphere

Resource reference list

Take a course

Courses for administrators

Courses for developers

WebSphere education

Information roadmaps

Roadmap for WebSphere Application Server

Training and certification

WebSphere software platform training and certification (includes tutorials)

WebSphere Application Server software education

WebSphere software certification

WebSphere training - IBM IT Education Services

Education center

IBM Education Assistant

developerWorks

Education Center

WebSphere Application Server zone

Redbooks

WebSphere Redbooks Domain

Redbooks query for the WebSphere Application Server

Product information

WebSphere Application Server software

WebSphere Software

Third-party information

Java EE on the Oracle website

Tasks

The primary mission of the information center is to support your quick, easy, and successful performance of tasks.

Installing

Chapter 4, "Installing Load Balancer," on page 29

Configuring

Chapter 5, "Configuring Load Balancer," on page 61

Administering

Chapter 6, "Administering Load Balancer," on page 111

Tuning

Chapter 7, "Tuning Load Balancer," on page 175

Troubleshooting

Chapter 8, "Troubleshooting Load Balancer components," on page 185

Glossary

View terms and definitions

Community and Support

Join a community to discover and share information with other product users.

Content feeds

- IBM Redbooks - WebSphere**

- Articles and Tutorials for WebSphere (developerWorks)**

- Articles and Tutorials about Web services (developerWorks)**

- Downloads and Demos for WebSphere (developerWorks)**

Discussion spaces

- WebSphere forums and communities

- Newsgroup - [ibm.software.websphere.*](#) on Google groups

Share and discover

- developerWorks articles

- developerWorks forum

- developerWorks "Meet the Experts"

- Subscribe to Wiki updates

- TechNotes

- Techdocs - The Technical Sales Library

- Wiki - About WebSphere product information

- All wikis on developerWorks

- All groups on developerWorks

- WebSphere Technical Exchange Webcasts

- WebSphere events

Community links

- Follow WASInfo4U on Twitter

- developerWorks blogs

- Global WebSphere Community forum – WebSphere Application Server

Support

- Support Technotes

- Ten steps to getting support

Feedback

- Participate in product survey to help make IBM products easier to install and use
- Help make IBM products easier to install and use.

- Send feedback to the documentation team.

© Copyright IBM Corporation 2002, 2012.

This information center is Built on Eclipse™. (www.eclipse.org)

Chapter 3. Product overview

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

When used with Web servers, Load Balancer can help maximize the potential of your site by providing a powerful, flexible, and scalable solution to peak-demand problems. If visitors to your site can not get through at times of greatest demand, use Load Balancer to automatically find the optimal server to handle incoming requests, thus enhancing your customers' satisfaction and your profitability.

What are the advantages to using Load Balancer?

The number of users and networks connected to the global Internet is growing exponentially. This growth is causing scalability problems that can limit users' access to popular sites. Currently, network administrators are using numerous methods to try to maximize access. With some of these methods, you can choose a different server at random if an earlier choice is slow or not responding. This approach is cumbersome, annoying, and inefficient. Another method is standard round-robin, in which the domain name server selects servers in turn to handle requests. This approach is better, but still inefficient because it forwards traffic without any consideration of the server workload. In addition, even if a server fails, requests continue to be sent to it. The need for a more powerful solution has resulted in Load Balancer. It offers numerous benefits over earlier and competing solutions:

- **Scalability:** As the number of client requests increases, you can add servers dynamically, providing support for tens of millions of requests per day, on tens or even hundreds of servers.
- **Efficient use of equipment:** Load balancing ensures that each group of servers makes optimum use of its hardware by minimizing the hot-spots that frequently occur with a standard round-robin method.
- **Easy integration:** Load Balancer uses standard TCP/IP protocols. You can add it to your existing network without making any physical changes to the network. It is simple to install and configure.
- **Low overhead** Using a simple MAC level forwarding method, the Dispatcher component looks at the inbound client-to-server flows only. It does not need to see the outbound server-to-client flows. This significantly reduces its impact on the application compared with other approaches and can result in improved network performance.
- **Content-based routing (using the CBR component)** In conjunction with Caching Proxy, the CBR component has the ability to proxy HTTP and HTTPS (SSL) requests to specific servers based on the content requested. For example, if a request contains the string "/cgi-bin/" in the directory portion of the URL, and the server name is a local server, CBR can direct the request to the best server in a set of servers specifically allocated to handle cgi requests.

- **High availability:** The Dispatcher component offers built-in high availability, utilizing a backup machine that remains ready at all times to take over load balancing if the primary server machine fails. When one of the servers fails, requests continue to be serviced by the other server. This process eliminates any server as a single point of failure and makes the site highly available.
- **Client to server affinity:** The affinity feature maps a client IP address to a back-end server, providing a higher level of efficiency by decreasing the memory and CPU utilization when compared to traditional connection forwarding.
- **Load balancing a private network:** You can set up Dispatcher and the TCP server machines using a private network. This configuration can reduce the contention on the public or external network that can affect performance.
- **Learn how you can manage servers using Load Balancer.**

Load Balancer balances traffic among your servers through a unique combination of load balancing and management software. All client requests sent to the Dispatcher machine are directed to the “best” server according to weights that are set dynamically. You can use the default values for those weights or change the values during the configuration process.

Dispatcher can also detect a failed server and forward traffic around it. Dispatcher supports HTTP, FTP, SSL, SMTP, NNTP, IMAP, POP3, Telnet, SIP, and any other TCP based application.

Load Balancer is the key to stable, efficient management of a large, scalable network of servers. You can link many individual servers into what seems to be a single, virtual server. Your site is presented as a single IP address to the world. Dispatcher functions independently of a domain name server; all requests are sent to the IP address of the Dispatcher machine.

Dispatcher brings distinct advantages in balancing traffic load to clustered servers, resulting in stable and efficient management of your site.

Related tasks:

“Managing servers” on page 24

You can load balance traffic to existing server topologies without changing the physical configuration of the machines or how clients will connect to your site.

Related reference:

“Types of cluster, port, and server configurations for Dispatcher” on page 26

There are many ways that you can configure Dispatcher to support your site.

What is new in this release

Load Balancer for IBM WebSphere Application Server contains a number of new features.

Note: The new features that are described in this topic are current as of the original release of the product version. Changes in this documentation that are related to service releases are marked with fix pack icons.

The new features in this release are:

“Content Based Routing (CBR) component” on page 16

Content based routing provides the ability to balance the load based on the content of the client requests, such as URI.

“Site Selector component” on page 19

Site Selector provides the ability to balance the load on servers within a local or wide area network using a DNS round-robin approach or a more

advanced user-specified approach. Site Selector works in conjunction with a name server to map DNS names to IP addresses.

“Using NAT forwarding to forward traffic” on page 152

Using Network Address Translation (NAT) Dispatcher capability removes the limitation for the backend servers to be located on a locally attached network.

“Configuring collocated servers for the Dispatcher component” on page 80

Load Balancer can reside on the same machine as a server for which it is load balancing requests. This is commonly referred to as collocating a server. It is supported on Linux and AIX only.

Components of Load Balancer that are available to use

This topic highlights the components of Load Balancer that are available for use.

Load Balancer consists of the following three components that you can use separately or together to provide superior load-balancing results:

- You can use the Dispatcher component by itself to balance the load on servers within a local area network or wide area network using a number of weights and measurements that are dynamically set by Dispatcher. This component provides load balancing at a level of specific services, such as HTTP, FTP, SSL, NNTP, IMAP, POP3, SMTP, SIP, and Telnet. It does not use a domain name server to map domain names to IP addresses.
- For both HTTP and HTTPS (SSL) protocol, you can use the Content Based Routing (CBR) component to load balance based on the content of the client request. A client sends a request to Caching Proxy, and Caching Proxy sends the request to the appropriate server. The chosen server is the result of matching the URL to a specified rule.
- You can use the Site Selector component to balance the load on servers within a local or wide area network using a DNS round-robin approach or a more advanced user-specified approach. Site Selector works in conjunction with a name server to map DNS names to IP addresses.

For more information on the Dispatcher, CBR, and Site Selector, components, see

Related concepts:

“Dispatcher component”

Learn about the Dispatcher component of Load Balancer.

“Content Based Routing (CBR) component” on page 16

CBR works with Caching Proxy to proxy client requests to specified HTTP or HTTPS (SSL) servers. Use CBR to manipulate caching details for faster web document retrieval with low network bandwidth requirements. CBR and Caching Proxy examines HTTP requests by using specified rule types.

“Site Selector component” on page 19

Site Selector acts as a name server that works with other name servers in a domain name system to load balance among a group of servers by using measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers based on the domain name used for a client's request.

Chapter 3, “Product overview,” on page 9

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Dispatcher component

Learn about the Dispatcher component of Load Balancer.

The Dispatcher component balances traffic among your servers through a unique combination of load balancing and management software. Dispatcher can also detect a failed server and forward traffic around it. Dispatcher supports HTTP, FTP, SSL, SMTP, NNTP, IMAP, POP3, Telnet, SIP, and any other TCP or stateless UDP-based application.

All client requests that are sent to the Dispatcher machine are directed to the best server according to weights that are set dynamically. You can use the default values for those weights or change the values during the configuration process.

Dispatcher offers three forwarding methods:

- **MAC forwarding method:** Use this forwarding method to load balance the incoming requests among the servers in same subnet. The servers respond directly to the client, bypassing the Dispatcher.
- **NAT forwarding method:** This forwarding method removes the limitation of the servers being required to be in same subnet as the Dispatcher. With the NAT forwarding method, Dispatcher load balances the incoming requests to the servers, the servers respond to the Dispatcher and the Dispatcher then forwards the response to the client.
- **Encapsulation forwarding:** This forwarding method is midway between the NAT and MAC forwarding methods. First, encapsulation forwarding enables the servers to be in a subnet different from the Dispatcher machine. Second, servers respond directly to the client, bypassing the Dispatcher. Dispatcher supports two

encapsulation protocols: GRE and IPIP. Not all operating systems support both GRE and IPIP protocols, hence their support is operating system dependent.

The Dispatcher component is the key to stable, efficient management of a large, scalable network of servers. With Dispatcher, you can link many individual servers into what seems to be a single, virtual server. Your site is presented as a single IP address to the world. Dispatcher functions independently of a domain name server; all requests are sent to the IP address of the Dispatcher machine.

Dispatcher brings distinct advantages in balancing traffic load to clustered servers, resulting in stable and efficient management of your site.

Dispatcher component features:

High availability

To use Dispatcher to remove single point-of-failure limitations in your network, read the [Configuring High availability](#) topic.

Client to server affinity

To ensure that the client uses the same server for multiple connections, read the [Optimize connections with client-to-server affinity](#) for the Dispatcher component topic.

Rules-based load balancing

To run Dispatcher on the same machine as one of the managed servers that you are load balancing, read the [Configuring collocated servers](#) for the Dispatcher component topic.

Wide area load balancing

You can achieve wide area load balancing through two methods: Encapsulation Forwarding and Network Address Translation (NAT).

To load balance to remote servers by using Dispatcher's Encapsulation Forwarding, read the [Use encapsulation forwarding to forward traffic across network segments](#) topic.

To load balance to remote servers by using Dispatcher's NAT forwarding method, read the [Using NAT forwarding to forward traffic](#) topic.

Alerts To generate alerts when servers are marked up or down, read the [Sample scripts to generate alerts and record server failure](#).

Collocation

To run Dispatcher on the same machine as one of the managed servers that you are load balancing, read the [Using collocated servers](#) for the Dispatcher component topic.

Functions that provide load balancing

The primary functions of Dispatcher interact with each other and your server configuration to balance network traffic in your environment. Dispatcher consists of the following functions:

- The **dsserver** command handles requests from the command line to the executor, manager, and advisors.
- The executor supports port-based load balancing of TCP connections. It is able to forward connections to servers based on the type of request received (for example, HTTP, FTP, SSL). The executor always runs when the Dispatcher component is being used for load balancing.
- The manager sets weights that are used by the executor that is based on:

- Internal counters in the executor
- Feedback from the servers that are provided by the advisors
- Feedback from a system-monitoring program, such as Metric Server or WLM.
- Using the manager is optional; however, if the manager is not used, load balancing is performed by using weighted round-robin scheduling, based on the current server weights, and advisors are not available
- The advisors query the servers and analyze results by protocol before calling the manager to set weights as appropriate. Currently there are advisors available for the following protocols: HTTP, FTP, SSL, SMTP, NNTP, IMAP, POP3, SIP, and Telnet.

Dispatcher also offers advisors that do not exchange protocol-specific information, such as the DB2[®] advisor that reports on the health of DB2 servers and the ping advisor that reports whether the server responds to a ping. For a complete list of advisors, Read the List of advisors topic. You also have the option of writing your own advisors (see Creating a custom advisor).

Tip: Using the advisors is optional, but recommended.

Configure and manage the executor, advisors, and manager by using the command line (**dscontrol**) or the graphical user interface (**ladmin**).

The three key functions of Dispatcher (executor, manager, and advisors) interact to balance and dispatch the incoming requests between servers. Along with load balancing requests, the executor monitors the number of new connections, active connections, and connections in a finished state. The executor also does garbage collection of completed or reset connections and supplies this information to the manager.

The manager collects information from the executor, the advisors, and a system-monitoring program, such as Metric Server. Based on the information the manager receives, it adjusts how the server machines are weighted on each port and gives the executor the new weighting for use in its balancing of new connections.

The advisors monitor each server on the assigned port to determine the server's response time and availability and then give this information to the manager. The advisors also monitor whether a server is up or down. Without the manager and the advisors, the executor does round-robin scheduling that is based on the current server weights.

Managing local servers with Dispatcher

You can install the Dispatcher machine without making any physical changes to the network. After a client request is directed to the optimal server by the Dispatcher, the response is sent directly from server to client with no involvement by the Dispatcher when you are using the MAC forwarding method.

. Physical representation of the site that uses an Ethernet network configuration.

Managing local and remote servers with Dispatcher

Wide area support in Dispatcher enables you to use both local and remote servers (servers on different subnets). The figure shows a configuration where one local Dispatcher (Dispatcher 1) serves as the entry point for all requests. It distributes

these requests among its own local servers (ServerA, ServerB, ServerC) and to the remote Dispatcher (Dispatcher 2), which load balances to its local servers (ServerG, ServerH, ServerI). When you use Dispatcher's NAT forwarding method or GRE support, wide area support with Dispatcher can also be achieved without using a Dispatcher at the remote site (where ServerD, ServerE, and ServerF are located).

- . Physical representation of a site that uses an Ethernet network configuration.

Related concepts:

Chapter 3, “Product overview,” on page 9

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

“Plan for Dispatcher” on page 64

You can use the Dispatcher component by itself to balance the load on servers within a local area network or wide area network, using a number of weights and measurements that are dynamically set by Dispatcher.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Optimize connections with client-to-server affinity for the Dispatcher component” on page 166

The Dispatcher affinity feature maps a client IP address to a back-end server. Affinity is established once a packet's destination IP address matches the cluster, the destination port matches the Load Balancer port, and the source IP address matches.

“Configuring collocated servers for the Dispatcher component” on page 80

Learn about advanced features for configuring the Dispatcher component with collocated servers.

“Using NAT forwarding to forward traffic” on page 152

Using Network Address Translation (NAT) Dispatcher capability removes the limitation for the backend servers to be on a locally attached network. With the NAT forwarding method, Dispatcher load balances the incoming request to the server. The server returns the response to Dispatcher. The Dispatcher machine then returns the response to the client.

“Use encapsulation forwarding to forward traffic across network segments” on page 154

Use encapsulation forwarding when the back-end server is not located on the same network segment or if you are using virtualization technology and need to forward packets that are otherwise unable to be forwarded.

Related reference:

“List of advisors” on page 114

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

“Sample scripts to generate alerts and record server failure” on page 174

Load Balancer provides user exits that trigger scripts that you can customize. You can create the scripts to perform automated actions, such as alerting an Administrator when servers are marked down by the manager or simply record the event of the failure.

Content Based Routing (CBR) component

CBR works with Caching Proxy to proxy client requests to specified HTTP or HTTPS (SSL) servers. Use CBR to manipulate caching details for faster web document retrieval with low network bandwidth requirements. CBR and Caching Proxy examines HTTP requests by using specified rule types.

CBR gives you the ability to specify a set of servers that handle a request that is based on regular expression matching of the content of the request. Because CBR can specify multiple servers for each type of request, the requests can be load balanced for optimal client response. CBR also detects when one server in a set fails, and stops routing requests to that server. The load-balancing algorithm that is used by the CBR component is identical to the proven algorithm used by the Dispatcher component.

When a request is received by Caching Proxy, it is checked against the rules that are defined in the CBR component. If a match is found, then one of the servers that are associated with that rule is chosen to handle the request. Caching Proxy then continues its normal processing to proxy the request to the chosen server.

CBR has the same function as Dispatcher, except for high availability and a few other configuration commands.

Caching Proxy must be installed, configured, and running on the same machine before CBR can begin load balancing client requests.

With the CBR component, you can provide the following advantages to your clients:

- Load balance client requests for different types of content to sets of servers.
- Improve response time by optimally dividing your site's content among your web servers.
- Ensure uninterrupted client traffic during server failure by allowing multiple servers to be assigned to each type of content.

Features of Content Based Routing:

- Collocation
- CBR with multiple instances of Caching Proxy
- Provide content-based routing for SSL connections
- Server partitioning
- Rules-based load balancing
- Client to server affinity
- High availability
- Binary logging
- Alerts

Functions that provide load balancing

The primary functions of CBR interact with each other and Caching Proxy to balance network traffic in your environment. CBR consists of the following functions:

- The **cbrserver** command handles requests from the command line to the executor, manager, and advisors.
- The executor supports content-based load balancing of HTTP, HTTPS, and FTP connections. The executor always runs when the CBR component is being used for load balancing.
- The manager sets weights that are used by the executor, which is based on:
 - Internal counters in the executor
 - Feedback from the servers that are provided by the advisors

- Feedback from a system-monitoring program, such as Metric Server or WLM.
- Using the manager is optional; however, if the manager is not used, load balancing is performed by using weighted round-robin scheduling that is based on the current server weights, and advisors are not available
- The advisors query the servers and analyze results by protocol before it calls the manager to set weights as appropriate. Currently there are advisors available for the following protocols: HTTP, HTTPS, and FTP. You also have the option of writing your own advisors (read the Creating a custom advisor topic). Using the advisors is optional but recommended.

Configure and manage the executor, advisors, and manager, by using the command line (**cbrcontrol**) or the graphical user interface (**lbadmin**).

The three key functions of CBR (executor, manager, and advisors) and Caching Proxy interact to balance and dispatch the incoming requests between servers. Along with load balancing requests, the executor monitors the number of new connections, active connections, and connections in a finished state. The executor also does garbage collection of completed or reset connections and supplies this information to the manager.

The manager and advisor functions are similar to that of the Dispatcher.

Managing local servers with CBR

The CBR component uses Caching Proxy to forward client requests (HTTP or HTTPS) to the servers based on the content of the URL. This figure shows a logical representation of a site in which CBR is being used to proxy some content from local servers.

Related concepts:

Chapter 3, “Product overview,” on page 9

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

“Plan for Content Based Routing” on page 95

This topic describes what the network planner might consider when installing and configuring the CBR component with Caching Proxy.

Related tasks:

“Creating a custom advisor” on page 125

A custom advisor is a small piece of Java code, provided as a class file, that is called by the Load Balancer base code to determine the load on a server. The base code provides all necessary administrative services, including starting and stopping an instance of the custom advisor, providing status and reports, recording history information in a log file, and reporting advisor results to the manager component.

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Site Selector component

Site Selector acts as a name server that works with other name servers in a domain name system to load balance among a group of servers by using measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers based on the domain name used for a client's request.

A client submits a request for resolution of a domain name to a name server within its network. Name server forwards the request to the Site Selector machine. Site Selector then resolves the domain name to the IP address of one of the servers that is configured under the site name. Site Selector returns the IP address of the selected server to the name server. The name server returns the IP address to the client.

Metric Server is a system monitoring component of Load Balancer that must be installed in each load-balanced server within your configuration. Using Metric Server, Site Selector can monitor the level of activity on a server, detect when a server is the least heavily loaded, and detect a failed server. The load is a measure of how hard the server is working. By customizing system metric script files, you can control the type of measurements that are used to measure the load. You can configure Site Selector to suit your environment, considering such factors as frequency of access, the total number of users, and types of access (for example, short queries, long-running queries, or CPU-intensive loads).

Features of Site Selector:

- Collocation
- High availability
- Client to server affinity
- Rules-based load balancing

- Wide area load balancing
- Alerts

Functions that provide load balancing

The primary functions of Site Selector interact with each other and your server configuration to balance network traffic in your environment. Site Selector consists of the following functions:

- The **ssserver** command handles requests from the command line to the name server, manager, and advisors.
- The **nameserver** command works with other name servers in a domain name system to load balance among a group of servers. The name server always runs when the Site Selector component is being used for load balancing.
- The manager sets weights that are used by the executor, which is based on:
 - Internal counters in the executor
 - Feedback from the servers that are provided by the advisors
 - Feedback from a system-monitoring program, such as Metric Server or WLM.
- The advisors query the servers and analyze results by protocol before it calls the manager to set weights as appropriate. To configure and manage the name server, advisors, and manager, use the command line (**sscontrol**) or the graphical user interface (**1admin**).

Managing local and remote servers with Site Selector and Metric Server

A client submits a request for resolution of a domain name to a client name server. The client name server forwards the request through the DNS to the Site Selector machine (Path 1). Site Selector then resolves the domain name to the IP address of one of the servers. Site Selector returns the IP address of the selected server to the client name server. The name server returns the IP address to the client. After the client receives the IP address of the server, the client routes application requests directly to the selected server (Path 2).

Important: In this example, the Metric Server provides system load information to the Site Selector machine. The Metric Server agent is installed on each backend server. Use Metric Server with Site Selector; otherwise Site Selector can use only a round-robin selection method for load balancing.

This figure illustrates a site in which the Site Selector component is used to answer requests. Server1, Server2, and Server3 are local. Server4, Server5, and Server6 are remote.

Related concepts:

Chapter 3, “Product overview,” on page 9

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

“Plan for Site Selector” on page 85

Site Selector works with a domain name server to load balance among a group of servers that uses measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers that are based on the domain name that is used for a client request.

Related tasks:

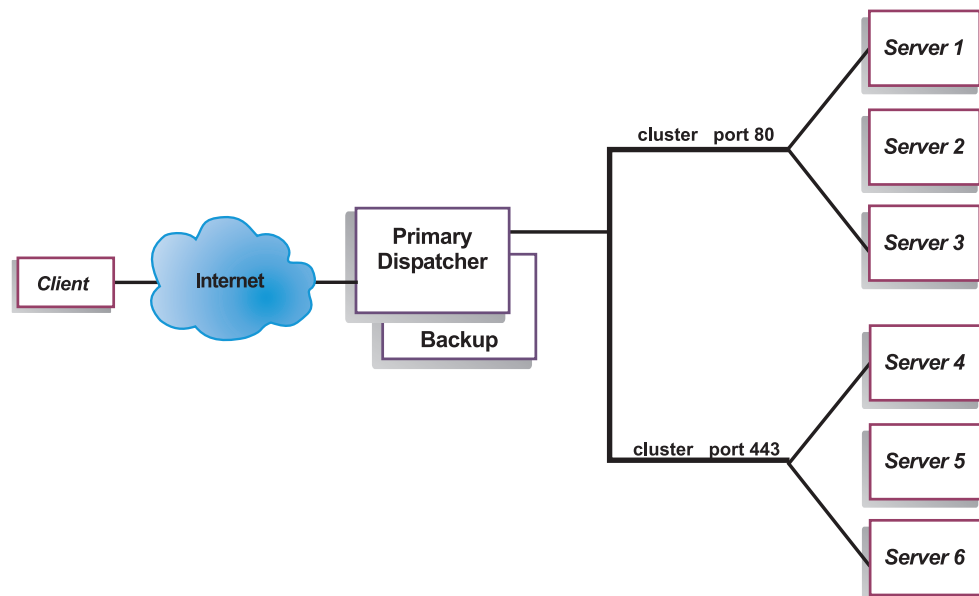
“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

High availability with Load Balancer

The Dispatcher component offers a built-in high availability feature, eliminating Dispatcher as a single point of failure from your network. This feature involves the use of a second Dispatcher machine that monitors the main, or primary, machine and stands by to take over the task of load balancing should the primary machine fail at any time.

Functioning in conjunction with content hosts, such as WebSphere Application Server, the Load Balancer Dispatcher component enables you to enhance your network’s availability and scalability. Load Balancer is used by enterprise networks and is installed between the Internet and the enterprise’s back end servers.

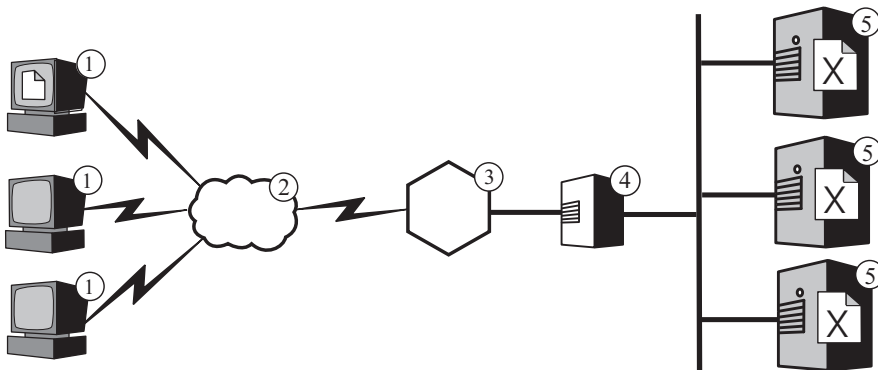


Load Balancer acts as the enterprise's single point-of-presence on the Internet, even if the enterprise uses multiple back end servers because of high demand or a large amount of content. Availability is achieved through load balancing multiple content hosts and failover support.

Load balancing multiple content hosts

You can satisfy high demand by duplicating content on multiple hosts, but then you need a way to balance the load among them. Domain Name Service (DNS) can provide basic round-robin load balancing, but there are several situations in which it does not perform well.

A more sophisticated solution for load balancing multiple content hosts is to use the Dispatcher component that is shown in the Load balancing multiple content hosts figure.



. Load balancing multiple content hosts

Legend:

- 1 - Client
- 2 - Internet
- 3 - Router/Gateway
- 4 - Dispatcher
- 5 - Content host

In this configuration, all of the content hosts (the machines marked 5) store the same content. They are defined to form a load-balanced cluster, and one of the network interfaces of the Load Balancer machine (4) is assigned a host name and IP address dedicated to the cluster. When an end user working on one of the machines marked 1 requests file X, the request crosses the Internet (2) and enters the enterprise's internal network through its Internet gateway (3). The Dispatcher intercepts the request because its URL is mapped to the Dispatcher's host name and IP address. The Dispatcher determines which of the content hosts in the cluster is currently best able to service the request, and forwards the request to that host, which returns file X directly to the client (that is, file X does not pass through Load Balancer).

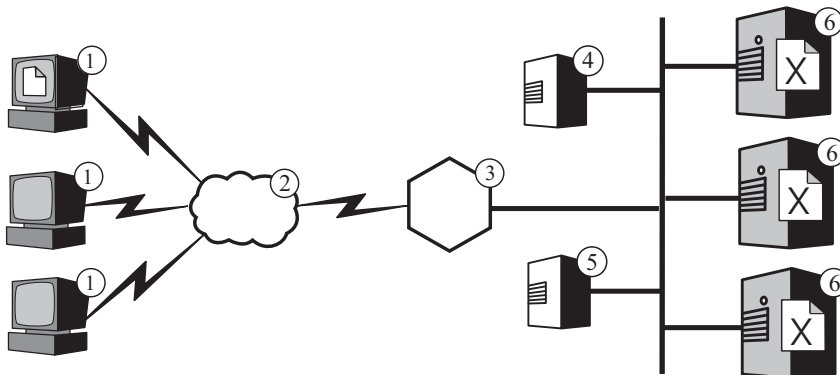
By default, the Dispatcher uses weighted round-robin load balancing, and it addresses many of DNS's inadequacies. Unlike DNS, it tracks whether a content host is unavailable or inaccessible and does not continue to direct clients to an unavailable content host. Further, it considers the current load on the content hosts by tracking new, active, and finished connections. You can further optimize load

balancing by activating Load Balancer's optional advisor and manager components, which track a content host's status even more accurately and incorporate the additional information into the load-balancing decision process. The manager enables you to assign different weights to the different factors used in the decision process, further customizing load balancing for your site.

Failover support

Load Balancer acts as a single point-of-presence for your enterprise's content hosts. This is beneficial because you advertise the cluster host name and address in DNS, rather than the host name and address of each content host, which provides a level of protection against casual attacks and provides a unified feel for your enterprise's web site. To further enhance web site availability, configure another Load Balancer to act as a backup for the primary Load Balancer, as shown in the figure that follows. If one Load Balancer fails or becomes inaccessible due to a network failure, end users can still reach the content hosts.

Legend: 1--Client 2--Internet 3--Router/Gateway 4--Primary Dispatcher 5--Backup



. Using a primary and backup Load Balancer to make Web content highly available

Dispatcher 6--Content host

Legend:

- 1 - Client
- 2 - Internet
- 3 - Router/Gateway
- 4 - Primary Dispatcher
- 5 - Backup Dispatcher
- 6 - Content host

In the normal case, a browser running on one of the machines marked 1 directs its request for a file X to the cluster host name that is mapped to the primary Load Balancer (4). The Dispatcher routes the request to the content host (6) selected on the basis of the Dispatcher's load-balancing criteria. The content host sends file X directly to the browser, routing it through the enterprise's gateway (3) across the Internet (2) but bypassing Load Balancer. The backup Dispatcher (5) does not perform load balancing as long as the primary one is operational. The primary and backup Dispatchers track each other's status by periodically exchanging messages called heartbeats. If the backup Dispatcher detects that the primary has failed, it automatically takes over the responsibility for load balancing by intercepting requests directed to the primary's cluster host name and IP address.

Related concepts:

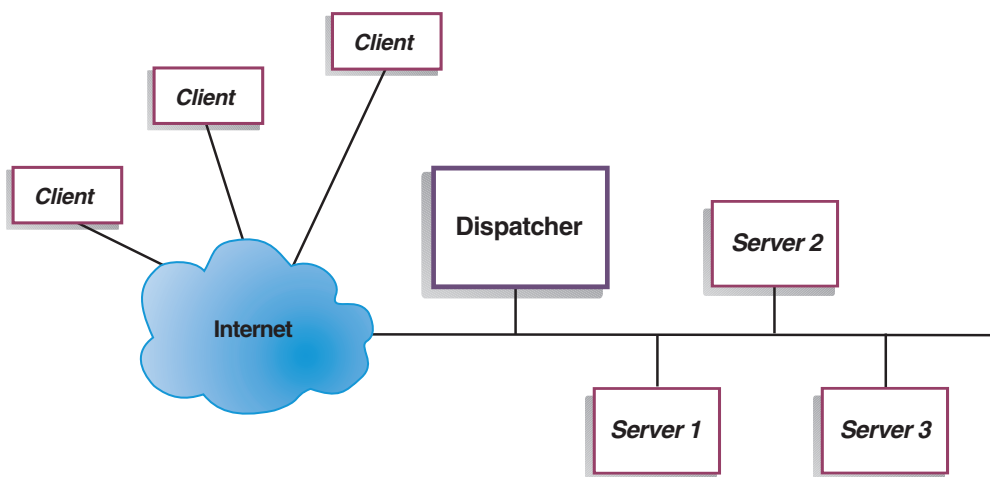
“Managing servers”

You can load balance traffic to existing server topologies without changing the physical configuration of the machines or how clients will connect to your site.

Managing servers

You can load balance traffic to existing server topologies without changing the physical configuration of the machines or how clients will connect to your site.

The figure below shows a physical representation of the site using an Ethernet network configuration.

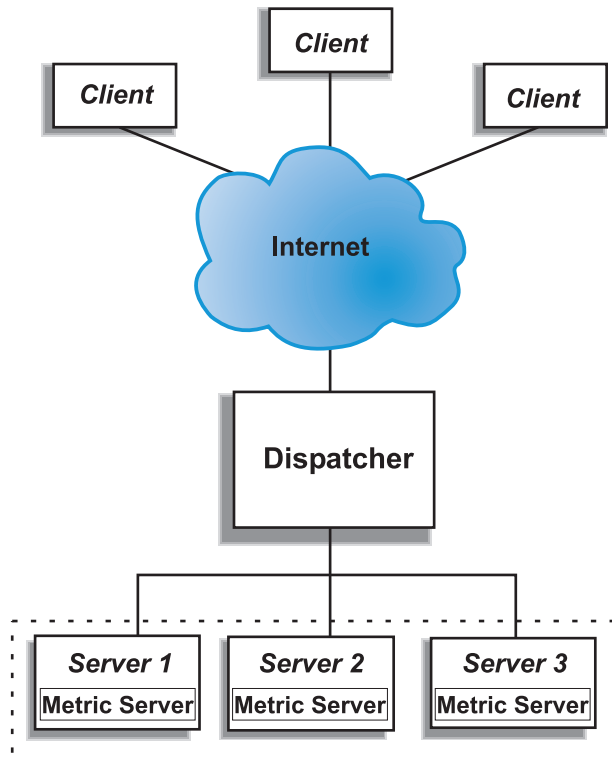


The Dispatcher machine can be installed without making any physical changes to the network. After a client request is directed to the optimal server by the Dispatcher, the response is then sent directly from server to client with no involvement by the Dispatcher.

Read “Types of cluster, port, and server configurations for Dispatcher” on page 26 for examples of the different types of configurations you can use with Load Balancer.

Managing servers with Load Balancer and Metric Server

The figure below illustrates a site in which all servers are on a local network. The Dispatcher component is used to forward requests, and the Metric Server is used to provide system load information to the Dispatcher machine.



In this example, the Metric Server daemon is installed on each back end server. You can use Metric Server with the Dispatcher component.

For more information on Metric Server refer to “Getting advice with Metric Server” on page 121.

Related concepts:

Chapter 3, “Product overview,” on page 9

Load Balancer is a software solution for distributing incoming client requests across servers. It boosts the performance of servers by directing TCP/IP session requests to different servers within a group of servers; in this way, it balances the requests among all the servers. This load balancing is transparent to users and other applications. Load Balancer is useful for applications such as e-mail servers, World Wide Web servers, distributed parallel database queries, and other TCP/IP applications.

“High availability with Load Balancer” on page 21

The Dispatcher component offers a built-in high availability feature, eliminating Dispatcher as a single point of failure from your network. This feature involves the use of a second Dispatcher machine that monitors the main, or primary, machine and stands by to take over the task of load balancing should the primary machine fail at any time.

Related tasks:

“Getting advice with Metric Server” on page 121

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

Related reference:

“Types of cluster, port, and server configurations for Dispatcher”

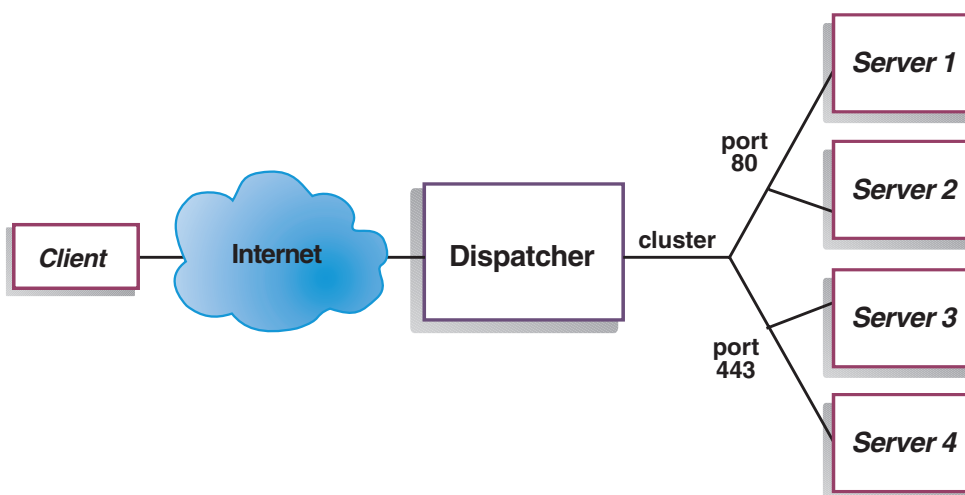
There are many ways that you can configure Dispatcher to support your site.

Types of cluster, port, and server configurations for Dispatcher

There are many ways that you can configure Dispatcher to support your site.

1 cluster with 2 ports

If you have only one host name for your site to which all of your customers will connect, you can define a single cluster of servers. For each of these servers, configure a port through which Load Balancer communicates.

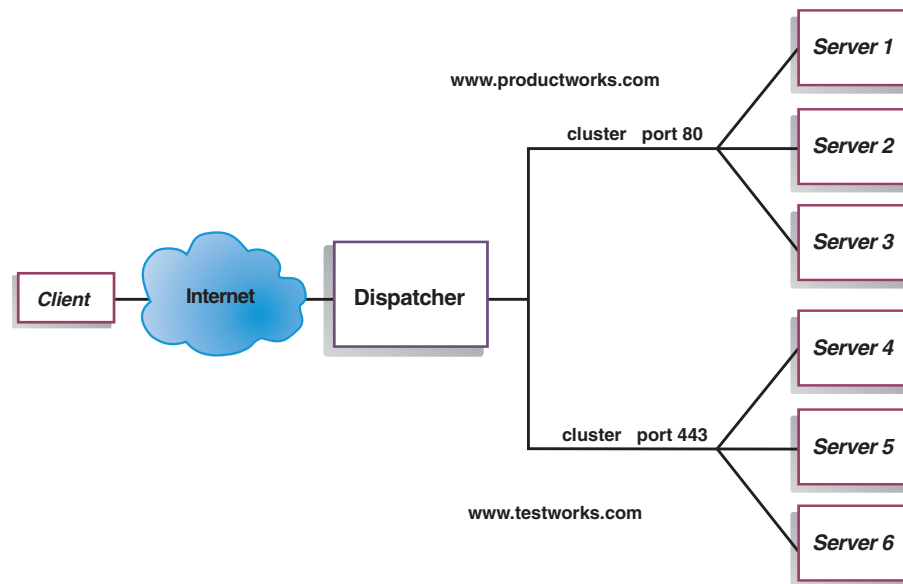


In this example for the Dispatcher component, one cluster is defined at www.productworks.com. This cluster has two ports: port 80 for HTTP and port 443 for SSL. A client making a request to <http://www.productworks.com> (port 80)

goes to a different server than a client requesting `https://www.productworks.com` (port 443).

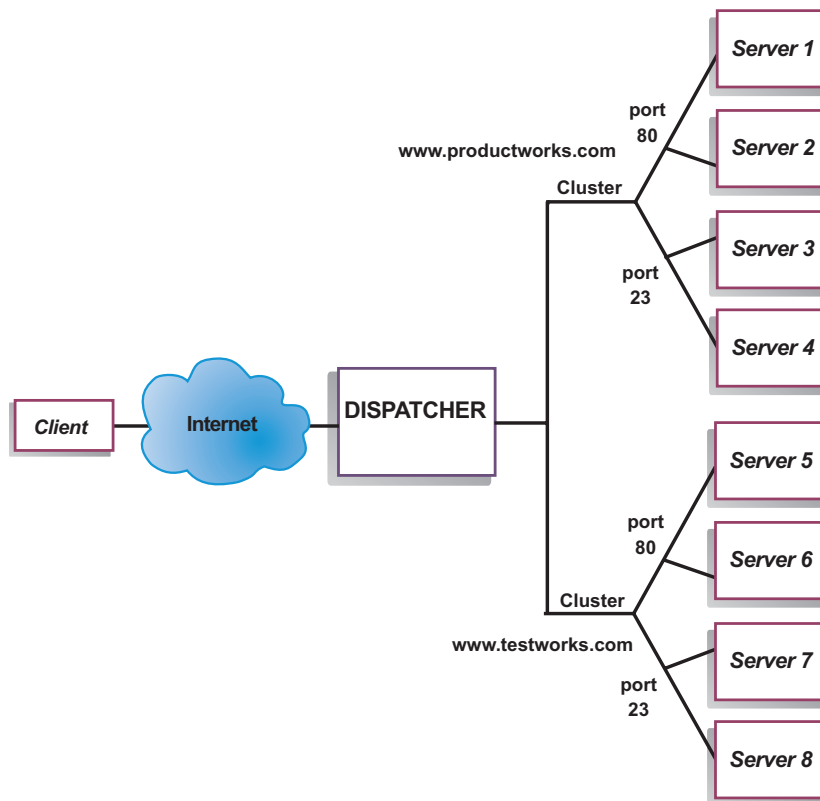
2 clusters, each with 1 port

Another way of configuring Load Balancer might be appropriate if you have a very large site with many servers dedicated to each protocol supported. In this case, you might want to define a cluster for each protocol with a single port but with many servers.



In this example for the Dispatcher component, two clusters are defined: `www.productworks.com` for port 80 (HTTP) and `www.testworks.com` for port 443 (SSL). A third way of configuring Load Balancer might be necessary if your site does content hosting for several companies or departments, each one coming into your site with a different URL. In this case, you might want to define a cluster for each company or department and then define any ports to which you want to receive connections at that URL, as shown in the configuration for 2 clusters, each with two ports.

2 clusters, each with 2 ports



In this example for the Dispatcher component, two clusters are defined with port 80 for HTTP and port 23 for Telnet for each of the sites at www.productworks.com and www.testworks.com.

Related tasks:

“Quick start configuration” on page 61

This quick start example shows how to configure three locally attached workstations to load-balance web traffic between two web servers.

Chapter 4. Installing Load Balancer

This section describes how to establish load balancing capability in new and existing environments. The information includes planning, preparing for, completing, and maintaining product installations.

Setting up IBM Installation Manager

Install and configure IBM Installation Manager so that you can download and install the product and updates.

About this task

There are two main steps for this task:

- Install IBM Installation Manager on your system.
- Configure Installation Manager to install and maintain the Load Balancer product.

Procedure

- **Install IBM Installation Manager on your system.**

If you already have IBM Installation Manager installed, you can proceed to Configure Installation Manager to install and maintain the Load Balancer product..

1. Obtain the necessary files for your system.

Note: Be aware of the following:

- The downloadable version of Installation Manager is preconfigured to use the remote repositories. If you use the local repositories that are included in the download, you must configure the location of the repositories in Installation Manager.
- If you want to access the most recent Load Balancer installation files, use the remote web-based repositories.
- If you use local repositories, you can also leave in the preconfigured remote repositories, because they will provide post-shipment service updates when you are connected to the internet. Deselect the remote repositories in the Repositories panel of IBM Installation Manager if you use the local repositories.

There are four options for obtaining the files and installing the Installation Manager and Load Balancer product:

- Use the copy of Installation Manager and the product files from your product disk:
 - a. Insert the product disk for Websphere Load Balancer for IPv4 and IPv6, and copy the product repository to a temporary location in your system.
 - b. Insert the Installation Manager Disk and proceed to the installation step.
- Download a file from the Passport Advantage® site and use web-based or local installation:

- a. If you are a licensed client, download and unpack a compressed file that contains the preconfigured Installation Manager and the necessary product repositories.
 - b. Install Installation Manager on your local system and use it to install Load Balancer from either the web-based repositories or the unpacked local product repositories.
- Download a file from the Installation Manager for WebSphere web site, and perform a web-based installation:
 - a. Download and unpack a compressed file that contains the Installation Manager that is appropriate for your operating system from the IBM Installation Manager download website.
 - b. Install Installation Manager on your local system and use it to install the product from the web-based repositories.
 - Download files from both the Installation Manager for WebSphere web site and the product web site, and perform a local installation:
 - a. Download and unpack the compressed files that contain the Installation Manager and the product repositories from the product web site.
 - b. Install Installation Manager on your local system and use it to install the product from the unpacked local product repositories.
2. Unpack the compressed file that contains Installation Manager into a temporary folder.
Use a folder name like IM_temp.
 3. Run the install command for IBM Installation Manager.
 - For an interactive installation, run the following command as an administrator from the temporary directory:


```
./install
install.exe
```

The installer opens the Install Packages window.
 - To install IBM Installation Manager silently, run the following command as an administrator from the temporary directory:


```
./installc --launcher.ini silent-install.ini -log log_file_path_and_name
installc --launcher.ini silent-install.ini -log log_file_path_and_name
```

After installation is complete, configure your repositories. For more information on installing Installation Manager silently, see the topic on silently installing packages using the Installation Manager installer that is in the information center for IBM Installation Manager.
 4. Ensure that the latest version of the Installation Manager package is selected, and click **Next**.
 5. Accept the terms in the license agreements, and click **Next**.
The program creates the directory for your installation.
 6. Click **Next**.
 7. Review the summary information, and click **Install**. Review the results:
 - If the installation is successful, the program displays a message indicating that the installation is successful.
 - If the installation is not successful, click **View Log File** to troubleshoot the problem.

8. If you are using local repositories to install and maintain the product, configure each local repository in Installation Manager.

If you are using the web-based repositories to install and maintain the product, go to the topic on “Installing the product with the GUI” on page 36 or “Installing the product silently” on page 38. Otherwise, perform the following actions:

- a. Unpack the compressed file containing the repository to a directory on your system.
- b. Add the local repository to your Installation Manager preferences:
 - 1) In the top menu, click **File > Preferences**.
 - 2) Select **Repositories**.
 - 3) Click **Add Repository**.
 - 4) Enter the path to the repository.config file in the local directory into which you unpacked the repository files. For example, enter the following:
`/var/repositories/edge85/local-repositories`
`C:\repositories\edge85\local-repositories`
- c. Click **OK**.
- d. Click **Apply**.
- e. Click **OK**.
- f. Close Installation Manager. Click **File > Exit**.

- **Configure Installation Manager to install and maintain the Load Balancer product.**

Note: Be aware of the following:

- If you want to access the most recent installation files for the Load Balancer product, then use the remote web-based repositories.
 - If you use local repositories, you can also leave in the preconfigured remote repositories, because they will provide post-shipment service updates when you are connected to the internet. Deselect the remote repositories in the Repositories panel of IBM Installation Manager if you use the local repositories.
1. Download the necessary compressed files from the web to your system.
There are three options for downloading and installing the product:
 - Download a file from the Passport Advantage site and use a web-based or local installation:
 - a. If you are a licensed client, download and unpack a compressed file containing the necessary product repositories.
 - b. Use Installation Manager to install the product from either the web-based repositories or the unpacked local product repositories.
 - Use Installation Manager to install the product from the web-based repositories.
 - Download files from the product web site, and use a local installation:
 - a. Download the compressed files that contain the product repositories from the product web site.
 - b. Use Installation Manager to install the product from the unpacked local product repositories.
 2. Add the web-based repositories or configure the local repositories.
Perform one of the following actions:

- If you are using the web-based repositories, add them to your Installation Manager preferences.
 - a. In the top menu, click **File > Preferences**.
 - b. Select **Repositories**.
 - c. Add the product repository.
 - 1) Click **Add Repository**.
 - 2) Enter the URL for the product repository. The product repository for Edge Components is:
`http://www.ibm.com/software/repositorymanager/V85WASNDEdge`
 - 3) Click **OK**.
 - d. Add the service repository.
 - 1) Click **Add Repository**.
 - 2) Enter the URL for the service repository: The service repository might or might not be the same as the product repository.
 - 3) Click **OK**.
 - e. Deselect any locations that are listed in the Repositories panel that you will not use.
 - f. Click **Apply**.
 - g. Click **OK**.
 - h. Close Installation Manager. Click **File > Exit**.
 - i. Install the product using the GUI, or install the product silently.
- If you are using local repositories to install and maintain the product, configure each local repository in Installation Manager.
 - a. Unpack the compressed file containing the repository to a directory on your system.
 - b. Add the local repository to your Installation Manager preferences.
 - 1) In the top menu, click **File > Preferences**.
 - 2) Select **Repositories**.
 - 3) Click **Add Repository**.
 - 4) Enter the path to the repository.config file in the local directory into which you unpacked the repository files. For example:
`/var/repositories/edge85/local-repositories`
`C:\repositories\edge85\local-repositories`
 - c. Click **OK**.
 - d. Deselect locations in the Repositories panel that you do not want to use.
 - e. Click **Apply**.
 - f. Click **OK**.
 - g. Close Installation Manager. Click **File > Exit**.

Related tasks:

“Installing the product with the GUI” on page 36

Use IBM Installation Manager graphical user interface (GUI) to install the Load Balancer product.

“Installing the product silently” on page 38

Use IBM Installation Manager from the command line to install Websphere Load Balancer for IPv4 and IPv6 silently.

Related reference:

 Installing packages silently by using the installer

Installing the product on distributed operating systems

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Before you begin

Note: Installation Manager is a single installation program that can use remote or local software flat-file repositories to install, modify, or update new WebSphere Application Server products. It determines and shows available packages—including products, fix packs, interim fixes, and so on—checks prerequisites and interdependencies, and installs the selected packages. You also use Installation Manager to easily uninstall the packages that it installed.

Restrictions:

- Do not use the same response files that are used with previous versions of Websphere Load Balancer for IPv4 and IPv6 to install or uninstall Version 8.5 silently. Use response files that are based on Installation Manager to install, update, or uninstall Version 8.5 and later.
- When you install a product using Installation Manager with local repositories, the installation takes a significantly longer amount of time if you use a compressed repository file directly without extracting it.
Before you install a product using local repositories, extract the compressed repository file to a location on your local system before using Installation Manager to access it.
- Installation Manager console mode, which is included in Installation Manager Version 1.4.3 and later, does not work with WebSphere Application Server Version 8.5 offerings on systems other than z/OS®.

Important: Do not transfer the content of a repository in non-binary mode and do not convert any content on extraction.

Tip: Although almost all of the instructions in this section of the information center will work with earlier versions of IBM Installation Manager, the information here is optimized for users who have installed or upgraded to Installation Manager Version 1.5 or later.

Important: Installation Manager can install any fix-pack level of the product directly without installing the interim fix packs; in fact, Installation Manager installs the latest level by default. For example, you can skip fix-pack levels and go from Version 8.5.0.1 directly to Version 8.5.0.5. Keep in mind, however, that later you can not roll back to any level that was skipped. If you directly install to Version 8.5.0.5, for example, you cannot roll back to Version 8.5.0.4. If you skip

from Version 8.5.0.1 to Version 8.5.0.5, you can only roll back to Version 8.5.0.1. You should plan your installations accordingly.

About this task

Perform one of these procedures to install or uninstall the product using Installation Manager.

Procedure

- “Installing the product with the GUI” on page 36.
- “Installing the product silently” on page 38.
- “Uninstalling Load Balancer from distributed operating systems using the GUI” on page 45
- “Uninstalling Load Balancer from distributed operating systems using response files” on page 45
- “Uninstalling Load Balancer if it was installed without Installation Manager” on page 48

Results

Notes on logging and tracing:

- One way to view the logs is to open Installation Manager and go to **File > View Log**. You can open an individual log file by selecting it in the table and then clicking the **Open log file** icon.
- Logs are located in the logs directory of Installation Manager's application data location. For example:

- **Administrative installation:**

```
C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager
```

- **Administrative installation:**

```
/var/ibm/InstallationManager
```

- The main log files are time-stamped XML files in the logs directory, and they can be viewed using any standard web browser.
- The log.properties file in the logs directory specifies the level of logging or tracing that Installation Manager uses. To turn on tracing for the Websphere Load Balancer for IPv4 and IPv6 plug-ins, for example, create a log.properties file with the following content:

```
com.ibm.ws=DEBUG
com.ibm.cic.agent.core.Engine=DEBUG
global=DEBUG
```

Restart Installation Manager as necessary, and Installation Manager outputs traces for the Websphere Load Balancer for IPv4 and IPv6 plug-ins.

Notes on troubleshooting:

- When you attempt to launch Installation Manager from a DVD that was mounted using the CD-ROM file system (CDFS) on an HP-UX operating system, it might fail to launch and point to a log file that contains an exceptions similar to one of the following:

```
java.util.zip.ZipException: Exception in opening zip file:
org.osgi.framework.BundleException: Exception in
org.eclipse.update.internal.configurator.ConfigurationActivator.start()
or bundle org.eclipse.update.configurator.
```

This issue might be caused by Installation Manager reaching the upper limit of number of descriptors that can be opened on a CDFS-mounted device. This upper limit is determined by the value for the tunable kernel parameter `ncdnode`, which specifies the maximum number of CDFS nodes that can be in memory at any given time. To resolve the problem, change the `ncdnode` system kernel setting to 250. If the problem persists, increase the setting.

- By default, some HP-UX systems are configured to not use DNS to resolve host names. This could result in Installation Manager not being able to connect to an external repository.

You can ping the repository, but `nslookup` does not return anything.

Work with your system administrator to configure your machine to use DNS, or use the IP address of the repository.

- In some cases, you might need to bypass existing checking mechanisms in Installation Manager.
 - On some network file systems, disk space might not be reported correctly at times; and you might need to bypass disk-space checking and proceed with your installation.

To disable disk-space checking, specify the following system property in the `config.ini` file in `IM_install_root/eclipse/configuration` and restart Installation Manager:

```
cic.override.disk.space=sizeunit
```

where *size* is a positive integer and *unit* is blank for bytes, k for kilo, m for megabytes, or g for gigabytes. For example:

```
cic.override.disk.space=120 (120 bytes)
cic.override.disk.space=130k (130 kilobytes)
cic.override.disk.space=140m (140 megabytes)
cic.override.disk.space=150g (150 gigabytes)
cic.override.disk.space=true
```

Installation Manager will report a disk-space size of `Long.MAX_VALUE`. Instead of displaying a very large amount of available disk space, N/A is displayed.

- To bypass operating-system prerequisite checking, add `disableOSPrereqChecking=true` to the `config.ini` file in `IM_install_root/eclipse/configuration` and restart Installation Manager.

If you need to use any of these bypass methods, contact IBM Support for assistance in developing a solution that does not involve bypassing the Installation Manager checking mechanisms.

- For more information on using Installation Manager, read the IBM Installation Manager Information Center.

Read the release notes to learn more about the latest version of Installation Manager. To access the release notes, complete the following task:

- Go to the documentation subdirectory in the directory where Installation Manager is installed, and open the `readme.html` file.
- Click **Start > Programs > IBM Installation Manager > Release Notes**.

- If a fatal error occurs when you try to install the product, take the following steps:
 1. Make a backup copy of your current product installation directory in case IBM support needs to review it later.

2. Use Installation Manager to uninstall everything that you have installed under the product installation location (package group). You might run into errors, but they can be safely ignored.
3. Delete everything that remains in the product installation directory.
4. Use Installation Manager to reinstall the product to the same location or to a new one.

Related tasks:

“Installing the product with the GUI”

Use IBM Installation Manager graphical user interface (GUI) to install the Load Balancer product.

“Installing the product silently” on page 38

Use IBM Installation Manager from the command line to install Websphere Load Balancer for IPv4 and IPv6 silently.

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

“Uninstalling Load Balancer from distributed operating systems using response files” on page 45

Uninstall Load Balancer before upgrading to a newer version or if you think that the current installation is corrupted. You can use IBM Installation Manager to uninstall Load Balancer using response files.

“Uninstalling Load Balancer if it was installed without Installation Manager” on page 48

If you did not use IBM Installation Manager to install the product, you cannot use Installation Manager to uninstall it. Uninstall the product using the packaging and installation tools on your system.

Installing the product with the GUI

Use IBM Installation Manager graphical user interface (GUI) to install the Load Balancer product.

Before you begin

For information on hardware and software requirements, including supported browsers, refer to the WebSphere Application Server detailed system requirements.

If you have not installed or configured IBM Installation Manager, refer to “Setting up IBM Installation Manager” on page 29.

Two versions of Load Balancer cannot be installed on the same machine. If you have an earlier version installed, uninstall that copy before installing the current version. Refer to “Uninstalling the product” on page 44 for more information.

Procedure

1. Log in as a user root or an administrator, and start Installation Manager.
2. Click **Install**.

If you are prompted to authenticate, use the IBM ID and password that you used to register on the program web site. Installation Manager will search its defined repositories for available packages.

3. Select **IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6** and the most recent version number.

Note: If you already installed the product using Installation Manager, Installation Manager displays a message that you are not allowed to install more than one instance in a system.

4. Click **Next**.

Note: If you installed a previous version of the product using installers other than Installation Manager, Installation Manager will not allow you to proceed with the installation. Uninstall the older version of the product, and proceed with the installation using Installation Manager. Refer to “Uninstalling Load Balancer if it was installed without Installation Manager” on page 48 if you have installed earlier versions using installers other than IBM Installation Manager.

5. Accept the terms in the license agreements, and click **Next**.

6. Review the installation location.

The Installation Location panel appears. This panel also displays the shared resources directory and disk-space information.

- You are not allowed to change the default installation location. The folder for the installation location on the system should be empty before you proceed.
- Specify the installation root directory for the product binary files.

Note: Be aware of the following restrictions:

- You cannot continue if you delete the default target location and leave any installation-directory fields empty.
- Do not use symbolic links as the destination directory.
- Do not use a semicolon in the directory name.
- The maximum path length on the Windows 2000, Windows Server 2008, Windows Vista, and Windows 7 operating systems is 60 characters.

7. Click **Next**.

8. Select the languages for which translated content should be installed, and click **Next**.

English is always selected.

9. Select any features that you want to install, and click **Next**.

Choose from the following optional features:

Dispatcher

The Dispatcher component performs load balancing for servers within a local area network (LAN) or wide area network (WAN) for all Internet services, such as HTTP, FTP, HTTPS, and Telnet. For HTTP services, Dispatcher can perform load balancing of servers based on the URL content of the client request.

Metric Server

The Metric Server component runs as a daemon on a load-balanced server. Metric Server:

- Provides information about system loads to the Load Balancer components
- Reports on the health of the servers

Content Based Routing

Content Based Routing proxies requests to specific servers depending on the content of the request for HTTP and SSL traffic.

Site Selector

Site Selector works with a name server to load balance among a group of servers.

Tip: You can run the Installation Manager later to modify this installation and add or remove features.

10. If you are using a 64-bit operating system that has a corresponding 32-bit image for the product, choose to install the 32-bit version of the product or the 64-bit version and click **Next**.

This panel displays only if you are using a 64-bit operating system.

11. Review the summary information, and click **Install**.

Review the results:

- If the installation is successful, the program displays a message indicating that the installation is successful. The program might also display important post-installation instructions.
- If the installation is not successful, click **View Log File** to troubleshoot the problem.

12. Click **Finish**.

13. Close Installation Manager.

Click **File > Exit**.

What to do next

The installation process does not add the command directories for Load Balancer into the PATH environment variable. To run Load Balancer commands from the system root, add the command directories to the PATH environment variable.

Note: If you installed previous installation of the Dispatcher component, be aware that Websphere Load Balancer for IPv4 versions of the Dispatcher component used commands in the `/usr/bin` directory, which might be in the PATH variable. Websphere Load Balancer for IPv4 and IPv6 has commands in the `install_root/bin` directory, so ensure that the directory entries point to the appropriate directories for the **dsserver** and **dscontrol** commands.

Related tasks:

“Installing the product silently”

Use IBM Installation Manager from the command line to install Websphere Load Balancer for IPv4 and IPv6 silently.

“Installing fix packs” on page 52

Use IBM Installation Manager to update the Load Balancer product to a later version.

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

Related information:

WebSphere Application Server detailed system requirements

Installing the product silently

Use IBM Installation Manager from the command line to install Websphere Load Balancer for IPv4 and IPv6 silently.

Before you begin

For information on hardware and software requirements, including supported browsers, refer to the following Web page: WebSphere Application Server detailed system requirements.

If you have not installed or configured IBM Installation Manager, refer to “Setting up IBM Installation Manager” on page 29.

Two versions of Load Balancer cannot be installed on the same machine. If you have an earlier version installed, uninstall that copy before installing the current version. Refer to “Uninstalling the product” on page 44 for more information.

About this task

With IBM Installation Manager, you can work with response files to install the product silently in a variety of ways. You can:

- Generate a new response file by hand or by taking an example and modifying it.
- Record a response file using the GUI as described in the following procedure.

Procedure

1. Optional: Record a response file to install the product.

On one of your systems, perform the following actions to record a response file that will install the product.

- a. From a command line, navigate to the directory in which you installed Installation Manager, and then change directories to the `eclipse` subdirectory.
- b. Start Installation Manager from the command line with the `-record` option.

For example, run the following command as an administrator:

```
./IBMIM -skipInstall /var/temp/imRegistry -record /var/temp/install_response_file.xml  
IBMIM.exe -skipInstall "C:\temp\imRegistry" -record C:\temp\install_response_file.xml
```

Note: When you record a new response file, specify the `-skipInstall` parameter. This parameter streamlines the installation process.

- No files are actually installed, which speeds up the recording.
- If you use a temporary data location with the `-skipInstall` parameter, Installation Manager writes the installation registry to the specified data location while recording. When you start Installation Manager again without the `-skipInstall` parameter, you then can use your response file to install the files to the real installation registry.

The `-skipInstall` operation should not be used on the actual agent data location that Installation Manager will use. Use a clean, writable location, and reuse that location for future recording sessions.

For more information on recording a response file, read the topic on recording a response file with Installation Manager in the Installation Manager information center.

- c. Add the appropriate repositories to your preferences.
 - 1) In the menu, click **File > Preferences**.
 - 2) Select **Repositories**.
 - 3) Click **Add Repository**.

- 4) Enter the path to the repository.config file in the remote web-based repository or the local directory into which you unpacked the repository files.

For example, you can add one or both of the following types:

Remote repositories

<https://www.ibm.com/software/rational/repositorymanager/repositories/websphere>
<http://public.dhe.ibm.com/software/websphere/repositories>

The remote repository for Websphere Load Balancer for IPv4 and IPv6 is:

<http://www.ibm.com/software/repositorymanager/com.ibm.websphere.EDGEIPV4IPV6.v85>

Local repositories

`/var/repositories/edge85/local-repositories`

`C:\repositories\edge85\local-repositories`

- 5) Click **OK**.
- d. Click **Apply**.
- e. Click **OK**.
- f. Click **Install**. If you are prompted to authenticate, use the IBM ID and password that you registered with on the program's web site. Installation Manager will search the defined repositories for available packages.
- g. Select **IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6** and the most recent version number.

Note: If you already installed the product using Installation Manager, Installation Manager displays a message that you are not allowed to install more than one instance in a system. You cannot collocate two versions of Load Balancer.

- h. Click **Next**.

Note: If you installed a previous version of the product using installers other than Installation Manager, Installation Manager will detect it and prompt you to uninstall that version product and then proceed with the installation. Uninstall the older version of the product, and proceed with the installation using Installation Manager. Refer to “Uninstalling Load Balancer if it was installed without Installation Manager” on page 48.

- i. Accept the terms in the license agreements, and click **Next**.
- j. Review the installation location. The Installation Location panel will appear. This panel also displays the shared resources directory and disk-space information.
 - You are not allowed to change the default installation location. The folder for the installation location on the system should be empty before you proceed.
 - Specify the installation root directory for the product binaries, which are also referred to as the core product files or system files.

Note: Be aware of the following restrictions:

- You cannot continue if you delete the default target location and leave any installation-directory fields empty.
- Do not use symbolic links as the destination directory.
- Symbolic links are not supported.
- Do not use a semicolon in the directory name.

- The maximum path length on the Windows 2000, Windows Server 2008, Windows Vista, and Windows 7 operating systems is 60 characters.
- k. Click **Next**.
- l. Select the languages for which translated content should be installed, and click **Next**.
English is always selected.
- m. Select any features that you want to install, and click **Next**.
Choose from the following optional features:

Dispatcher

The Dispatcher component performs load balancing for servers within a local area network (LAN) or wide area network (WAN) for all Internet services, such as HTTP, FTP, HTTPS, and Telnet. For HTTP services, Dispatcher can perform load balancing of servers based on the URL content of the client request.

Metric Server

The Metric Server component runs as a daemon on a load-balanced server. Metric Server:

- Provides information about system loads to the Load Balancer components
- Reports on the health of the servers

Content Based Routing

Content Based Routing proxies requests to specific servers depending on the content of the request for HTTP and SSL traffic.

Site Selector

Site Selector works with a name server to load balance among a group of servers.

Tip: You can run the Installation Manager later to modify this installation and add or remove features.

- n. If you are using a 64-bit operating system that has a corresponding 32-bit image for the product, choose to install the 32-bit version of the application server or the 64-bit version, and click **Next**. This panel displays only if you are using a 64-bit operating system.
- o. Review the summary information, and click **Install**. Review the results:
 - If the installation is successful, the program will display a message indicating that the installation is successful. The program might also display important post-installation instructions.
 - If the installation is not successful, click **View Log File** to troubleshoot the problem.
- p. Click **Finish**.
- q. Close Installation Manager.
Click **File > Exit**.
- r. Optional: If you are using an authenticated remote repository, create a keyring file for silent installation.
 - 1) From a command line, navigate to the eclipse subdirectory in the directory in which you installed Installation Manager.
 - 2) Start Installation Manager from the command line with the -record option. For example, run the following command as an administrator:

```

./IBMIM -skipInstall /var/temp/imRegistry -keyring /var/IM/im.keyring
        -password password
        -record /var/temp/keyring_response_file.xml
IBMIM.exe -skipInstall "C:\temp\imRegistry" -keyring C:\IM\im.keyring
        -password password
        -record C:\temp\keyring_response_file.xml

```

- 3) When you are prompted for your credentials for the authenticated remote repository, enter the correct credentials, and save them.
 - 4) Close Installation Manager. Click **File > Exit**. For more information on keyring files, read the topic on silently installing from authenticated repositories with Installation Manager in the Installation Manager information center.
2. Use the response files to install the product silently.
 - a. Use the response file to install the product silently.

From the command line, on each of the systems on which you want to install the product, change to the eclipse subdirectory in the directory in which you installed Installation Manager, and install the product silently. For example, enter the following command as an administrator:

```

./IBMIM --launcher.ini silent-install.ini -input /var/temp/install_response_file.xml
        -log /var/temp/install_log.xml
IBMIMc.exe --launcher.ini silent-install.ini -input C:\temp\install_response_file.xml
        -log C:\temp\install_log.xml

```

Note: The program might write important post-installation instructions to standard output. Read the topic on silently installing packages using Installation Manager in the Installation Manager information center for more information.

- b. Optional: Use the response file to install the keyring silently.

From the command line, for each of the systems on which you want to install the product, change to the eclipse subdirectory in the directory in which you installed Installation Manager, and install the keyring silently. For example, enter the following command as an administrator:

```

./IBMIM --launcher.ini silent-install.ini -input /var/temp/install_response_file.xml
        -log /var/temp/install_log.xml
        -keyring /var/IM/im.keyring -password password
IBMIMc.exe --launcher.ini silent-install.ini -input C:\temp\install_response_file.xml
        -log C:\temp\install_log.xml
        -keyring C:\IM\im.keyring -password password

```

Example

The following is an example of a response file for silently installing the product. This installation includes all the features into the C:\Program Files\IBM\edge\ulb directory and uses a local repository that is located in the C:\repositories\edge85\local-repositories directory.

```

<?xml version="1.0" encoding="UTF-8"?>
<agent-input acceptLicense='true'>
<server>
<repository location='C:\repositories\edge85\local-repositories' />
</server>
<profile id='Load Balancer for IPV6 v8.5'
        installLocation='C:\Program Files\IBM\WebSphere\edge\ulb'>
<data key='eclipseLocation' value='C:\Program Files\IBM\WebSphere\edge\ulb' />
<data key='user.import.profile' value='false' />
<data key='cic.selector.nl' value='cs,ja,es,zh_HK,en,zh_TW,pt_BR,it' />
</profile>
<install modify='false'>

```

```

<offering id='com.ibm.websphere.EDGEIPV4IPV6.v85' version='8.5.0.0'
  profile='IBM WebSphere Edge Components: Load Balancer for IPV4 and IPV6 V8.5'
  features='Dispatcher,Load Balancer,Metric Server,Content Based Routing,Site Selector' />
</install>
<preference name='com.ibm.cic.common.core.preferences.eclipseCache'
  value='C:\Program Files\IBM\IMShared' />
<preference value='30' name='com.ibm.cic.common.core.preferences.connectTimeout' />
<preference value='30' name='com.ibm.cic.common.core.preferences.readTimeout' />
<preference value='0'
  name='com.ibm.cic.common.core.preferences.downloadAutoRetryCount' />
<preference value='false' name='offering.service.repositories.areUsed' />
<preference value='false' name='com.ibm.cic.common.core.preferences.ssl.nonsecureMode' />
<preference value='false'
  name='com.ibm.cic.common.core.preferences.http.disablePreemptiveAuthentication' />
<preference value='true'
  name='com.ibm.cic.common.core.preferences.preserveDownloadedArtifacts' />
<preference value='false' name='PassportAdvantageIsEnabled' />
<preference value='false'
  name='com.ibm.cic.common.core.preferences.searchForUpdates' />
</agent-input>

```

Tip: To disable remote searches for updates in the response file, set the following preferences to false:

offering.service.repositories.areUsed

Used for searching remote repositories for updates to installed offerings

com.ibm.cic.common.core.preferences.searchForUpdates

Used for searching for updates to Installation Manager

For example, change the response file to read:

```

<preference value='false' name='offering.service.repositories.areUsed' />
<preference value='false' name='com.ibm.cic.common.core.preferences.searchForUpdates' />

```

You can find more details on silent preference keys in the information center for IBM Installation Manager. Here are some examples of changes that you could make to manipulate this response file to perform alternative actions:

- Alter the location of the local repository by updating the repository location with the desired path. For example, replace the following line:

```
<repository location='C:\repositories\edge85\local-repositories' />
```

with this:

```
<repository location='/usr/edge85/local-repositories' />
```

- Install from the live remote repository instead of a local repository by replacing the repository location. For example, replace the following line:

```
<repository location='C:\repositories\edge85\local-repositories' />
```

with this:

```
<repository location='http://www.ibm.com/software/repositorymanager/com.ibm.websphere.EDGEIPV4IPV6.v85
```

- Remove optional features by removing the entry for each desired feature in the offering from the comma-separated list. For example, to install without an optional feature, replace the following line:

```

<offering id='com.ibm.websphere.EDGEIPV4IPV6.v85'
  version='8.5.0.0'
  profile='IBM WebSphere Edge Components: Load Balancer for IPV4 and IPV6 V8.5'
  features='Dispatcher,Load Balancer,Metric Server,Content Based Routing,Site Selector' />

```

with this:


```
<offering id='com.ibm.websphere.EDGELBIPV4IPV6.v85'  
  version='8.5.0.0'  
  profile='IBM WebSphere Edge Components: Load Balancer for IPV4 and IPV6 V8.5'  
  features='Dispatcher,Load Balancer,Content Based Routing,Site Selector'>
```

The new line indicates that the Metric Server will not be installed.

What to do next

The installation process does not add the command directories for Load Balancer into the PATH environment variable. To run Load Balancer commands from the system root, add the command directories to the PATH environment variable.

Note: If you installed previous installation of the Dispatcher component, be aware that Websphere Load Balancer for IPv4 versions of the Dispatcher component used commands in the `/usr/bin` directory, which might be in the PATH variable. Websphere Load Balancer for IPv4 and IPv6 has commands in the `install_root/bin` directory, so ensure that the directory entries point to the appropriate directories for the `dsserver` and `dscontrol` commands.

Related tasks:

“Installing the product on distributed operating systems” on page 33
You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.




“Installing fix packs” on page 52

Use IBM Installation Manager to update the Load Balancer product to a later version.

“Uninstalling the product”

You can use IBM Installation Manager to uninstall the Load Balancer product.

Related reference:

-  [WebSphere Application Server detailed system requirements](#)
-  [Recording a response file with Installation Manager in the Installation Manager](#)
-  [Installing packages silently from authenticated repositories with Installation Manager](#)

Uninstalling the product

You can use IBM Installation Manager to uninstall the Load Balancer product.

What to do next

Note: If you are migrating from a previous version of Load Balancer or reinstalling the operating system before installation, you can save any of your previous configuration files or script files for Load Balancer. After installation:

- Place your configuration files in the `installation_root/configurations/dispatcher` directory.
- Place your script files in the `installation_root/servers/bin` directory so that you can run them.

Uninstalling Load Balancer from distributed operating systems using the GUI

Uninstall Load Balancer before upgrading to a newer version or if you think that the current installation is corrupted. Use the IBM Installation Manager GUI to uninstall the product.

Procedure

1. Uninstall the product.
 - a. Ensure that you have stopped all the executors and the servers.
 - b. Start Installation Manager.
 - c. Click **Uninstall**.
 - d. In the Uninstall Packages panel, select **IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6**, and click **Next**.
 - e. Review the summary information.
 - f. Click **Uninstall**.
 - If the uninstallation is successful, the program displays a message that indicates success.
 - If the uninstallation is not successful, click **View log** to troubleshoot the problem.
 - g. Click **Finish**.
 - h. Click **File > Exit** to close Installation Manager.
2. Optional: Uninstall IBM Installation Manager.

Note: Before you can uninstall IBM Installation Manager, you must uninstall all of the packages that were installed by Installation Manager. Read the topic on uninstalling Installation Manager in the Installation Manager information center.

Related tasks:

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

“Installing fix packs” on page 52

Use IBM Installation Manager to update the Load Balancer product to a later version.

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Related reference:

 [Uninstalling Installation Manager](#)

Uninstalling Load Balancer from distributed operating systems using response files

Uninstall Load Balancer before upgrading to a newer version or if you think that the current installation is corrupted. You can use IBM Installation Manager to uninstall Load Balancer using response files.

Before you begin

Optional: Perform or record the installation of Installation Manager and installation of the product to a temporary installation registry on one of your

systems so that you can use this temporary registry to record the uninstallation without using the standard registry where Installation Manager is installed.

Read the following for more information:

- “Setting up IBM Installation Manager” on page 29
- “Installing the product silently” on page 38

About this task

Using Installation Manager, you can work with response files to uninstall the product in a variety of ways. You can use a response file by:

- Writing a new response file by hand.
- Modifying an existing response file or sample.
- Recording a new response file using the GUI as described in this topic

Procedure

1. Stop all the processes that are related to the Load Balancer product.
2. Optional: Record a response file to uninstall the product.

On one of your systems, perform the following actions to record a response file that will uninstall the product:

- a. From a command line, change to the eclipse subdirectory in the directory where you installed Installation Manager.
- b. Start Installation Manager from the command line using the `-record` option.

For example, run the following as an administrator:

```
./IBMIM -skipInstall /var/temp/imRegistry -record /var/temp/uninstall_response_file.xml  
IBMIM.exe -skipInstall "C:\temp\imRegistry" -record C:\temp\uninstall_response_file.xml
```

Tip: If you choose to use the `-skipInstall` parameter with a temporary installation registry that you created, Installation Manager uses the temporary installation registry while recording the response file. When the `-skipInstall` parameter is specified, no product packages are installed or uninstalled. All of the actions that you perform in Installation Manager simply update the installation data that is stored in the specified temporary registry. Do not use the `-skipInstall` on the actual agent data location used by Installation Manager. Use a clean writable location, and reuse that location for future recording sessions.

After the response file is generated, it can be used to uninstall the product, remove the product files, and update the standard installation registry. For more information, read the topic on recording a response file with Installation Manager in the Installation Manager information Center.

- c. Click **Uninstall**.
- d. In the Uninstall Packages window, select **IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6** and click **Next**.
- e. Review the summary information.
- f. Click **Uninstall**.
 - If the uninstallation is successful, the program displays a message that indicates success.
 - If the uninstallation is not successful, click **View log** to troubleshoot the problem.
- g. Click **Finish**.

- h. Click **File > Exit** to close Installation Manager.
3. Use the response file to uninstall the product silently

From a command line on each of the systems from which you want to uninstall the product, change to the eclipse subdirectory in the directory where you installed Installation Manager and use the response file that you created to silently uninstall the product. For example, run the following command as an administrator:

```
./IBMIM --launcher.ini silent-install.ini
-input /var/temp/uninstall_response_file.xml
-log /var/temp/uninstall_log.xml

IBMIMc.exe --launcher.ini silent-install.ini
-input C:\temp\uninstall_response_file.xml
-log C:\temp\uninstall_log.xml
```

Refer to the Installation Manager Information Center for more information.

4. Optional: Uninstall IBM Installation Manager.

Important: Before you can uninstall IBM Installation Manager, you must uninstall all of the packages that were installed by Installation Manager. Read the topic on silently uninstalling Installation Manager in the Installation Manager Information Center for information about using the uninstall script to perform this procedure.

Example

Here is an example of a response file for silently uninstalling the product.

```
<?xml version="1.0" encoding="UTF-8" ?>
<agent-input>
<server>
  <repository location='C:\repositories\edge80\local-repositories' />
</server>
<profile installLocation='C:\Program Files\IBM\WebSphere\edge\ulb'>
  <data key='eclipseLocation' value='C:\Program Files\IBM\WebSphere\edge\ulb' />
  <data key='cic.selector.nl'
    value='en, fr, it, zh, ru, zh_TW, de, ja, pl, es, cs, hu, ko, pt_BR' />
</profile>
<uninstall modify='false'>
  <offering id='com.ibm.websphere.EDGELBIPV4IPV6.v85'
    version='8.5.0.0'
    profile='IBM WebSphere Edge Components: Load Balancer for IPV4 and IPV6 V8.5'
    features='Dispatcher,Load Balancer,Metric Server,Content Based Routing,Site Selector' />
</uninstall>
<preference name='com.ibm.cic.common.core.preferences.eclipseCache'
  value='C:\Program Files\IBM\IMShared' />
<preference value='30' name='com.ibm.cic.common.core.preferences.connectTimeout' />
<preference value='30' name='com.ibm.cic.common.core.preferences.readTimeout' />
<preference value='0' name='com.ibm.cic.common.core.preferences.downloadAutoRetryCount' />
<preference value='true' name='offering.service.repositories.areUsed' />
<preference value='false' name='com.ibm.cic.common.core.preferences.ssl.nonsecureMode' />
<preference value='false'
  name='com.ibm.cic.common.core.preferences.http.disablePreemptiveAuthentication' />
<preference value='true'
  name='com.ibm.cic.common.core.preferences.preserveDownloadedArtifacts' />
<preference value='false' name='PassportAdvantageIsEnabled' />
<preference value='false'
  name='com.ibm.cic.common.core.preferences.searchForUpdates' />
</agent-input>
```

Note: In this example, the offering tag specifies the version attribute. When you include the version attribute, you can use the uninstallation response file only for the version of the offering that you specify. Go to the Installation Manager information center for more information.

Related tasks:

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

“Installing fix packs” on page 52

Use IBM Installation Manager to update the Load Balancer product to a later version.

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Related reference:

 Recording a response file with Installation Manager

 Installation Manager information center

 Uninstalling Installation Manager silently

Uninstalling Load Balancer if it was installed without Installation Manager

If you did not use IBM Installation Manager to install the product, you cannot use Installation Manager to uninstall it. Uninstall the product using the packaging and installation tools on your system.

About this task

Uninstall Load Balancer before you upgrade to a newer version or if you think that the current installation is corrupted.

Procedure

- First, ensure that you have stopped all the executors and the servers.

Uninstall the product by issuing the following command:

```
installp -u package
```

.

- To uninstall the entire product, enter the following command:

```
installp -u ibmulb
```

Use the previous name when it is applicable, **intnd** for example.

- To uninstall specific file sets, list them instead of specifying the package name:

Table 1. AIX Installation Images

Package Name	Installation Image
Base	ibmulb-base-7.0.0-0.noarch.rte
Dispatcher	ibmulb-disp-7.0.0-0.noarch.rte
License	ibmulb-lic-7.0.0-0.noarch.rte
Metric Server	ibmulb-ms-7.0.0-0.noarch.rte
Native	ibmulb-7.0.0-0.ppc64.rte Substitute ppc for ppc64 when it is appropriate for your system.

Table 1. AIX Installation Images (continued)

Package Name	Installation Image
Messages	ibmulb-lang_ <i>language</i> .7.0.0-0.noarch.rte where <i>language</i> can be: <ul style="list-style-type: none"> • cs_CZ • en_US • de_DE • es_ES • fr_FR • hu_HU • it_IT • ja_JP • ko_KR • pl_PL • pt_BR • ru_RU • zh_CN • zh_TW

- First, ensure that you have stopped all the executors and the servers. Use the **swremove** command to uninstall the packages.
 - To uninstall all of the Load Balancer packages:
swremove ibmulb
 - To uninstall an individual package, the Dispatcher component for example, enter:
swremove ibmulb.disp

The package names are:

Table 2. HP-UX Installation Images

Package Name	Installation Image
Base	ibmulb-base-7.0.0-0.noarch.depot
Dispatcher	ibmulb-disp-7.0.0-0.noarch.depot
License	ibmulb-lic-7.0.0-0.noarch.depot
Metric Server	ibmulb-ms-7.0.0-0.noarch.depot
Native	ibmulb-native-7.0.0-0.parisc.depot Substitute ia64 for parisc when it is appropriate for your system.

Table 2. HP-UX Installation Images (continued)

Package Name	Installation Image
Messages	ibmulb-lang_ <i>language</i> .7.0.0-0.noarch.depot where <i>language</i> can be: <ul style="list-style-type: none"> • cs_CZ • en_US • de_DE • es_ES • fr_FR • hu_HU • it_IT • ja_JP • ko_KR • pl_PL • pt_BR • ru_RU • zh_CN • zh_TW

- First, ensure that all the executors and all the servers are stopped.

To uninstall the entire product, enter:

```
rpm -e pkgname
```

The package names are:

Table 3. Linux Installation Images

Package Name	Installation Image
Base	ibmulb-base-7.0.0-0.noarch.rpm
Dispatcher	ibmulb-disp-7.0.0-0.noarch.rpm
License	ibmulb-lic-7.0.0-0.noarch.rpm
Metric Server	ibmulb-ms-7.0.0-0.noarch.rpm
Native	ibmulb-native-7.0.0-0.i386.rpm where i386 can be <ul style="list-style-type: none"> • i386 • ppc • ppc64 • s390 • s390x • x86_64

Table 3. Linux Installation Images (continued)

Package Name	Installation Image
Messages	ibmulb-lang_ <i>language</i> .7.0.0-0.noarch.rpm where <i>language</i> can be: <ul style="list-style-type: none"> • cs_CZ • en_US • de_DE • es_ES • fr_FR • hu_HU • it_IT • ja_JP • ko_KR • pl_PL • pt_BR • ru_RU • zh_CN • zh_TW

- First, ensure that you have stopped all the executors and the servers. Then, to uninstall Load Balancer enter the **pkgrm** command. Enter the following:

`pkgrm package`

The package names are:

Table 4. Solaris Installation Images

Package Name	Installation Image
Base	ibmulb-base-7.0.0-0.noarch.pkg
Dispatcher	ibmulb-disp-7.0.0-0.noarch.pkg
License	ibmulb-lic-7.0.0-0.noarch.pkg
Metric Server	ibmulb-ms-7.0.0-0.noarch.pkg
Native	ibmulb-native-7.0.0-0.sparc.pkg Substitute sparcv9 for sparc when it is appropriate for your system.

Table 4. Solaris Installation Images (continued)

Package Name	Installation Image
Messages	ibmulb-lang_ <i>language</i> .7.0.0-0.noarch.pkg where <i>language</i> can be: <ul style="list-style-type: none"> • cs_CZ • en_US • de_DE • es_ES • fr_FR • hu_HU • it_IT • ja_JP • ko_KR • pl_PL • pt_BR • ru_RU • zh_CN • zh_TW

- To uninstall using the **Add/Remove Programs** option:
 1. Click **Start > Settings > Control Panel**.
 2. Double-click **Add/Remove Programs**.
 3. Select *IBM WebSphere Edge Components* (or the previous name—*IBM Edge Server* for example).
 4. Click **Change/Remove** button.

Related tasks:

“Installing the product on distributed operating systems” on page 33
You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Installing fix packs

Use IBM Installation Manager to update the Load Balancer product to a later version.

Before you begin

Make sure that the preferences for Installation Manager are configured for web-based or local repositories that contain the appropriate updates for the product.

What to do next

If you want to revert to a previous version, refer to the topic on rolling back an update for more information.

Related tasks:

“Installing the product on distributed operating systems” on page 33
You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

Installing fix packs on distributed operating systems using the GUI

You can update Load Balancer to a later version using the IBM Installation Manager GUI.

Before you begin

Make sure that the web-based or local service repository location is listed and checked or that the **Search service repositories during installation and updates** option is selected on the Repositories panel in your Installation Manager preferences. For more information on using service repositories with Installation Manager, read the IBM Installation Manager Information Center.

Procedure

1. Log on to your system.
2. Stop all the executors and servers that are related to Load Balancer.
3. Start Installation Manager.
4. Perform the update process.
 - a. Click **Update**.

Note: If you are prompted to authenticate, use the IBM ID and password that you registered with on the program's web site.

- b. Select the package group to update.

Tip: If you select **Update all**, Installation Manager will search all of the added and predefined repositories for updates to all of the package groups that it has installed. Use this feature only if you have full control over which fixes are contained in the targeted repositories. If you create and point to a set of custom repositories that include only the specific fixes that you want to install, you should be able to use this feature confidently. If you enable searching service repositories or install fixes directly from other live web-based repositories, then you might not want to select this option so that you can select only the fixes that you want to install for each offering on subsequent panels.

- c. Click **Next**.
- d. Select the version to which you want to update under the section for IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6.
- e. Select any fixes that you want to install.

Any recommended fixes are selected by default.

If there are recommended fixes, you can select the option to show only recommended fixes and hide non-recommended fixes.
- f. Click **Next**.
- g. Accept the terms in the license agreements, and click **Next**.
- h. Review the summary information, and click **Update**.

- If the installation is successful, the program displays a message indicating that installation is successful.
 - If the installation is not successful, click **View Log File** to troubleshoot the problem.
- i. Click **Finish**.
5. Click **File > Exit** to close Installation Manager.

What to do next

If you want to revert to a previous version, refer to the topic on rolling back an update for more information.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

Installing fix packs on distributed operating systems using the command line

You can use the IBM Installation Manager command line to update the product with fix packs.

Before you begin

Contact the IBM Software Support Center for information about upgrades. The most current information is available from the IBM Software Support Center and Fix Central.

IBM Installation Manager is used to apply product maintenance.

Note: As an alternative to the procedure that is described in this article, Installation Manager allows you to use the **updateAll** command in a response file or on the command line to search for and update all installed packages. Use this command only if you have full control over which fixes are contained in the targeted repositories. If you create and point to a set of custom repositories that include only the specific fixes that you want to install, you should be able to use this command confidently. If you enable searching service repositories or install fixes directly from other live web-based repositories, then you might not want to select this option so that you can select only the fixes that you want to install using the **-installFixes** option with the **install** command on the command line or the **installFixes** attribute in a response file.

About this task

Use the Installation Manager command line whenever you want to apply a new fix pack to your system.

Procedure

1. For a list of interim fixes and fix packs that are available and specific information about each fix, perform the following actions.
 - a. Go to Fix Central.
 - b. Select **WebSphere** as the product group.

- c. Select **WebSphere Load Balancer for IPv4 and IPv6** as the product.
 - d. Select **8.x** as the installed version.
 - e. Select your operating system as the platform, and click **Continue**.
 - f. Select **Browse for fixes**, and click **Continue**.
 - g. Click **More Information** under each fix to view information about the fix.
 - h. **Recommendation:** Make a note of the name of the fix pack that you would like to install.
2. Update the product with the fix pack using one of the following procedures.

- Access the live service repository that contains the fix pack, and use web-based updating.

Use Installation Manager on your local system to update the product with the fix pack from the live web-based service repositories.

- For the live service repositories, use the same URLs as those used for the generally available product-offering repositories during installation. These URLs are based on the following pattern:

```
http://www.ibm.com/software/repositorymanager/offering_ID
```

where *offering_ID* is the offering ID that you can find in WebSphere Application Server Version 8 product offerings for supported operating systems.

- These locations do not contain web pages that you can access using a web browser. They are remote web-based repository locations that you specify for Installation Manager so that it can maintain the product.

To install a fix pack from a service repository, perform the following actions:

- a. Log on to your system.
- b. If you do not already have an Installation Manager credentials file containing your IBM software user ID and password, create one that will allow you to access the repository.

Note: These are the credentials that you use to access protected IBM software websites.

For information on creating a keyring file for Installation Manager, read the Installation Manager Information Center.

Tip: When creating a keyring file, append `/repository.config` at the end of the URL location if the `imutilsc` command is unable to find the URL that is specified.

- c. Stop all servers and related applications on the installation that is being updated.
- d. Change to the `Installation_Manager_binaries/eclipse/tools` directory, where `Installation_Manager_binaries` is the installation root directory for the Installation Manager.
- e. Install the fix pack.

```
./imcl install offering_ID_offering_version
  -acceptLicense
  -installationDirectory product_installation_location
  -repositories repository_URL
  -keyring keyring_file -password password
imcl.exe install offering_ID_offering_version
  -acceptLicense
  -installationDirectory product_installation_location
  -repositories repository_URL
  -keyring keyring_file -password password
```

where *offering_ID* is the offering ID that is listed in WebSphere Application Server Version 8 product offerings for supported operating systems.

Tips:

- The *offering_ID* is the offering ID that is listed in WebSphere Application Server Version 8 product offerings for supported operating systems.
- The *offering_version*, which optionally can be attached to the offering ID with an underscore, is a specific version of the offering to install (8.5.0.20110503_0200 for example).
 - If *offering_version* is **not** specified, the latest version of the offering and **all** interim fixes for that version are installed.
 - If *offering_version* is specified, the specified version of the offering and **no** interim fixes for that version are installed.

The offering version can be found attached to the end of the offering ID with an underscore when you run the following command against the repository:

```
imcl listAvailablePackages -repositories source_repository
```

- You can also specify none, recommended or all with the `-installFixes` argument to indicate which interim fixes you want installed with the offering.
 - If the offering version is **not** specified, the `-installFixes` option defaults to all.
 - If the offering version is specified, the `-installFixes` option defaults to none.

f. **Optional:** List all installed packages to verify the installation:

```
./imcl listInstalledPackages -long  
imcl.exe listInstalledPackages -long
```

- Download the file that contains the fix pack from Fix Central, and use local updating.

You can download a compressed file that contains the fix pack from Fix Central. Each compressed fix-pack file contains an Installation Manager repository for the fix pack and usually has a .zip extension. After downloading the fix-pack file, use Installation Manager to update IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6 with the fix pack.

a. To download the fix pack, perform the following actions:

- 1) Go to Fix Central.
- 2) Select **WebSphere** as the product group.
- 3) Select **WebSphere Load Balancer for IPv4 and IPv6** as the product.
- 4) Select **8.x** as the installed version.
- 5) Select your operating system as the platform, and click **Continue**.
- 6) Select **Browse for fixes**, and click **Continue**.
- 7) Select the fix pack that you want to download, and click **Continue**.
- 8) Select your download options, and click **Continue**.
- 9) Click **I agree** to agree to the terms and conditions.
- 10) Click **Download now** to download the fix pack.
- 11) Transfer the compressed file in binary format to the system on which it will be installed.

- 12) Extract the compressed repository files to a directory on your system.
- b. To install a fix pack from a downloaded file, perform the following actions:
 - 1) Log on to your system.
 - 2) Stop all servers and related applications on the installation that is being updated.
 - 3) Change to the *Installation_Manager_binaries/eclipse/tools* directory, where *Installation_Manager_binaries* is the installation root directory for the Installation Manager.
 - 4) Install the fix pack.

```
./imcl install offering_ID_offering_version
  -acceptLicense
  -installationDirectory product_installation_location
  -repositories location_of_expanded_files
imcl.exe install offering_ID_offering_version
  -acceptLicense
  -installationDirectory product_installation_location
  -repositories location_of_expanded_files
```

Tips:

- The *offering_ID* is the offering ID that is listed in WebSphere Application Server Version 8 product offerings for supported operating systems.
- The *offering_version*, which optionally can be attached to the offering ID with an underscore, is a specific version of the offering to install (8.5.0.20110503_0200 for example).
 - If *offering_version* is **not** specified, the latest version of the offering and **all** interim fixes for that version are installed.
 - If *offering_version* is specified, the specified version of the offering and **no** interim fixes for that version are installed.

The offering version can be found attached to the end of the offering ID with an underscore when you run the following command against the repository:

```
imcl listAvailablePackages -repositories source_repository
```

- You can also specify none, recommended or all with the *-installFixes* argument to indicate which interim fixes you want installed with the offering.
 - If the offering version is **not** specified, the *-installFixes* option defaults to all.
 - If the offering version is specified, the *-installFixes* option defaults to none.

- 5) **Optional:** List all installed packages to verify the installation:

```
./imcl listInstalledPackages -long
imcl.exe listInstalledPackages -long
```

Related tasks:

“Installing fix packs on distributed operating systems using the GUI” on page 53
You can update Load Balancer to a later version using the IBM Installation Manager GUI.

“Uninstalling fix packs from distributed operating systems using the GUI”
You can roll back this product to an earlier version using the Installation Manager GUI.

“Installing the product on distributed operating systems” on page 33
You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Uninstalling the product” on page 44
You can use IBM Installation Manager to uninstall the Load Balancer product.

Uninstalling fix packs from distributed operating systems using the GUI

You can roll back this product to an earlier version using the Installation Manager GUI.

Before you begin

During the rollback process, Installation Manager must access files from the earlier version of the package. By default, these files are stored on your computer when you install a package. If you change the default setting or delete the saved files, Installation Manager requires access to the repository that was used to install the earlier version.

Important: Installation Manager can install any fix-pack level of the product directly without installing the interim fix packs; in fact, Installation Manager installs the latest level by default. For example, you can skip fix-pack levels and go from Version 8.5.0.1 directly to Version 8.5.0.5. Keep in mind, however, that later you can not roll back to any level that was skipped. If you directly install to Version 8.5.0.5, for example, you cannot roll back to Version 8.5.0.4. If you skip from Version 8.5.0.1 to Version 8.5.0.5, you can only roll back to Version 8.5.0.1.

Procedure

1. Ensure that you have stopped all the processes related to the Load Balancer product.
2. Start Installation Manager.
3. Perform the rollback process.
 - a. Click **Roll Back**.
 - b. Select the package group to roll back.
 - c. Click **Next**.

If you are prompted to authenticate, use the IBM ID and password that you registered with on the program web site.

- d. Select the version to which you want to roll back for the IBM Websphere Edge Components: Load Balancer for IPv4 and IPv6 entry.
- e. Click **Next**.
- f. Review the summary information, and click **Roll Back**.
 - If the rollback is successful, the program displays a message indicating that the rollback is successful.

- If the rollback is not successful, click **View Log File** to troubleshoot the problem.
- g. Click **Finish**.
4. Click **File > Exit** to close Installation Manager.

Related tasks:

“Installing fix packs” on page 52

Use IBM Installation Manager to update the Load Balancer product to a later version.

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Uninstalling the product” on page 44

You can use IBM Installation Manager to uninstall the Load Balancer product.

Directory conventions

References in product information to *install_root* and other directories infer specific default directory locations. This topic describes the conventions in use for IBM WebSphere Edge Components.

These file paths are default locations. You can install the product and other components in any directory where you have write access. You can create profiles in any valid directory where you have write access.

install_root - the root directory in which the product was installed

Load Balancer install paths include the following:

- Administration - /opt/IBM/WebSphere/Edge/ULB/admin
- Dispatcher - /opt/IBM/WebSphere/Edge/ULB/servers
- Content Based Routing- /opt/IBM/WebSphere/Edge/ULB/cbr/servers
- Site Selector- /opt/IBM/WebSphere/Edge/ULB/ss/servers
- Metric Server - /opt/IBM/WebSphere/Edge/ULB/ms
- Documentation - /opt/IBM/WebSphere/Edge/ULB/docs/

install_root - the root directory in which the product was installed

Load Balancer install paths include the following:

- Administration - C:\Program Files\IBM\WebSphere\Edge\ULB\admin
- Dispatcher - C:\Program Files\IBM\WebSphere\Edge\ULB\servers
- Content Based Routing- C:\Program Files\IBM\WebSphere\Edge\ULB\cbr\servers
- Site Selector- C:\Program Files\IBM\WebSphere\Edge\ULB\ss\servers
- Metric Server - C:\Program Files\IBM\WebSphere\Edge\ULB\ms
- Documentation - C:\Program Files\IBM\WebSphere\Edge\ULB\docs

Chapter 5. Configuring Load Balancer

This section focuses on configuring production environments and realistic test environments.

Configuring Dispatcher

Configure Dispatcher for production environments and realistic test environments.

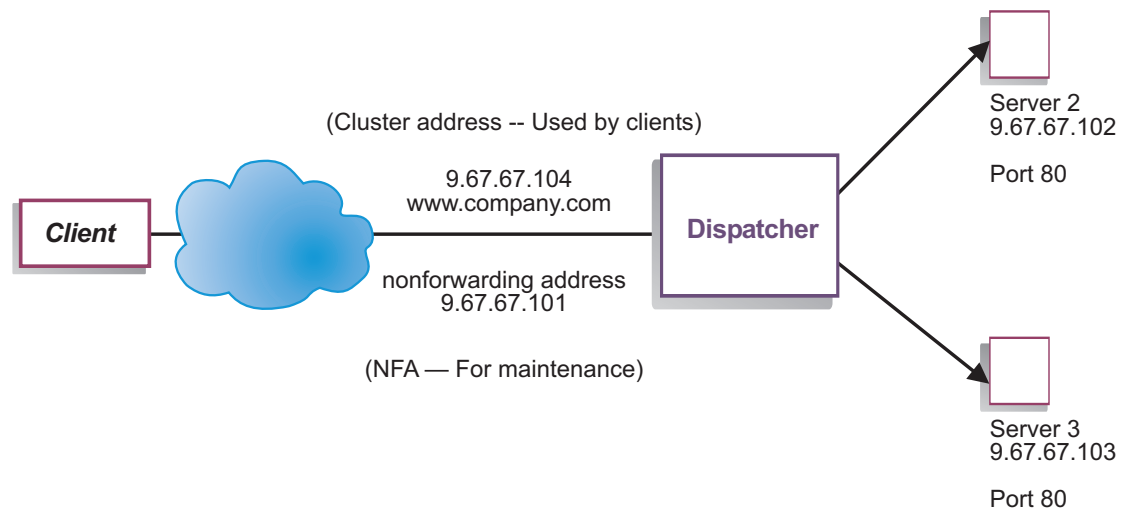
This section provides information on a quick start configuration, planning considerations, and describes the methods of configuring the Load Balancer Dispatcher component.

Quick start configuration

This quick start example shows how to configure three locally attached workstations to load-balance web traffic between two web servers.

Before you begin

For the quick start example, you need three workstations and four IP addresses. One workstation for the Dispatcher; the other two workstations are the web servers. Each web server requires one IP address. The Dispatcher workstation requires two addresses: the non-forwarding address (NFA), and the cluster address (the address that is load balanced) that you provide to clients to access your website.



For more information about the different ways Load Balancer can be set up, read “Types of cluster, port, and server configurations for Dispatcher” on page 26 to help you design your topology.

About this task

Use this configuration method for a quick way to establish a connection between servers and the dispatcher machine. This method does not include configuring advisors or tuning performance. For a full configuration, read “Configuring the

Dispatcher machine” on page 69 and “Configuring the server machines” on page 73.

Procedure

1. Prepare your servers.
 - a. For this locally attached configuration example, set up your workstations on the same LAN segment. Ensure that network traffic between the three machines does not pass through any routers or bridges.
 - b. Configure the network adapters of the three workstations. For this example, it is assumed you have the following network configuration, and each of the workstations contains only one standard Ethernet network interface card:

Table 5. Sample network configuration

Workstation	Name	IP address
1	server1.Intersplashx.com	9.67.67.101
2	server2.Intersplashx.com	9.67.67.102
3	server3.Intersplashx.com	9.47.47.103
Netmask = 255.255.255.0		

- c. 3. Ensure that all the servers can communicate with each other.
 - 1) Ensure that server1.Intersplashx.com can ping both server2.Intersplashx.com and server3.Intersplashx.com.
 - 2) Ensure that server2.Intersplashx.com and server3.Intersplashx.com can ping server1.Intersplashx.com.
 - d. Ensure that content is identical on the 2 web servers (Server 2 and Server 3). You can complete this step by replicating data on both workstations, using a shared file system such as NFS, AFS[®], or DFS[™], or by any other means appropriate for your site.
 - e. Ensure that web servers on server2.Intersplashx.com and server3.Intersplashx.com are operational. Use a web browser to request pages directly from http://server2.Intersplashx.com and http://server3.Intersplashx.com.
 - f. Obtain another valid IP address for this LAN segment. This IP address is the address that you provide to clients who want to access your site. For example:


```
Name= www.Intersplashx.com
IP=9.47.47.104
```
 - g. Configure the two web server workstations to accept traffic for www.Intersplashx.com. Add an alias for www.Intersplashx.com to the loopback interface on server2.Intersplashx.com and server3.Intersplashx.com:
 - ifconfig lo0 alias www.Intersplashx.com netmask 255.255.255.255
 - ifconfig lo0:1 plumb www.Intersplashx.com netmask 255.255.255.0 up
 - For other operating systems, see your operating system's instructions in “Aliasing the network interface card or loopback device” on page 76.
 - h. Delete any extra route that was created as a result of aliasing the loopback interface. See Step 2 in Configuring the server machines.
 - i.
2. Configure Load Balancer with the command line, the GUI, or the configuration wizard.

- Configuring with the command line:
 - a. Start the **dsserver** on Dispatcher:
 - Run the following command as root user:


```
dsserver
```
 - **dsserver** runs as a service that starts automatically
 - b. Start the executor function of Dispatcher by entering the following command:


```
dscontrol executor start
```
 - c. Add the cluster address to the Dispatcher configuration:


```
dscontrol cluster add www.Intersplashx.com
```
 - d. Add the HTTP protocol port to the Dispatcher configuration:


```
dscontrol port add www.Intersplashx.com@80
```
 - e. Add each of the web servers to the Dispatcher configuration:


```
dscontrol server add www.Intersplashx.com@80@server2.Intersplashx.com
dscontrol server add www.Intersplashx.com@80@server3.Intersplashx.com
```
 - f. Start the manager function of Dispatcher:


```
dscontrol manager start
```

Dispatcher now load balances based on server performance.

- g. Start the advisor function of Dispatcher:


```
dscontrol advisor start http 80
```

Dispatcher now makes sure that client requests are not sent to a failed web server.

- Configuring with the configuration wizard:
 - a. Start the **dsserver** on Dispatcher:
 - Run the following command as root user: **dsserver**
 - **dsserver** runs as a service that starts automatically
 - b. Start the wizard function of Dispatcher:


```
dswizard
```

The wizard guides you step-by-step through the process of creating a basic configuration for the Dispatcher component. It asks questions about your network and guides you through the setup of a cluster for Dispatcher to load balance the traffic for a group of servers. The configuration wizard contains the following panels:

 - Introduction to the wizard
 - What is going to happen
 - Preparing for the setup
 - Defining a cluster
 - Adding a port
 - Adding a server
 - Starting an advisor
 - Server setup
- Configuring with the GUI:
 - At a command prompt, enter the following command:


```
lbadmin
```
 - Click **Start > Programs > IBM WebSphere > Edge Components > IBM Load Balancer > Load Balancer**.

3. Test your configuration.
 - a. From a web browser, go to `www.Intersplashx.com`. If a page displays, the configuration is working.
 - b. Reload the page in the web browser.
 - c. Look at the results of the following command:


```
dscontrol server report www.Intersplashx.com@80@
```

The total connections column of the 2 servers should add up to 2.

Related tasks:

“Configuring the Dispatcher machine” on page 69

Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

“Configuring the server machines” on page 73

Related reference:

“Types of cluster, port, and server configurations for Dispatcher” on page 26

There are many ways that you can configure Dispatcher to support your site.

Plan for Dispatcher

You can use the Dispatcher component by itself to balance the load on servers within a local area network or wide area network, using a number of weights and measurements that are dynamically set by Dispatcher.

Dispatcher consists of the following functions:

- The **dsserver** command handles requests from the command line to the executor, manager, and advisors.
- The **executor** command supports port-based load balancing of TCP and UDP connections. It is able to forward connections to servers based on the type of request received (for example, HTTP, FTP, SSL, and so forth). The executor always runs when the Dispatcher component is being used for load balancing.
- The manager sets weights used by the executor based on:
 - Internal counters in the executor
 - Feedback from a system-monitoring program, such as Metric Server or WLM.

Using the manager is optional. However, if the manager is not used, load balancing is performed using weighted round-robin scheduling based on the current server weights, and advisors are not available.

- The advisors query the servers and analyze results by protocol before calling the manager to set weights as appropriate. Currently there are advisors available for the following protocols: HTTP, FTP, SSL, SMTP, NNTP, IMAP, POP3, SIP, and Telnet.

Dispatcher also offers advisors that do not exchange protocol-specific information, such as the DB2 advisor that reports on the health of DB2 servers and the ping advisor that reports whether the server responds to a ping.

Using the advisors is optional but recommended.

- To configure and manage the executor, advisors, and manager, use the command line (**dscontrol**) or the graphical user interface (**lbadmin**).
- A sample configuration file is provided to use for configuration and administration of the Dispatcher machine. After you have installed the product, find this file in the following directory:
 - `/opt/ibm/edge/lb/servers/samples`
 - `<install_root>ibm\edge\lb\servers\samples`

- The SNMP subagent enables an SNMP-based management application to monitor the status of the Dispatcher.

The three key functions of Dispatcher (executor, manager, and advisors) interact to balance and dispatch the incoming requests between servers. Along with load balancing requests, the executor monitors the number of new connections, active connections, and connections in a finished state. The executor also does garbage collection of completed or reset connections and supplies this information to the manager.

The manager collects information from the executor, the advisors, and a system-monitoring program, such as Metric Server. Based on the information the manager receives, it adjusts how the server machines are weighted on each port and gives the executor the new weighting for use in its balancing of new connections.

The advisors monitor each server on the assigned port to determine the server's response time and availability and then give this information to the manager. The advisors also monitor whether a server is up or down. Without the manager and the advisors, the executor does round-robin scheduling based on the current server weights.

Methods of configuration for Dispatcher

There are four methods you can use to configure Load Balancer: the command line, scripts, the graphical user interface (GUI), and the configuration wizard. The Information Center assumes use of the command line.

There are four basic methods of configuring the Dispatcher:

- "Command line"
- "Scripts" on page 66
- Graphical User Interface (GUI)
- "The configuration wizard" on page 67.

Command line

This is the most direct means of configuring the Dispatcher. The command parameter values must be entered in English characters. The only exceptions are host names (used in cluster, server, and high availability commands) and file names (used in file commands).

To start the Dispatcher from the command line:

1. Start dsserver:

- From the command prompt, issue the following:
dsserver

To stop the service, type:

dsserver stop

- Click **Start > Control Panel > Administrative Tools > Services**. Right-click **IBM Dispatcher (ULB)** and select **Start**. To stop the service, follow the same steps and select **Stop**.
2. Issue the Dispatcher control commands you want in order to set up your configuration. The command is `dscontrol`. For more information about commands, see "Commands" on page 226.

You can use a minimized version of the `dscontrol` command parameters by typing the unique letters of the parameters. For example, to get help on the file save command, you can type

```
dscontrol he f
```

instead of

```
dscontrol help file
```

To start up the command line interface issue `dscontrol` to receive an `dscontrol` command prompt. To end the command line interface issue the commands `exit` or `quit`.

Scripts

You can enter commands for configuring Dispatcher into a configuration script file and run them together. See Sample Load Balancer configuration files. To quickly run the content of a script file (for example, `myscript`), use either of the following commands:

- To update the current configuration, run the following executable commands from your script file:

```
dscontrol file appendload myscript
```

- To completely replace the current configuration, run the following executable commands from your script file:

```
dscontrol file newload myscript
```

To save the current configuration into a script file (for example, `savescript`), run the following command:

```
dscontrol file save savescript
```

This command will save the configuration script file in the `install_root/servers/configurations/dispatcher` directory.

GUI

To start the GUI, follow these steps:

1. Ensure **dsserver** is running:
 - Run the following as root user:

```
dsserver
```
 - **dsserver** runs as a service that starts automatically.
2. Start the Load Balancer GUI:
 - Run the following:

```
lbadmin
```
 - Click **Start > Programs > IBM WebSphere > Edge Components > IBM Load Balancer > Load Balancer**

To configure the Dispatcher component from the GUI, you must first select **Dispatcher** in the tree structure. You can start the executor and manager after you connect to a **Host**. You can also create clusters containing ports and servers, and start advisors for the manager.

The GUI can be used to do anything that you would do with the `dscontrol` command. For example, to define a cluster using the command line, you would enter `dscontrol cluster add cluster` command. To define a cluster from the GUI,

right-click **Executor**, then in the popup menu, left-click **Add Cluster**. Enter the cluster address in the popup window, then click **OK**.

Pre-existing Dispatcher configuration files can be loaded using the **Load New Configuration** (for completely replacing the current configuration) and **Append to Current Configuration** (for updating the current configuration) options presented in the **Host** popup menu. You should save your Dispatcher configuration to a file periodically using the **Save Configuration File As** option also presented in the **Host** popup menu. The **File** menu located at the top of the GUI will allow you to save your current host connections to a file or restore connections in existing files across all Load Balancer components.

In order to run a command from the GUI:

1. Highlight the Host node from the GUI tree and select **Send command...** from the Host pop-up menu.
2. In the command entry field, type the command that you want to run, for example:
`executor report`

The results and history of the commands run in the current session and appear in the window provided.

You can access Help by clicking the question mark icon in the upper right corner of the Load Balancer window.

- **Help: Field level** -- describes each field, default values
- **Help: How do I** -- lists tasks that can be done from that screen
- **InfoCenter** -- provides centralized access to product information

The configuration wizard

The wizard guides you step by step through the process of creating a basic configuration for the Dispatcher component. You will be asked questions about your network. You will be guided through the setup of a cluster for Dispatcher to load balance traffic between a group of servers.

If you are using the configuration wizard, follow these steps:

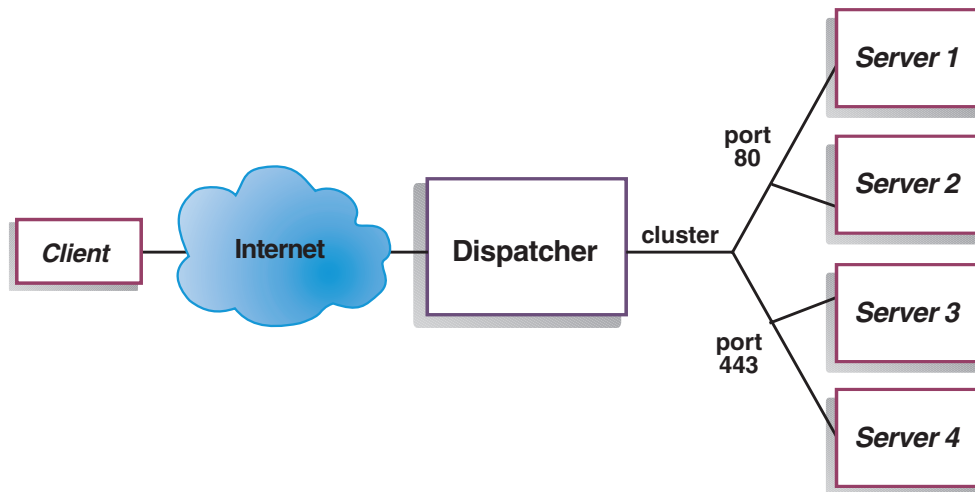
1. Start the **dserver** on Dispatcher:
 - Run the following as root user:
`dserver`
 - **dserver** runs as a service that starts automatically.
2. Start the wizard function of Dispatcher by issuing the following command:
`dswizard`

Types of cluster, port, and server configurations for Dispatcher

There are many ways that you can configure Dispatcher to support your site.

1 cluster with 2 ports

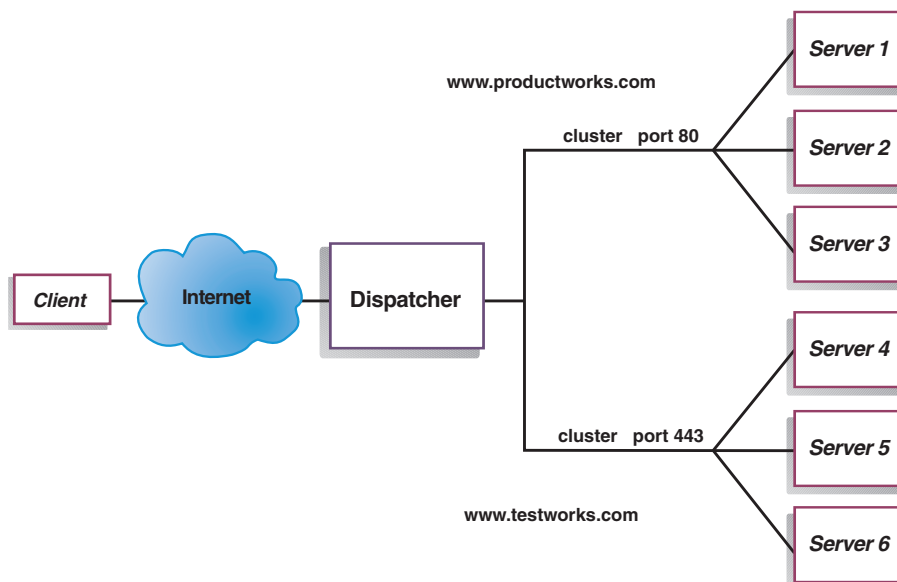
If you have only one host name for your site to which all of your customers will connect, you can define a single cluster of servers. For each of these servers, configure a port through which Load Balancer communicates.



In this example for the Dispatcher component, one cluster is defined at `www.productworks.com`. This cluster has two ports: port 80 for HTTP and port 443 for SSL. A client making a request to `http://www.productworks.com` (port 80) goes to a different server than a client requesting `https://www.productworks.com` (port 443).

2 clusters, each with 1 port

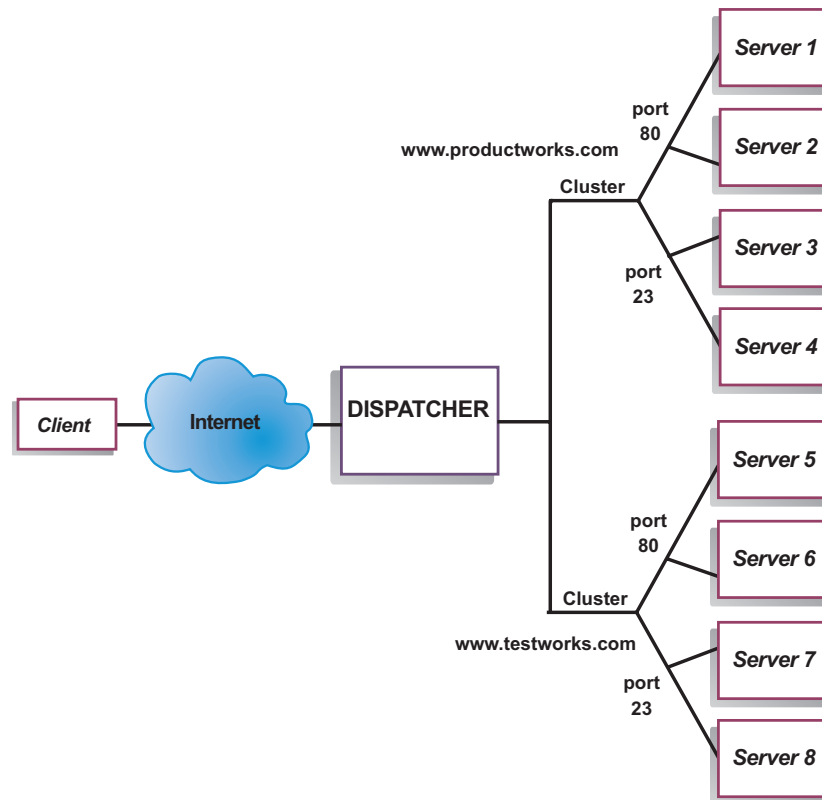
Another way of configuring Load Balancer might be appropriate if you have a very large site with many servers dedicated to each protocol supported. In this case, you might want to define a cluster for each protocol with a single port but with many servers.



In this example for the Dispatcher component, two clusters are defined: `www.productworks.com` for port 80 (HTTP) and `www.testworks.com` for port 443 (SSL). A third way of configuring Load Balancer might be necessary if your site does content hosting for several companies or departments, each one coming into your site with a different URL. In this case, you might want to define a cluster for each company or department and then define any ports to which you want to

receive connections at that URL, as shown in the configuration for 2 clusters, each with two ports.

2 clusters, each with 2 ports



In this example for the Dispatcher component, two clusters are defined with port 80 for HTTP and port 23 for Telnet for each of the sites at `www.productworks.com` and `www.testworks.com`.

Related tasks:

“Quick start configuration” on page 61

This quick start example shows how to configure three locally attached workstations to load-balance web traffic between two web servers.

Configuring the Dispatcher machine

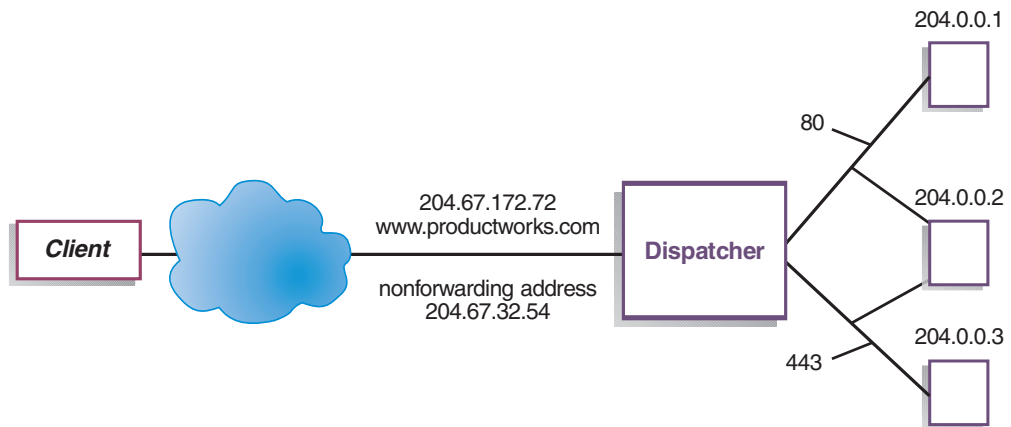
Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

Before you begin

Before setting up the Dispatcher machine, you must be the root user (for AIX, HP-UX, Linux, or Solaris systems) or the Administrator on Windows systems.

For the Dispatcher machine you will need at least two valid IP addresses.

The figure below shows an example of Load Balancer set up with a single cluster, two ports, and three servers.



About this task

You must configure the machine on which Load Balancer is installed before you can load balance traffic in your network environment.

Review the new features section for what's been added and improved in this release.

Note: Maximize your load-balancing by being aware of explicit addressing in links, and avoid it where possible.

If your pages specify links that point to individual servers for your site, you are in effect forcing a client to go to a specific machine, thus bypassing any load balancing function that might otherwise be in effect.

In general, the load balancing function works independently of the content of the sites on which you use the product, but link addressing is one area in which the site content can be important. Decisions that you make regarding the site's content can have a significant impact upon the Load Balancer's efficiency. Use the address of Load Balancer in any links contained in your pages. Note that the kind of addressing used might not always be apparent if your site uses automated programming that dynamically creates HTML.

Note: Be aware of the following restrictions:

- Load Balancer does not support passive FTP.
- **Solaris back-end server:** There is no support for load balancing IPv6 traffic to backend Solaris 5.8 servers. On Solaris 5.8, there is an incompatibility with a MAC-forwarded IPv6 packet and the Solaris IPv6 stack. When the cluster is configured on a Solaris 5.8 backend server using the `ifconfig lo0` (loopback) command, the packet arrives at the Solaris 5.8 node, but is not accepted. However, you can use Load Balancer for IPv4 and IPv6 installations to load balance IPv4 traffic to backend Solaris 5.8 servers.
- **z/OS back-end servers:** There is no support for load-balancing IPv6 traffic to backend z/OS servers. However, you can load balance IPv4 traffic to back-end z/OS servers using Load Balancer for IPv4 and IPv6 installations.

Procedure

1. Start the server function. To start the server function, type the following at a command prompt:

dsserver

Note: A default configuration file (default.cfg) gets automatically loaded when starting **dsserver**. If the user decides to save the Dispatcher configuration in default.cfg, then everything saved in this file is automatically loaded next time **dsserver** gets started.

2. Start the executor function.

- a. Optional: If you are using IPv6 addresses, enable the processing of IPv6 packets.

Prior to starting the executor (dscontrol executor start), the following must be issued from the command line as root:

- autoconf6

To enable uninterrupted processing of IPv6 packets, even after a system reboot, edit the etc/rc.tcpip file and uncomment the following line, and add the -A flag:

```
start usr/bin/autoconf6 " " -A
```

- modprobe ipv6
- netsh interface ipv6 install

These commands enable processing of IPv6 packets in the respective operating systems. Issue this command only once. Thereafter, you can start and stop the executor as often as you need. If you do not issue the command to enable processing of IPv6 packets on these systems, the executor will not start.

Using the **ifconfig** command, IPv6 addresses must be plumbed and an interface configured in order for Dispatcher to inspect IPv6 packets. If you do not issue these commands, the executor will start, but no IPv6 packets can be viewed. Prior to starting the executor (dscontrol executor start), issue the following from the command line as root:

- ifconfig *device* inet6 up
- Change the device to your device name, and change the IPv6 IP address and prefix to your address and prefix values:

```
ifconfig device inet6 plumb  
ifconfig device inet6 address/prefix up
```

- b. To start the executor function, enter the dscontrol executor start command. You may also change various executor settings at this time.

3. Optional: Define the non-forwarding address if it is different from the host name. The non-forwarding address is used to connect to the machine for administrative purposes, such as using Telnet or SMTP to this machine.

By default, this address is the hostname. To define the non-forwarding address, enter the following command, or edit the configuration file:

```
dscontrol executor set nfa IP_address
```

where *IP_address* is either the symbolic name or the IP address.

4. Define a cluster and set cluster options. Dispatcher will balance the requests sent to the cluster address to the servers configured on the ports for that cluster. The cluster is either the symbolic name, the dotted decimal address, or the special address 0.0.0.0 that defines a wildcard cluster. Wildcard clusters can be used to match multiple IP addresses for incoming packets to be load balanced.

- a. To define a cluster, issue the dscontrol cluster add command:

```
dscontrol cluster add cluster
```

- b. To set cluster options, use the `dscontrol cluster set` command. Issue the following command:


```
dscontrol cluster set options
```
- c. If you use a qeth/OSA implementation in layer 3 mode on Linux on z/OS, perform the following additional configuration steps to setup Load Balancer.

Note: These steps are not required if you run qeth/OSA in layer 2 mode.

- 1) 1. Configure the cluster address using `ip` or `ifconfig` command:

```
ip -version addr add cluster_address/prefix_length dev device
```

For example:

```
ip -4 addr add 12.42.38.125/24 dev eth0
ip -6 addr add 3ffe:34::24:45/64 dev eth0
```

- 2) 2) Add an iptables rule to drop incoming packets destined to the cluster address:

For IPv4 addresses:

```
iptables -t filter -A INPUT -d cluster_address -j DROP
```

For IPv6 addresses:

```
ip6tables -t filter -A INPUT -d cluster_address -j DROP
```

For example:

```
iptables -t filter -A INPUT -d 12.42.38.125 -j DROP
ip6tables -t filter -A INPUT -d 3ffe:34::24:45 -j DROP
```

- 3) 3) To undo the above configuration, use the following commands:

```
ip -version addr del cluster_address/prefix_length dev device
iptables -t filter -D INPUT -d cluster_address -j DROP
ip6tables -t filter -D INPUT -d cluster_address -j DROP
```

5. Define ports and set port options with the `dscontrol port add` command. You must define and configure all servers for a port.

- a. Define a port. Enter the following command:

```
dscontrol port add cluster@port
```

- *cluster* is either the symbolic name or the IP address
- *port* is the number of the port you are using for that protocol

- b. Change various port settings. Read “`dscontrol port`” on page 244 for more information on this command and the available options.

Note: You can select a new option for the selection algorithm that Load Balancer uses to route traffic:

- **conn+affinity:** Specifies that server selection is based on an existing connection. For new connections, the server selection is based on affinity.

You can also edit the sample configuration file or use the GUI.

6. Define the load-balanced server machines. To define a load-balanced server machine, enter the following command:

```
dscontrol server add cluster@port@server
```

You can also edit the sample configuration file or use the GUI. *Cluster* is either the symbolic name or the IP address, and *port* is the number of the port you are using for that protocol. You must define more than one server to a port on a cluster in order to perform load balancing.

- a. **Configure IPv6 link-local address:** With IPv6 addressing, each machine in the Load Balancer configuration must have an IPv6 link-local address. The

link-local address is the address used for neighbor discovery traffic for IPv6, and without this address on the Load Balancer machine and on the back-end servers neighbor discovery does not occur, and the machines are not known to each other. Load Balancer for IPv6 cannot forward traffic without a link-local IPv6 address configured on an interface of each machine in the Load Balancer configuration.

- b. Optional: **Bind-specific servers:** If the Dispatcher component is load balancing to bind-specific servers, then the servers must be configured to bind to the cluster address. Because the Dispatcher forwards packets without changing the destination IP address, when the packets reach the server, the packets will still contain the cluster address as the destination. If a server has been configured to bind to an IP address other than the cluster address, then the server will be unable to accept requests destined for the cluster.

To determine if the server is bind specific, issue the `netstat -an` command and look for the `server@port`. If the server is not bind specific, the result from this command will be `0.0.0.0@80`. If the server is bind specific, you will see an address such as `192.168.15.103@80`.

7. Optional: Start the manager function. The manager function improves load balancing. To start the manager, enter the `dscontrol manager start` command, edit the sample configuration file, or use the GUI. For example:

```
dscontrol manager start
```

8. Optional: Start the advisor function. The advisors give the manager more information about the ability of the load-balanced server machines to respond to requests. An advisor is specific to a protocol. For example, to start the HTTP advisor, issue the following command:

```
dscontrol advisor start http port
```

For a list of advisors along with their default ports, see “List of advisors” on page 114.

- a. Set cluster proportions, as required. If you start advisors, you may modify the proportion of importance given to advisor information being included in the load balancing decisions. To set the cluster proportions, issue the `dscontrol cluster set cluster proportions` command. For more information, see “Tuning the proportion of importance given to status information” on page 178

9. Configure the server machines.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Configuring the server machines”

Configuring the server machines

About this task

You need to configure each of the back-end server machines in your topology before Load Balancer will be able to properly work in your environment.

Dispatcher will only load balance across servers that allow the loopback adapter to be configured with an additional IP address, for which the back-end server will never respond to ARP (address resolution protocol) requests.

Procedure

1. Configure the loopback method.

- Alias the network interface card or loopback device.
- If you are connecting Load Balancer to a server that runs IBM i, create a virtual IP address with the address you want to use for the cluster IP by running the Add TCP/IP Interface command. For example:

```
ADDTCPIFC INTNETADR('cluster_IP_address') LIND(*VIRTUALIP) SUBNETMASK(*HOST)
```

When the server starts the virtual TCP/IP interface, the server will accept packets for that address.

Note: Do not enable proxy ARP for the Virtual IP Address; don't specify the PREIFC parameter on the command or enable proxy through the GUI configuration. This would prevent multiple systems from simultaneously using the same cluster IP address.

- For most types of Linux, you can follow the instructions in the topic on aliasing the loopback device.

If you are using certain types of Linux operating systems, you may need to use an alternative method for aliasing the network card and loopback devices. Some versions of Linux systems issue ARP responses for any IP address configured on the machine on any interface present on the machine. If your system issues an ARP response, then follow the steps in the “Configuring loopbacks with alternative methods” on page 79 topic.

2. Check for an extra route. On some operating systems, a default route may have been created and needs to be removed. If problems are encountered with routing after aliasing, remove the alias and add it back using a different netmask.

Note: Any extra routes should be ignored on Windows 2003.

- Check for an extra route on Windows operating systems with the following command:

```
route print
```

For example:

- a. After route print is entered, a table similar to the following example will be displayed. This example shows finding and removing an extra route to cluster 9.67.133.158 with a default netmask of 255.0.0.0.

Active Routes:

Network Address	Netmask	Gateway Address	Interface	Metric
0.0.0.0	0.0.0.0	9.67.128.1	9.67.133.67	1
9.0.0.0	255.0.0.0	9.67.133.158	9.67.133.158	1
9.67.128.0	255.255.248.0	9.67.133.67	9.67.133.67	1
9.67.133.67	255.255.255.255	127.0.0.1	127.0.0.1	1
9.67.133.158	255.255.255.255	127.0.0.1	127.0.0.1	1
9.255.255.255	255.255.255.255	9.67.133.67	9.67.133.67	1
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1
224.0.0.0	224.0.0.0	9.67.133.158	9.67.133.158	1
224.0.0.0	224.0.0.0	9.67.133.67	9.67.133.67	1
255.255.255.255	255.255.255.255	9.67.133.67	9.67.133.67	1

- b. Find your cluster address under the “Gateway Address” column. If you have an extra route, the cluster address will appear twice. In the example given, the cluster address (9.67.133.158) appears in row 2 and row 8.
- c. Find the network address in each row in which the cluster address appears. You need one of these routes and will need to delete the other route, which is extraneous. The extra route to be deleted is the one whose

network address begins with the first digit of the cluster address, followed by three zeroes. In the example shown, the extra route is the one in row two, which has a network address of 9.0.0.0:

Network Address	Netmask	Gateway Address	Interface	Metric
9.0.0.0	255.0.0.0	9.67.133.158	9.67.133.158	1

- Check for an extra route on all Linux and UNIX systems with the following command:

```
netstat -nr
```

3. Delete any extra route. You must delete the extra route. Use the command for your operating system shown below:

- `route delete -net network_address cluster_address`
- `route delete cluster_address cluster_address`
- `route delete network_address cluster_address`

Note: You must delete the extra route every time you reboot the server.

To delete the extra route as shown in the \triangle Active Routes \triangle example above, enter:

```
route delete 9.0.0.0 9.67.133.158
```

On Windows 2003, it is not possible to delete routes. Any extra routes should be ignored on Windows 2003. If problems are encountered with routing after aliasing, remove the alias and add it back using a different netmask.

4. Configure the workload partition (WPAR) as a backend server. Run all of the following commands on the machine that hosts the WPAR.

- a. Enable WPAR specific routing. Enter the following command:

```
chwpwr -i -I rtdest=default rtgateway=<Default_Gateway> <WPAR_name>
```

- b. Configure the IP address for the cluster on the loopback interface of the WPAR.

```
chwpwr -N interface=lo0 address=<cluster_IP_address> netmask=255.255.255.0 <WPAR_name>
```

5. Verify that the server is properly configured. To verify if a backend server is properly configured, perform the following steps from a different machine on the same subnet when the Load Balancer is not running and cluster is unconfigured:

- a. Issue the arp command. For example:

```
arp -d cluster
```

- b. Issue the ping command. For example:

```
ping cluster
```

There should be no response. If there is a response to the ping, ensure that you did not ifconfig the cluster address to the interface. Ensure that no machine has a published arp entry to the cluster address.

- c. Ping the back-end server, then immediately issue the arp -a command For example:

```
arp -a
```

In the output from the command, you should see the MAC address of your server. Issue the command:

```
arp -s cluster server_MAC_address
```

- d. Ping the cluster. You should get a response. Issue a http, telnet, or other request that is addressed to the cluster that you expect your back-end server to handle. Ensure that it works properly.

- e. Issue the arp -d command. For example:

```
arp -d cluster
```


- f. Ping the cluster. There should be no response. If there is a response, issue an arp cluster instruction to get the MAC address of the machine that is not configured correctly, and repeat steps 1 through 6.

Related tasks:

Chapter 5, “Configuring Load Balancer,” on page 61

This section focuses on configuring production environments and realistic test environments.

“Configuring the Dispatcher machine” on page 69

Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

Chapter 4, “Installing Load Balancer,” on page 29

This section describes how to establish load balancing capability in new and existing environments. The information includes planning, preparing for, completing, and maintaining product installations.

Aliasing the network interface card or loopback device

About this task

To alias the loopback device on servers that are being load-balanced, you must use the operating system's adapter configure commands. For the load-balanced server machines to work, you must set, or preferably alias, the loopback device, which is often called lo0, to the cluster address. By setting or aliasing the loopback device to the cluster address, the load balanced server machines will accept a packet that was addressed to the cluster address.

If you have an operating system that supports network interface aliasing (such as AIX, HP-UX, Linux, Solaris, or Windows systems), you should alias the loopback device to the cluster address. The benefit of using an operating system that supports aliases is that you have the ability to configure the load-balanced server machines to serve multiple cluster addresses.

Note: See “Configuring loopbacks with alternative methods” on page 79. If you have a server with an operating system that does not support aliases you must set the loopback device to the cluster address.

Procedure

- Use the following commands to alias the network interface and the loopback device (*interface_name*) for AIX systems :

- **For IPv4 addresses:**

```
ifconfig lo0 alias cluster_address netmask 255.255.255.255
```

- **For IPv6 addresses:**

```
ifconfig interface_name inet6 cluster_address/prefix_length alias
```

For example, to alias the loopback device on servers that are being load-balanced:

```
ifconfig lo0 inet6 2002:4a::541:56/128 alias
```

- Use these commands to alias the network interface and the loopback device (*interface_name*) for HP-UX systems:

- **For IPv4 addresses:**

```
ifconfig lo0:1 cluster_address up
```

- **For IPv6 addresses:**

```
ifconfig interface_name:alias inet6 cluster_address up prefix prefix_length
```


For example, to alias the loopback device on servers that are being load-balanced:

```
ifconfig lo0:1 inet6 3ffe:34::24:45 up prefix 128
```

Note: When using bind-specific server applications that bind to a list of IP addresses that do not contain the server's IP, use the `arp publish` command instead of `ifconfig` to dynamically set an IP address on the Load Balancer machine. For example:

```
arp -s <cluster> <Load Balancer's MAC address> pub
```

- Use these commands to alias the network interface and the loopback device (*interface_name*) for Linux systems:

- **For IPv6 or IPv4 addresses:**

```
ip -version addr add cluster_address/prefix_length dev lo
```

For example, to alias the loopback device on servers that are being load-balanced:

```
ip -6 addr add 3ffe:34::24:45/128 dev lo
ip -4 addr add 12.42.38.125/32 dev lo
```

Note: You can also use the `ifconfig` command. See below to alias the loopback device using the `ifconfig` command. Once you issue one of the configuration commands on your machine, it is important to consistently use the same configuration command (`ip` or `ifconfig`), or unexpected results can occur.

- **Using the `ifconfig` command:**

```
ifconfig lo:1 cluster_address netmask 255.255.255.255 up
```

- The following commands can be used to alias the network interface and the loopback device (*interface_name*) for Solaris systems.

- **For IPv4 addresses:**

```
ifconfig lo0:1 plumb cluster_address netmask netmask up
```

- **For IPv6 addresses:**

```
ifconfig interface_name inet6 addif cluster_address/prefix_length up
```

For example, to alias the loopback device on servers that are being load-balanced:

```
ifconfig lo0 inet6 addif 3ffe:34::24:45/128 up
```

Note: When using bind-specific server applications that bind to a list of IP addresses that do not contain the server's IP, use the `arp publish` command instead of `ifconfig` to dynamically set an IP address on the Load Balancer machine. For example:

```
arp -s <cluster> <Load Balancer's MAC address> pub
```

- Use these commands to alias the network interface and the loopback device for Windows operating systems.

1. Use the `ipconfig /all` command to determine the interface name for the loopback device. This command locates the connection with a description of the Microsoft Loopback Adapter. The following example is the output from the `ipconfig /all` command, where the Microsoft Loopback Adapter is Ethernet adapter Local Area Connection 2, so the connection is Local Area Connection 2:

```
Windows IP Configuration
```

```
Host Name . . . . . : ndserv10
Primary Dns Suffix . . . . . : rtp.raleigh.ibm.com
Node Type . . . . . : Unknown
IP Routing Enabled. . . . . : No
```

```
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : rtp.raleigh.ibm.com
```

Ethernet adapter Local Area Connection 2:

```
Connection-specific DNS Suffix . :
Description . . . . . : Microsoft Loopback Adapter
Physical Address. . . . . : 02-00-4C-4F-4F-50
DHCP Enabled. . . . . : No
IP Address. . . . . : 9.42.92.158
Subnet Mask . . . . . : 255.255.252.0
IP Address. . . . . : 9.42.92.159
Subnet Mask . . . . . : 255.255.252.0
IP Address. . . . . : 2002:92a:8f7a:162:9:42:92:160
IP Address. . . . . : 2002:92a:8f7a:162:9:42:92:159
IP Address. . . . . : fe80::4cff:fe4f:4f50%4
Default Gateway . . . . . :
DNS Servers . . . . . : 127.0.0.1
                        fec0:0:0:ffff::1%1
                        fec0:0:0:ffff::2%1
                        fec0:0:0:ffff::3%1
```

2. Add the cluster address to the loopback using the netsh command. For example:

```
netsh interface ipv6 add address "Local Area Connection 2"
    2002:92a:8f7a:162:9:42:92:161
```

Note: If you are using a high-availability configuration, and the machine is running as the primary machine, do not alias to the loopback device because this scenario prevents traffic to the cluster address from being routed by the Load Balancer machine.

3. Issue the following ipconfig /all command again, and you should see the address added on the loopback adapter. For example, issue the following command:

```
ipconfig /all
```

You should see output that is similar to the following:

Ethernet adapter Local Area Connection 2:

```
Connection-specific DNS Suffix . :
Description . . . . . : Microsoft Loopback Adapter
Physical Address. . . . . : 02-00-4C-4F-4F-50
DHCP Enabled. . . . . : No
IP Address. . . . . : 9.42.92.158
Subnet Mask . . . . . : 255.255.252.0
IP Address. . . . . : 9.42.92.159
Subnet Mask . . . . . : 255.255.252.0
IP Address. . . . . : 2002:92a:8f7a:162:9:42:92:161
IP Address. . . . . : 2002:92a:8f7a:162:9:42:92:160
IP Address. . . . . : 2002:92a:8f7a:162:9:42:92:159
IP Address. . . . . : fe80::4cff:fe4f:4f50%4
Default Gateway . . . . . :
DNS Servers . . . . . : 127.0.0.1
                        fec0:0:0:ffff::1%1
                        fec0:0:0:ffff::2%1
                        fec0:0:0:ffff::3%1
```

4. Enable forwarding for all the interfaces in the machine using the netsh interface ipv6 show interface command. Any interfaces listed with a name of Local Area Connection must have IP forwarding enabled. For example:

```
netsh interface ipv6>show interface
Querying active state...
```

```
Idx Met MTU State Name
```

```

---  ---  -----  -----  -----
6    2    1280  Disconnected  Teredo Tunneling Pseudo-Interface
5    0    1500  Connected     Local Area Connection
4    0    1500  Connected     Local Area Connection 2
3    1    1280  Connected     6to4 Pseudo-Interface
2    1    1280  Connected     Automatic Tunneling Pseudo-Interface
1    0    1500  Connected     Loopback Pseudo-Interface

```

```

netsh interface ipv6>set interface "Local Area Connection"
    forwarding=enabled
Ok.

```

```

netsh interface ipv6>set interface "Local Area Connection 2"
    forwarding=enabled
Ok.

```

5. Verify that the forward packets for each Local Area Connection is set to "Yes." Use the following commands:

```

netsh interface ipv6>show interface "Local Area Connection"

```

```

netsh interface ipv6>show interface "Local Area Connection 2"

```

Related concepts:

"Configuring loopbacks with alternative methods"

Related tasks:

"Configuring the Dispatcher machine" on page 69

Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

"Configuring the server machines" on page 73

Configuring loopbacks with alternative methods

About this task

Some versions of Linux systems issue ARP responses for any IP address configured on the machine on any interface present on the machine. It may also choose an ARP source IP address for ARP who-has queries based on all IP addresses present on the machine, regardless of the interfaces on which those addresses are configured. This causes all cluster traffic to be directed to a single server in an indeterminate manner.

With Dispatcher's forwarding method, a mechanism must be employed to ensure that cluster-addressed traffic can be accepted by the stacks of the back-end servers.

In most cases, you must alias the cluster address on the loopback; therefore, back-end servers must have the cluster aliased on the loopback. To ensure that Linux systems do not advertise addresses on the loopback, you can use any of these four solutions to make Linux systems compatible.

Procedure

- Use a kernel that does not advertise the addresses. This is the preferred option, as it does not incur a per-packet overhead and it does not require per-kernel reconfiguration.
 - Use the `arp_ignore` sysctl available in 2.4.25 and 2.6.5 and higher, but note that distributions sometimes backport features. Ensure that it is enabled before aliasing the cluster addresses with the commands:

```

# sysctl -w net.ipv4.conf.all.arp_ignore=3
net.ipv4.conf.all.arp_announce=2

```

Clusters must then be aliased with the following command:

```
# ip addr add $CLUSTER_ADDRESS/32 scope host dev lo
```

Note: When using `sysctl`, ensure that these settings survive reboot by adding the settings to the `install_root/etc/sysctl.conf` file.

- Use IP tables to redirect all incoming cluster traffic to the localhost. If you use this method, do not configure the loopback adapter with an alias. Instead, use the command:

```
# iptables -t nat -A PREROUTING -d $CLUSTER_ADDRESS -j REDIRECT
```

This command causes Linux systems to do destination NAT on each packet, converting the cluster address to the interface address. This method has about a 6.4% connections-per-second throughput penalty. This method works on any supported stock distribution; no kernel module or kernel patch+build+install is needed.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Configuring the Dispatcher machine” on page 69

Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

“Configuring the server machines” on page 73

Configuring collocated servers for the Dispatcher component

Learn about advanced features for configuring the Dispatcher component with collocated servers.

About this task

Load Balancer can be present on the same server for which it is load balancing requests. This configuration is commonly referred to as collocating a server. It is supported on Linux and AIX® only. It is supported for MAC forwarding only. Collocated server and High Availability record replication should not be used together because after takeover, the newly collocated application server will not have the connection state.

Important: A collocated server competes for resources with Load Balancer during times of high traffic. However, in the absence of overloaded machines, a collocated server offers a reduction in the total number of machines necessary to set up a load-balanced site. On Linux, the collocation feature depends on TUN module. Hence, before you configure collocation, ensure that the TUN module is supported on the Linux version you are using. Also, performance and function are dependent on the TUN module. To configure a server to be collocated, the `dscontrol` server command provides an option called `collocated` which can be set to **yes** or **no**. The default value is **no**. Alternatively, you can use the GUI to set `collocated` parameter under the configuration settings tab of the server object. The address of the server must be a valid IP address of a network interface card on the machine.

Tip:

- When high availability is configured, you must place the appropriate commands in `goScripts`.
- The neighbor state of a collocated server is always available.
- Manager and advisor works as usual with collocated server.
- Multiple clusters can have collocated server.

- Metric server should not be used with collocated server.
- In case you are using IPv4 or IPv6 clusters, the Load Balancer machine should be in both networks before you configure it.

Procedure

1. Configure a collocated server on Linux.

a. Configure a collocated server, as follows:

```
modprobe tun
modprobe ip_tables
```

b. Add or set the server with the preferred IP address and the collocated parameter that is set to **yes**, as follows:

```
dscontrol server add cluster@port@server collocated yes
dscontrol server set cluster@port@server collocated yes
```

Attention: Adding or setting the collocation parameter creates an extra virtual network interface named `ibmulb`, which is used by Load Balancer to forward the packets on same machine. This interface has data rate of 10 MB/s.

c. Add an iptable rule to drop IP packets that are destined to cluster IP. Complete this step so that the operating system does not respond to these packets and bypass the Load Balancer. You might add rules for each cluster IP. Choose the correct interface while you are adding the rules or collocation will not work.

```
iptables -A INPUT -i <interface> -d <cluster_ip>/32 -j DROP
ip6tables -A INPUT -i <interface> -d <cluster_ip>/128 -j DROP
```

Remove the iptable rules by using the following commands:

```
iptables -D INPUT -i <interface> -d <cluster_ip>/32 -j DROP(IPv4)
ip6tables -D INPUT -i <interface> -d <cluster_ip>/128 -j DROP(IPv6)
```

d. ARP suppression (IPv4 only)

```
sysctl -w net.ipv4.conf.all.arp_ignore=3 net.ipv4.conf.all.arp_announce=2
```

e. Alias cluster IP.

Attention: This step is not required for Linux on Z when interface is in layer3. In this case, cluster IP should be aliased on the layer3 interface with a prefix used for the primary IP.

```
ip addr add <cluster_ip>/32 dev ibmulb scope host
ip -6 addr add <cluster_ipv6>/128 dev ibmulb
```

Unalias the cluster IP from the loop back adapter by using following commands:

```
ip addr del <cluster_ip>/32 dev ibmulb scope host(IPv4)
ip -6 addr del <cluster_ipv6>/128 dev ibmulb(IPv6)
```

f. Disable Reverse Path Filtering on all interfaces (IPv4 only), as follows:

```
"sysctl -w net.ipv4.conf.<interface_name>.rp_filter=0 "
```

2. Configure a collocated server on AIX.

a. Add or set the server with your chosen IP address and the collocated parameter that is set to **yes**:

```
dscontrol server add cluster@port@server collocated yes
dscontrol server set cluster@port@server collocated yes
```

b. Configure or alias cluster IP on the loop back adapter:

```

ifconfig lo0 inet <cluster_ip> netmask 255.255.255.255 alias
ifconfig lo0 inet6 <cluster_ipv6>/128 alias

Remove alias from the loop back, as follows:
ifconfig lo0 inet <cluster_ip> netmask 255.255.255.255 -alias
ifconfig lo0 inet6 <cluster_ipv6>/128 -alias

```

Related tasks:

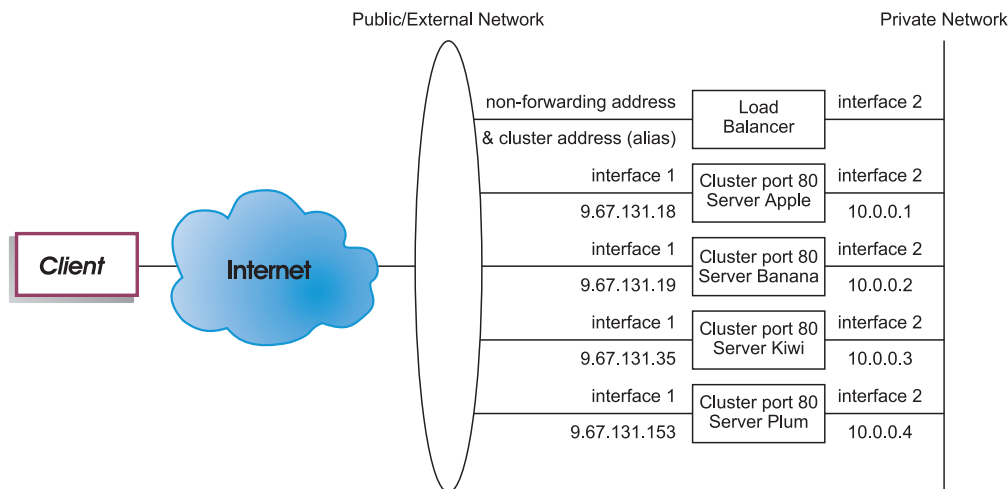
- “Configuring the Dispatcher machine” on page 69
- Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.
- “Configuring the server machines” on page 73

Load balancing a private network

You can set up Dispatcher and the TCP server machines using a private network. This configuration can reduce the contention on the public or external network that can affect performance.

Before you begin

To create a private network, each machine must have at least two LAN cards, with one of the cards connected to the private network. You must also configure the second LAN card on a different subnet.



About this task

When you configure a private network, the Load Balancer machine will then send the client requests to the TCP server machines through this network.

For AIX systems, this configuration can also take advantage of the fast speeds of the SP™ High Performance Switch if you are running Dispatcher and the TCP server machines on nodes in an SP Frame.

Procedure

The servers added using the `dscontrol server add` command must be added using the private network addresses. For example, a sample command could be coded as:

```
dscontrol server add cluster_address@80@10.0.0.1
```

not as

```
dscontrol server add cluster_address@80@9.67.131.18
```

Configuring Site Selector

Configure Site Selector for production environments and realistic test environments.

Quick start configuration for Site Selector

This quick start example shows how to create a site name configuration by using Site Selector to load balance traffic among a set of servers that are based on the domain name that is used on a client request.

Before you begin

. Simple Site Selector configuration.

For this quick start configuration example, you must meet the following conditions:

- Administrative access to the name server for your site
- Four servers (server1, server2, server3, server4) configured to the network and an extra server with the Site Selector component installed

Tip: If you collocate Site Selector on one of the load balanced servers, then you need four servers instead of five. However, collocation has an impact the performance of the load balanced servers.

About this task

With Site Selector, you can create a configuration by using the command line, the configuration wizard, or the graphical user interface (GUI). For this quick start example, configuration steps are demonstrated by using the command line.

Procedure

1. Prepare the servers. For this quick start example, the company site domain is mywebshop.com. Site Selector is responsible for a subdomain within mywebshop.com. Therefore, define a subdomain within mywebshop.com, for example: apps.mywebshop.com.

Site Selector is not a fully implemented DNS, such as BIND, and acts as a leafnode in a DNS hierarchy. Site Selector is authoritative for the apps.mywebshop.com subdomain. Subdomain apps.mywebshop.com includes the following site names: marketing.apps.mywebshop.com and developer.apps.mywebshop.com.

- a. Update the company site's domain name server, as shown in the figure. Create a name server record in the named.data file for the subdomain (apps.mywebshop.com) where Site Selector is the authoritative name server: apps.mywebshop.com. IN NS siteselector.mywebshop.com
- b. Ensure that the fully qualified host name or site does not resolve in the current domain name system.
- c. Install Metric Server on the servers (server1, server2, server3, server4) that you intend to have Site Selector load balance. Read the Getting advice with Metric Server topic to learn how to use Metric Server with the Site Selector component.

2. With Site Selector, you can create a configuration by using the command line, the configuration wizard, or the graphical user interface (GUI). For this quick start example, configuration steps are demonstrated using the command line.

Important: The parameter values must be entered in English characters. The only exceptions are parameter values for host names and file names.

- Configuring with the command line:
 - a. Start **ssserver** from the *install_root/ss/server/bin* directory:
 - Run the following command as root user:

```
ssserver
```
 - Start **ssserver** (IBM Site Selector) from the Services panel: Start > Settings (for Windows 2000) > Control Panel > Administrative Tools > Services.
 - b. Start the name server on the Site Selector configuration. Enter the command

```
sscontrol nameserver start
```
 - c. Configure the site names (marketing.apps.mywebshop.com and developer.apps.mywebshop.com) on Site Selector:

```
sscontrol sitename add marketing.apps.mywebshop.com@
sscontrol sitename add developer.apps.mywebshop.com@
```
 - d. Add the servers to the Site Selector configuration. (Configure server1 and server2 to site name marketing.apps.mywebshop.com. Configure server3 and server4 to site name developer.apps.mywebshop.com):

```
sscontrol server add marketing.apps.mywebshop.com@server1+server2
sscontrol server add developer.apps.mywebshop.com@server3+server4
```
 - e. Start the manager function of Site Selector:

```
sscontrol manager start
```
 - f. Start the advisor function of Site Selector (HTTP advisor for marketing.apps.mywebshop.com and FTP advisor for developer.apps.mywebshop):

```
sscontrol advisor start http marketing.apps.mywebshop.com@80
sscontrol advisor start ftp developer.apps.mywebshop.com@21
```
 - g. Ensure that the Metric Server is started on each of the load-balanced servers.
- 3. Test your configuration.
 - a. From a client, that has a primary DNS configured as the name server responsible for mywebshop.com, try to ping one of the configured site names.
 - b. Connect to the application. For example:
 - Open a browser, request marketing.apps.mywebshop.com and a valid page is served
 - Open an FTP client to developer.apps.mywebshop.com and enter a valid user and password
 - c. Look at the results of the following command:

```
sscontrol server status marketing.apps.mywebshop.com@
sscontrol server status developer.apps.mywebshop.com@
```

The total hits entry of each server adds up to the ping and application request.

What to do next

You can also configure Site Selector with the graphical user interface (GUI) and the configuration wizard. Read the Configuring Site Selector and GUI: General Instructions topics for more information.

Related concepts:

“Methods of configuration for Site Selector” on page 88

There are four methods you can use to configure Site Selector: the command line, scripts, the graphical user interface (GUI), and the configuration wizard. This topic assumes use of the command line.

Related tasks:

“Getting advice with Metric Server” on page 121

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

Related reference:

“GUI: General instructions” on page 209

Load Balancer graphical user interface.

Plan for Site Selector

Site Selector works with a domain name server to load balance among a group of servers that uses measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers that are based on the domain name that is used for a client request.

CAUTION:

The DNS queries that Site Selector supports are type A queries only. Any other query types results in a return code of NOTIMPL (Not Implemented). If an entire domain is delegated to Site Selector, ensure that the domain receives type A queries only.

. Example of a DNS environment.

When you set up a subdomain for Site Selector within your DNS environment, ensure that Site Selector has authority over its own subdomain. For example, your company is granted authority over `company.com` domain, as shown in the Example of a DNS environment figure. Within the company, there are several subdomains. Site Selector would have authority for `siteload.company.com`, while the DNS servers would still maintain authority for `atlanta.company.com` and `boston.company.com`.

For a company name server to recognize Site Selector as having authority for the `siteload` subdomain, add a name server entry to its named data file. For example, on AIX systems, a name server entry would look like the following example:

```
siteload.company.com IN NS siteselector.company.com
```

Where *siteselector.company.com* is the host name of the Site Selector server. Equivalent entries must be made in any other named database files for use by DNS servers.

A client submits a request for resolution of a domain name to a name server within its network. Name server forwards the request to the Site Selector server. Site Selector then resolves the domain name to the IP address of one of the servers that is configured under the site name. Site Selector returns the IP address of the

selected server to the name server. Name server returns the IP address to the client. (Site Selector acts as a non-recursive (leaf node) name server, and it returns an error if it does not resolve the domain name request.)

Refer to the figure in the Site Selector component topic, which illustrates a site in which Site Selector is used with a DNS system to load balance across local and remote servers.

Site Selector consists of the following functions:

- **sssserver** handles request from the command line to the name server, manager, and advisors.
- The name server function supports the load balancing of incoming name server requests. You must start the name server function for Site Selector to begin providing DNS resolution. Site Selector listens on port 53 for incoming DNS requests. If the requesting site name is configured, then Site Selector returns a single server address (from a set of server addresses) associated with the site name.
- The manager sets weights that are used by the name server that is based on:
 - Feedback from the servers that are provided by the advisors
 - Feedback from a system-monitoring program, such as Metric Server.

Using the manager is optional. However, if the manager is not used, load balancing uses weighted round-robin scheduling that is based on the current server weights, and advisors are not available.

- The Metric Server is a system monitoring component of Load Balancer that you install on the backend server. If you collocate Load Balancer on a server that is being load balanced, then you would install Metric Server on the Load Balancer server.

With Metric Server, Site Selector can monitor the level of activity on a server, detect when a server is the least heavily loaded, and detect a failed server. The load is a measure of how hard the server is working. The system Site Selector administrator controls the type of measurement that is used to measure the load. You can configure Site Selector to suit your environment, considering such factors as frequency of access, the total number of users, and types of access (for example, short queries, long-running queries, or CPU-intensive loads).

Load balancing is based on server weights. For Site Selector, there are four proportions that the manager uses to determine weights:

- CPU
- memory
- port
- system

CPU and memory values are all supplied by Metric Server, which is recommended with the Site Selector component. Read the Getting advice with Metric Server topic to learn how to use Metric Server with the Site Selector component.

- The advisors query the servers and analyze results by protocol before calling the manager to set weights as appropriate. It might not make sense to use some of these advisors in a typical configuration. You also have the option of writing your own advisors. Using the advisors is optional, but recommended. Read the List of advisors topic to learn about the advisors that are provided with Load Balancer.

- To configure and manage the name server, advisors, Metric Server, and manager, use the command line (**sscontrol**) or the graphical user interface (**ladmin**).

The four key functions of Site Selector (name server, manager, Metric Server, and advisors) interact to balance and resolve the incoming requests between servers.

TTL

Using DNS-based load balancing requires that caching of name resolutions be disabled. The TTL (time to live) value determines the effectiveness of DNS-based load balancing. TTL determines how long another name server caches the resolved response. Small TTL values allow for subtle changes in the server or network load to be realized more quickly. However, disabling caching requires that clients contact the authoritative name server for every name resolution request, thus potentially increasing the client latency. When you choose a TTL value, carefully consider the impact that disabled-caching has on an environment. Also, be aware that DNS-based load balancing is potentially limited by client-side caching of name resolutions.

Configure TTL by using the **sscontrol sitename [add | set]** command.

Network Proximity feature

Network proximity is the calculation of each server's nearness to the requesting client. To determine network proximity, the Metric Server agent (which must be present on each load-balanced server) sends a ping to the client IP address and returns the response time to Site Selector. Site Selector uses the proximity response in the load-balancing decision. Site Selector combines the network proximity response value with the weight from the manager to create a combined final weight value for the server.

Use of the network proximity feature with Site Selector is optional.

The Site Selector provides the following network proximity options that can be set per site name:

- **Cache life:** The amount of time a proximity response is valid and saved in the cache.
- **Proximity percent:** The importance of the proximity response versus the health of the server (as input from the manager weight).
- **Wait for all:** Determines whether to wait for all proximity (ping) responses from the servers before responding to the client request.

If set to **yes**, the Metric Server pings the client to obtain the proximity response time. Name server waits for all Metric Servers to respond or for a timeout to occur. Then, for each server, the name server combines the proximity response time with the weight the manager calculated to create a combined weight value for each server. Site Selector supplies the client with the server IP address with the best combined weight. It is expected that most client name servers have a 5-second timeout. Site Selector tries to respond before that time out is exceeded.

If set to **no**, a name resolution is provided to the client based on the current manager weights. Then, the Metric Server pings the client to obtain the proximity response time. The name server caches the response time that it receives from the Metric Server. When the client returns for a second request, the name server combines the current manager weight with the cached ping

response value for each server to obtain the server with the best combined weight. Site Selector returns the IP address for this server to the client for its second request.

Set network proximity options with the **sscontrol sitename [add | set]** command. Read the Commands reference topic to learn about Site Selector commands.

Related concepts:

“Site Selector component” on page 19

Site Selector acts as a name server that works with other name servers in a domain name system to load balance among a group of servers by using measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers based on the domain name used for a client's request.

Related tasks:

“Getting advice with Metric Server” on page 121

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

“List of advisors” on page 114

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

“sscontrol sitename” on page 270

Use the **sscontrol sitename** command to configure a site name.

Methods of configuration for Site Selector

There are four methods you can use to configure Site Selector: the command line, scripts, the graphical user interface (GUI), and the configuration wizard. This topic assumes use of the command line.

There are four basic methods of configuring the Site Selector component:

- “Command line”
- “Scripts” on page 89
- Graphical User Interface (GUI)
- “The configuration wizard” on page 90.

Command line

The command line is the most direct means of configuring Site Selector. The command parameter values must be entered in English characters. The only exceptions are host names (used, for example, in site name and server commands) and file names.

To start the Site Selector from the command line:

1. Start **sserver**:

- Go to the `<install_root>/ss/servers/bin` directory and issue the following command from the command prompt:
sserver

To stop the service, type:

```
sserver stop
```

- Click **Start > Control Panel > Administrative Tools > Services**. Right-click **IBM Site Selector** and select **Start**. To stop the service, follow the same steps and select **Stop**.
2. Issue the Site Selector control commands you want to set up your configuration. The command is **sscontrol**.

You can use a minimized version of the **sscontrol** command parameters by typing the unique letters of the parameters. For example, to get help on the file save command, you can type

```
sscontrol he f
```

instead of

```
sscontrol help file
```

To start the command-line interface issue **sscontrol** to receive an **sscontrol** command prompt. To end the command-line interface issue the `exit` or `quit` commands.

Scripts

You can enter commands for configuring Site Selector into a configuration script file and run them together. See Sample Load Balancer configuration files. To quickly run the content of a script file (for example, `myscript`), use one of the following commands:

- To update the current configuration, run the following executable commands from your script file:

```
sscontrol file appendload myscript
```
- To completely replace the current configuration, run the following executable commands from your script file:

```
sscontrol file newload myscript
```

To save the current configuration into a script file (for example, `savescript`), run the following command:

```
sscontrol file save savescript
```

This command saves the configuration script file in the `install_root/servers/configurations/ss` directory.

GUI

Start the GUI, as follows:

1. Ensure **ssserver** is running:
 - Run the following command as root user:

```
ssserver
```
 - **ssserver** runs as a service that starts automatically.
2. Start the Site Selector GUI:
 - Run the following command:

```
lbadmin
```
 - Click **Start > Programs > IBM WebSphere > Edge Components > IBM Load Balancer > Load Balancer**

To configure the Site Selector component from the GUI, you must first select **Site Selector** in the tree structure. You can start the executor and manager after you connect to a **Host**. You can also create clusters that contain ports and servers, and start advisors for the manager.

You can use the GUI to do anything that you would do with the **sscontrol** command. For example, to define a cluster by using the command line, you would enter `dscontrol cluster add cluster` command. To define a cluster from the GUI, right-click **Executor**, then in the menu, left-click **Add Cluster**. Enter the cluster address in the window, then click **OK**.

You can load pre-existing Site Selector configuration files by using the **Load New Configuration** (for completely replacing the current configuration) and **Append to Current Configuration** (for updating the current configuration) options that are presented in the **Host** menu. Save your Site Selector configuration to a file periodically by using the **Save Configuration File As** option that is also presented in the **Host** menu. The **File** menu that is at the top of the GUI enables you to save your current host connections to a file or restore connections in existing files across all Load Balancer components.

Run a command from the GUI, as follows:

1. Highlight the host node from the GUI tree and select **Send command...** from the host menu.
2. In the command entry field, enter the command that you want to run, for example:
`executor report`

The results and history of the commands run in the current session and display in the window provided.

You can access help by clicking the question mark icon in the upper right corner of the Load Balancer window.

- **Help: Field level** describes each field, default values
- **Help: How do I** lists tasks that can be done from that screen
- **InfoCenter** provides centralized access to product information

The configuration wizard

The wizard guides you step by step through the process of creating a basic configuration for the Site Selector component. You are asked questions about your network and you are guided through the setup of a cluster for Site Selector to load balance traffic between a group of servers.

If you are using the configuration wizard, follow these steps:

1. Start the **ssserver** on Site Selector:
 - Run the following command as root user:
`ssserver`
 - **ssserver** runs as a service that starts automatically.
2. Start the wizard function of Site Selector by issuing the following command:
`sswizard`

Related tasks:

“Setting up Site Selector”

Set up the Site Selector component in production environments and realistic test environments.

Setting up Site Selector

Set up the Site Selector component in production environments and realistic test environments.

Before you begin

Before you set up the Site Selector server, you must be the root user (for AIX, HP-UX, Linux, or Solaris systems) or the Administrator (for Windows systems).

About this task

You need an unresolvable fully qualified host name to use as a site name for a group of servers that you set up. The site name is the name that the clients use to access your site (such as `www.yourcompany.com`). Site Selector load-balances traffic for this site name among the group of servers by using DNS

Procedure

1. Start the server function by typing **ssserver** at the command prompt. A default configuration file, `default.cfg`, automatically loads when you start **ssserver**. If you decide to save the configuration in the `default.cfg` file, everything that is saved in this file automatically loads the next time **ssserver** starts.
2. Start the name server by typing **sscontrol nameserver start** at a command prompt. Optionally, start the name server by using the `bindaddress` keyword to bind only to the specified address.

3. Define a site name and set site name options. Site Selector balances the requests that are sent for the site name to the corresponding servers configured to it.

The site name is an unresolvable host name that the client requests. The site name must be a fully qualified domain name (for example, `www.dnsdownload.com`). When a client requests this site name, one of the server IP addresses associated with the site name is returned.

To define a site name, issue the following command:

```
sscontrol sitename add sitename
sscontrol sitename set sitename option value
```

4. Define load balanced server workstations. The server workstations run the applications that you want to load balance. The server is the symbolic name or dotted decimal address of the server machine. To define a server on the site name from step 3, issue the following command:

```
sscontrol server add sitename@server
```

Define more than one server under a site name to perform load balancing.

5. Optional: Start the manager function (optional). The manager function enhances load balancing. Before starting the manager function, ensure that the metric server is installed in all the load-balanced machines.

To start the manager, issue the following command:

```
sscontrol advisor start http sitename@port
```

6. Optional: Start the advisor function (optional). The advisors give the manager more information about the ability of the load balanced server machines to respond to requests. An advisor is specific to a protocol. The Load Balancer

supplies many advisors. For example, to start the HTTP advisor for a specific site name, issue the following command:

```
sscontrol advisor start http sitename@port
```

7. Optional: Define system metric (optional).
8. Set site name proportions as required. If you start advisors, you might modify the proportion of importance that is given to advisor (port) information that is included in the load balancing decisions. To set the site name proportions, issue the following command:

```
sscontrol sitename set sitename proportions
```

Read the Tuning the proportion of importance that is given to status information to learn how the manager uses ratios to determine the importance of status information that is coming from advisors and Load Balancer.

What to do next

Set up the Metric Server on all machines that Site Selector is load balancing. Read the Getting advice with Metric Server topic to learn how to use Metric Server with the Site Selector component.

Related tasks:

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

“Configuring the server machines” on page 73

“Getting advice with Metric Server” on page 121

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

“Tuning the proportion of importance given to status information” on page 178

The manager uses ratios to determine the importance of status information coming from advisors and Load Balancer. You can change the default ratios that the manager uses to weight this information.

Configuring Content Based Routing

Configure Content Based Routing (CBR) for production environments and realistic test environments.

This section provides information on a quick start configuration, planning considerations, and describes the methods of configuring the Load Balancer CBR component.

Quick start configuration for Content Based Routing

This quick start example shows how to configure three locally attached workstations using Content Based Routing (CBR) along with Caching Proxy, to load-balance web traffic between two web servers.

Before you begin

Attention: For simplicity, this example illustrates servers on the same LAN segment; however, with CBR, there is no restriction for using servers on the same LAN.

You need three workstations and four IP addresses for this quickstart example. One workstation is the CBR machine; the other two workstations are the web servers. Each web server requires one IP address. The CBR workstation requires one actual addresses and one address to be load balanced.

Simple local CBR configuration

For more information on the different ways Load Balancer can be setup, read the Types of cluster, port, and server configurations for Content Based Routing topic to help you design your topology.

Caching Proxy should be installed on the same server to use CBR. To configure Caching Proxy for CBR, read step one of the Setting up the Content Based Routing component topic.

About this task

Use this configuration method for a quick way to establish a connection between servers and the CBR machine. This method does not include configuring advisors or tuning performance. For a full configuration, read Configuring the Load Balancer machine and Configuring the server machines topics.

Procedure

1. Prepare your servers.
 - a. For this example, set up your workstations on the same LAN segment. Ensure that network traffic between the three machines does not have to pass through any routers or bridges.
 - b. Configure the network adapters of the three workstations. For this example, we will assume you have the following network configuration:

Table 6. Sample network configuration

Workstation	Name	IP Address
1	server1.mywebsite.com	9.27.27.101
2	server2.mywebsite.com	9.27.27.102
3	server3.mywebsite.com	9.47.47.103
Netmask = 255.255.255.0		

Each of the workstations contains only one standard Ethernet network interface card.

- c. Ensure that all the servers can communicate with each other.
 - 1) Ensure that server1.mywebsite.com can ping both server2.mywebsite.com and server3.mywebsite.com.
 - 2) Ensure that server2.mywebsite.com and server3.mywebsite.com can ping server1.mywebsite.com.
- d. Ensure that web servers on server2.mywebsite.com and server3.mywebsite.com are operational. Use a web browser to request pages directly from http://server2.mywebsite.com (for example, .../member/index.html) and http://server3.mywebsite.com (for example, .../guest/index.html).
- e. Obtain another valid IP address for this LAN segment. This is the cluster address you will provide to clients who wish to access your site. For this example we will use:

Name= www.mywebsite.com
IP=9.27.27.104

2. Configure CBR using the command line, the GUI, or the configuration wizard. With CBR, you can create a configuration by using the command line, the configuration wizard, or the graphical user interface (GUI). For this quick start example, configuration steps are demonstrated using the command line.

Important: You must enter the parameter values in English characters. The only exceptions are parameter values for host names and file names.

- Configuring with the command line:
 - a. Start **cbserver**:
 - Go to the `<install_root>/cbr/servers/bin` directory and run the following command as root user:
`cbserver`
 - Start **cbserver** (Content Based Routing) from the Services panel: Start > Settings (for Windows 2000) > Control Panel > Administrative Tools > Services.
 - b. Start the executor function of CBR. Enter the command
`cbrcontrol executor start`
 - c. Start Caching Proxy. You can start caching proxy any time after you start the executor function.
 - Run the following command:
`ibmproxy`
 - from the Services panel: Start > Settings (for Windows 2000) > Control Panel > Administrative Tools > Services.
 - d. Add the cluster (the host name and web site to which clients connect) to the CBR configuration:
`cbrcontrol cluster add www.mywebsite.com`
 - e. Add the cluster address (9.27.27.104) for the web site to the network interface card on the CBR machine.
 - f. Add the http protocol port to the CBR configuration:
`cbrcontrol port add www.mywebsite.com@80`
 - g. Add each of the web servers to the CBR configuration:
`cbrcontrol server add www.mywebsite.com@80@server2.mywebsite.com cbrcontrolserver add www.mywebsite.com@80@server3.mywebsite.com`
 - h. Add content rules to your CBR configuration. A content rule defines how a URL request is distinguished and sent to one of the servers or server sets:
`cbrcontrol rule add www.mywebsite.com@80@memberRule type content pattern uri=*/member/*`
`cbrcontrol rule add www.mywebsite.com@80@guestRule type content pattern uri=*/guest/*`
In this example, using the content rule, client requests to `www.mywebsite.com` are sent to a different server, based on a directory in their URI request path.
 - i. Add servers to your rules:
`cbrcontrol rule useserver www.mywebsite.com@80@memberRule server2.mywebsite.com`
`cbrcontrol rule useserver www.mywebsite.com@80@guestRule server3.mywebsite.com`
CBR will now load balance based on content-based rule. A client with a URL request containing `/member/` is directed to `server2.mywebsite.com`. A client with a URL request containing `/guest/` is directed to `server3.mywebsite.com`.
 - j. Start the manager function of CBR:

```
cbrcontrol manager start
```

- k. Start the advisor function of CBR:

```
cbrcontrol advisor start http 80
```

CBR will now make sure that client requests are not sent to a failed web server.

Your basic configuration with locally attached servers is now complete.

3. Test your configuration.

- a. From a web browser, enter `http://www.mywebsite.com/member/index.htm`. If a page displays, the configuration is working.
- b. Reload the page in the web browser.
- c. Look at the results of the following command:

```
cbrcontrol server report www.mywebsite.com@80@
```

The total connections column of the two servers should add up to 2.

What to do next

You can also configure CBR with the graphical user interface (GUI) and the configuration wizard. Read the Methods of configuration for Content Based Routing topic for more information.

Related concepts:

“Methods of configuration for Content Based Routing” on page 98

Learn about the methods that you can use to configure Content Based Routing (CBR). The information center assumes use of the command line.

Related reference:

“Types of cluster, port, and server configurations for Content Based Routing” on page 102

There are many ways that you can configure Content Based Routing (CBR) to support your site.

“Setting up the Content Based Routing component” on page 105

Set up the Content Based Routing (CBR) component in production environments and realistic test environments.

“GUI: General instructions” on page 209

Load Balancer graphical user interface.

“Configuring the Dispatcher machine” on page 69

Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

“Configuring the server machines” on page 73

Plan for Content Based Routing

This topic describes what the network planner might consider when installing and configuring the CBR component with Caching Proxy.

The CBR component load balances HTTP and SSL traffic with Caching Proxy to proxy the request. With CBR, you can load balance servers that you configure from your CBR configuration file with **cbrcontrol** commands.

CBR is similar to Dispatcher in its component structure. CBR consists of the following functions:

- **cbrserver** handles requests from the command line to the executor, manager, and advisors.

- The executor supports load balancing of client requests. You must start the executor to use the CBR component.
- The manager sets weights that are used by the executor that is based on:
 - Internal counters in the executor
 - Feedback from the servers that are provided by the advisors
 - Feedback from a system-monitoring program, such as Metric Server.

Using the manager is optional. However, if the manager is not used, load balancing uses weighted, round-robin scheduling, which is based on the current server weights. Advisors are not available.

- The advisors query the servers and analyze results by protocol before calling the manager to set weights as appropriate. It might not make sense to use some of these advisors in a typical configuration. You also have the option of writing your own advisors. Using the advisors is optional, but recommended. Load Balancer provides a Caching Proxy (**cachingproxy**) advisor. Read the List of advisors topic for information about advisors in Load Balancer.
- To configure and manage the executor, advisors, and manager, use the command line (**cbrcontrol**) or the graphical user interface (**lbadmin**).

The three key functions of CBR, **executor**, **manager**, and **advisors**, interact to balance and dispatch the incoming requests between servers. Along with load balancing requests, the executor monitors the number of new connections and active connections and supplies this information to the manager.

Load balancing requests for different types of content

The CBR component gives you the ability to specify a set of servers that handle a request that is based on regular expression matching the content of the client request. You can use CBR to partition your site so that different content or application services can be served by different sets of servers. This partitioning is not apparent to clients that access your site.

Dividing your site content for better response time

One way to divide your site is by assigning some servers to handle only CGI requests, and assigning another set of servers to handle all other requests. Dividing your site stops compute intensive CGI scripts from slowing down the servers for normal HTML traffic, allowing clients to get better overall response time. Using this scheme, you might also assign more powerful workstations for normal requests. This assignment gives clients better response time without the expense of upgrading all your servers. You might also assign more powerful workstations for CGI requests.

Another possibility for partitioning your site might be to direct clients who are accessing pages that require registration to one set of servers, and all other requests to a second set of servers. Partitioning keeps casual browsers of your site from tying up resources that might be used by clients who are committed to your registration. You can also use more powerful workstations to service those clients who are registered.

You might also combine the previous methods for even more flexibility, and improved service.

Providing backup of web server content

Because CBR you can specify multiple servers for each type of request with CBR, you can load balance the requests for optimal client response. By allowing multiple servers to be assigned to each type of content, you are protected if one workstation or server fails. CBR recognizes the failure and continues to load balance client requests to the other servers in the set.

Multiple Caching Proxy processes to improve processor utilization

Caching Proxy communicates with a CBR process through its plug-in interface. CBR must be running on the local machine for this communication to work. Because there are two separate processes, multiple instances of the Caching Proxy can run and work with a single instance of CBR. This setup might be configured to segregate addresses or function between the Caching Proxies, or to improve the resource utilization of the machine by having several Caching Proxies handling client traffic. The proxy instances can be listening on different ports or binding to unique IP addresses on the same port, depending on what best suits the traffic requirements.

Rules-based load balancing with CBR

CBR, along with Caching Proxy, examines HTTP requests that use specified rule types. When Caching Proxy is running, it accepts client requests and queries the CBR component for the best server. Upon this query, CBR matches the request to a set of prioritized rules. When a rule is matched, an appropriate server is chosen from a preconfigured server set. Finally, CBR informs Caching Proxy which server was chosen and the request gets proxied there.

After you define a cluster to be load balanced, make sure that all requests to that cluster have a rule that will choose a server. If no rule is found that matches a particular request, the client receives an error page from Caching Proxy. The easiest way to ensure that all requests match some rule is to create an *always true* rule at a high priority number. Make sure that the servers that use this rule can handle all the requests that are not explicitly handled by the rules that have a lower-numbered priority. The lower-numbered priority rules are evaluated first.

Read the Configure rules to manage traffic to busy or unavailable servers topic to learn about rules-based load balancing.

Load balancing across fully secure (SSL) connections

CBR with Caching Proxy can receive SSL transmission from the client to the proxy (client-to-proxy side) and support transmission from the proxy to an SSL server (proxy-to-server side). By defining an SSL port on a server in the CBR configuration to receive the SSL request from the client, you can maintain a fully secure site that uses CBR to load balance across secure (SSL) servers.

In addition to other `ibmproxy.conf` file changes for CBR, another configuration statement must be added to the `ibmproxy.conf` SSL file for Caching Proxy to enable SSL encryption on the proxy-to-server side. The format must be as follows:

```
proxy uri_pattern url_pattern address
```

where:

- *uri_pattern* is a pattern to match (for example: `/secure/*`)

- *url_pattern* is a replacement URL (for example: `https://clusterA/secure/*`)
- *address* is the cluster address (for example: `clusterA`)

Load balancing client-to-proxy in SSL and proxy-to-server in HTTP

CBR with Caching Proxy can also receive SSL transmission from the client and then decrypt the SSL request before proxying the request to an HTTP server. For CBR to support client-to-proxy in SSL and proxy-to-server in HTTP, there is an optional keyword, `mapport`, on the `cbrcontrol` server command. Use this keyword when you must indicate that the port on the server is different from the incoming port from the client. The following example adds a port that uses the `mapport` keyword. The client's port is 443 (SSL) and the server's port is 80 (HTTP):

```
cbrcontrol server add cluster@443 mapport 80
```

The port number for `mapport` can be any positive integer value. The default is the port number value of the incoming port from the client.

Because CBR must be able to advise on an HTTP request for a server that is configured on port 443 (SSL), a special `ssl2http` advisor is provided. This advisor starts on port 443 (the incoming port from the client) and advises on the servers that are configured for that port. If there are two clusters that are configured and each cluster has port 443 and servers that are configured with a different `mapport`, then a single instance of the advisor can open the appropriate port accordingly. The following is an example of this configuration:

```
Executor
  Port:443
    Server1 mapport 80
    Server2 mapport 8080
  Port:443
    Server3 mapport 80
    Server4 mapport 8080
  Advisor ssl2http 443
```

Related tasks:

“Configure rules to manage traffic to busy or unavailable servers” on page 162
Use rules-based load balancing to fine tune when and why packets are sent to which servers. Load Balancer reviews any rules you add from first priority to last priority, stopping on the first rule that it finds to be true, then load balancing the traffic between any servers associated with the rule. It already balances the load based on the destination and port, but using rules expands your ability to distribute connections.

Related reference:

“List of advisors” on page 114
Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

Methods of configuration for Content Based Routing

Learn about the methods that you can use to configure Content Based Routing (CBR). The information center assumes use of the command line.

There are four basic methods of configuring the Content Based Routing component:

- “Command line” on page 99
- “Scripts” on page 100

- Graphical User Interface (GUI)
- “The configuration wizard” on page 102.

Restriction: To use CBR, Caching Proxy must be installed. Caching Proxy is a service that starts automatically by default after installation. Stop Caching Proxy before starting the CBR server function (**cbserver**) and modify the Caching Proxy service to start manually, rather than automatically.

- For Linux or UNIX systems: Stop Caching Proxy by finding its process identifier by using the **ps -ef | grep ibmproxy** command and then ending the process by using the **kill process_id** command.
- For Windows systems: Stop Caching Proxy from the Services panel.

Command line

The command line is the most direct means of configuring Content Based Routing. The command parameter values must be entered in English characters. The only exceptions are host names (used, for example, in cluster and server commands) and file names.

Start the CBR from the command line, as follows:

1. Start **cbserver**, as follows:

- Go to the `<install_root>/cbr/servers/bin` directory and issue the following command from the command prompt:

```
cbserver
```

To stop the service, enter the following command:

```
cbserver stop
```

- Click **Start > Control Panel > Administrative Tools > Services**. Right-click **IBM Content Based Routing** and select **Start**. To stop the service, follow the same steps and select **Stop**.
2. Issue the CBR control commands you want in order to set up your configuration. The command is **cbcontrol**. Read the Commands topic for more information about commands
 3. Start Caching Proxy by entering the **ibmproxy** command at a command prompt. You must start the executor prior to starting Caching Proxy.
On Windows platforms, click **Start > Settings** (for Windows 2000) **> Control Panel > Administrative Tools > Services**.

You can use a minimized version of the **cbcontrol** command parameters by typing the unique letters of the parameters. For example, to get help on the file save command, you can type

```
cbcontrol he f
```

instead of

```
cbcontrol help file
```

To start the command line interface issue **cbcontrol** to receive an **cbcontrol** command prompt. To end the command line interface issue the **exit** or **quit** commands.

Tip: On Windows platforms, the **dsserver** Dispatcher component starts automatically. If you are using only CBR and not the Dispatcher component, you can stop **dsserver** from starting automatically as follows:

1. In the Services window, right-click **IBM Dispatcher**.
2. Select **Properties**.
3. In the **Startup type** field, select **Manual**.
4. Click **OK**, and close the Services window.

Tip: When you configure Content Based Routing (CBR) from the operating system command prompt, rather than from the **cbrcontrol>>** prompt, be careful when you use the following characters:

- () right and left parentheses
- & ampersand
- | vertical bar
- ! exclamation point
- * asterisk

The operating system shell might interpret these characters as special characters and convert them to alternate text before **cbrcontrol** evaluates them.

The special characters are optional characters on the **cbrcontrol rule add** command and are used to specify a pattern for a content rule. For example, the following command might be valid only when you use the **cbrcontrol>>** prompt:

```
rule add 10.1.203.4@80@cbr_prod_rule_ek type content pattern uri=/nipoeck/*
```

For this same command to work at the operating system prompt, you must place double quotation marks (␣␣) around the pattern as follows:

```
cbrcontrol rule add 10.1.203.4@80@cbr_prod_rule_ek type content pattern "uri=/nipoeck/*"
```

If the quotation marks are not used, some of the pattern might be truncated when the rule is saved in CBR. Quotation marks are not supported when you use the **cbrcontrol>>** command prompt.

Scripts

You can go to the `<install_root>/cbr/servers/bin` directory and enter commands for configuring Content Based Routing into a configuration script file and run them together. To quickly run the content of a script file (for example, *myscript*), use either of the following commands:

- To update the current configuration, run the following executable commands from your script file:

```
cbcontrol file appendload myscript
```

- To completely replace the current configuration, run the following executable commands from your script file:

```
cbcontrol file newload myscript
```

Save the current configuration into a script file (for example, *savescript*) by running the following command:

```
cbcontrol file save savescript
```

This command saves the configuration script file in the `install_root/servers/configurations/cbr` directory.

GUI

To start the GUI, follow these steps:

1. Ensure **cbserver** is running, as follows:
 - Run the following as root user by issuing the following command from the `<install_root>/cbr/servers/bin` directory:
`cbserver`
 - **cbserver** runs as a service that starts automatically.
2. Start the GUI:
 - Run the following command:
`lbadmin`
 - Click **Start > Programs > IBM WebSphere > Edge Components > IBM Load Balancer > Load Balancer**
3. Start Caching Proxy, as follows:
 - **ibmproxy**
 - Go to the Services panel by clicking **Start > Settings (for Windows 2000) > Control Panel > Administrative Tools > Services**

To configure the CBR component from the GUI, you must first select **Content Based Routing** in the tree structure. You can start the executor and manager after you connect to a **Host**. You can also create clusters that contain ports and servers, and start advisors for the manager.

The GUI can be used to do anything that you would do with the `cbrcontrol` command. For example, to define a cluster with the command line, you would enter `cbcontrol cluster add cluster` command. To define a cluster from the GUI, right-click **Executor**, then in the menu, left-click **Add Cluster**. Enter the cluster address in the window, then click **OK**.

You can load pre-existing CBR configuration files by using the **Load New Configuration** (for completely replacing the current configuration) and **Append to Current Configuration** (for updating the current configuration) options that are presented in the **Host** menu. Save your CBR configuration to a file periodically by using the **Save Configuration File As** option that is also presented in the **Host** menu. The **File** menu that is at the top of the GUI allows you to save your current host connections to a file or restore connections in existing files across all Load Balancer components.

You can access Help by clicking the question mark icon in the upper right corner of the Load Balancer window.

- **Help: Field level** describes each field, default values
- **Help: How do I** lists tasks that can be done from that screen
- **InfoCenter** provides centralized access to product information

Run a command from the GUI, as follows:

1. Highlight the Host node from the GUI tree and select **Send command...** from the host menu.
2. In the command entry field, type the command that you want to run, for example:
`executor report`

The results and history of the commands run in the current session and display in the window provided.

The configuration wizard

If you are using the configuration wizard, follow these steps:

1. Go to the `<install_root>/cbr/servers/bin` directory and issue **cbserver** at the command prompt as root user or administrator.
2. Start the wizard function of Content Based Routing by issuing the **cbwizard** command or select the Configuration wizard from the **CBR component** menu, as presented in the GUI.
3. Start Caching Proxy to load balance HTTP or HTTPS (SSL) traffic, as follows:
 - **ibmproxy**
 - Start Caching Proxy on the Services panel by clicking **Start > Settings (for Windows 2000) > Control Panel > Administrative Tools > Services**.

The CBR wizard guides you step-by-step through the process of creating a basic configuration for the CBR component. The wizard asks you questions about your network and guides you as you set up a cluster that enables CBR to load balance traffic between a group of servers.

Related concepts:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

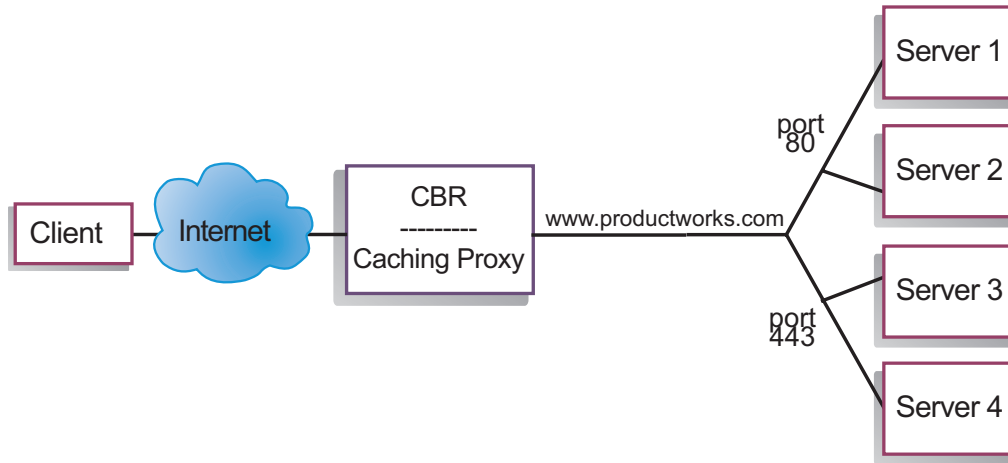
Types of cluster, port, and server configurations for Content Based Routing

There are many ways that you can configure Content Based Routing (CBR) to support your site.

There are many ways that you can configure CBR to support your site. If you have only one host name for your site to which all of your customers will connect, you can define a single cluster of servers. For each of these servers, you configure a port through which CBR communicates.

1 cluster with 2 ports

If you have only one host name for your site to which all of your customers will connect, you can define a single cluster of servers. For each of these servers, configure a port through which Load Balancer communicates.

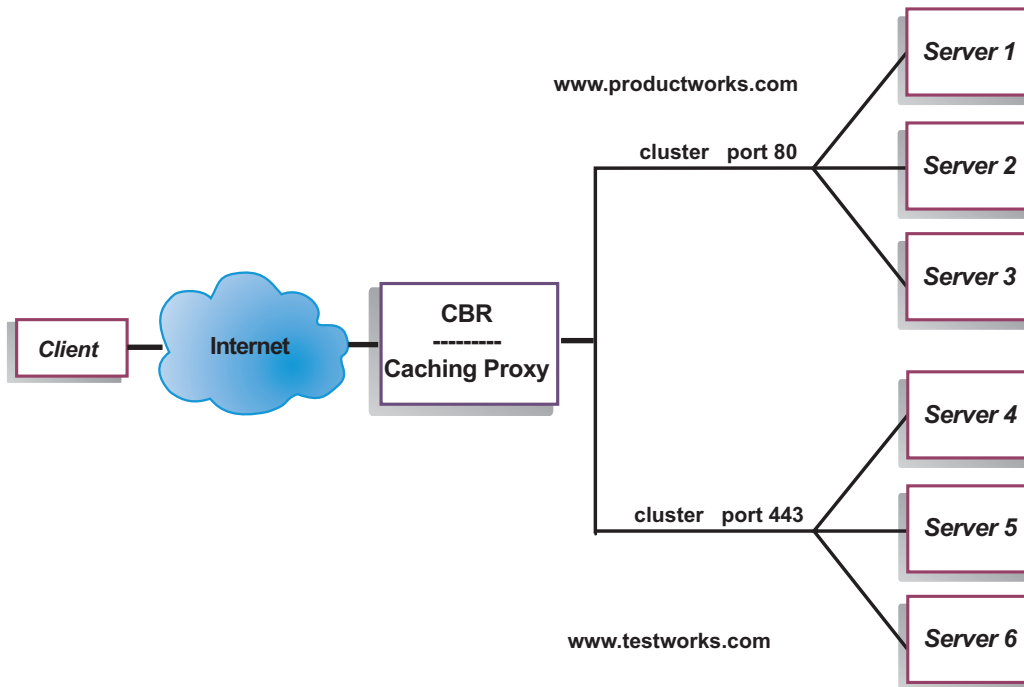


In this example for the CBR component, one cluster is defined at `www.productworks.com`. This cluster has two ports: port 80 for HTTP and port 443 for SSL. A client making a request to `http://www.productworks.com` (port 80) would go to a different server than a client requesting `https://www.productworks.com` (port 443).

Another way of configuring CBR would be appropriate if you have a very large site with many servers dedicated to each protocol supported. In this case, you might want to define a cluster for each protocol with a single port but with many servers, as shown in Figure 2.

2 clusters, each with 1 port

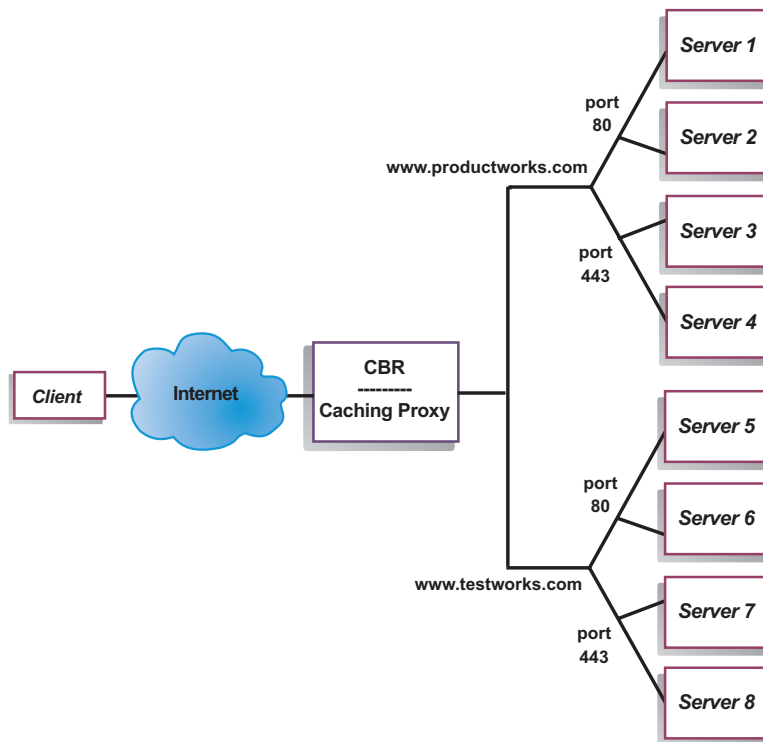
Another way of configuring Load Balancer might be appropriate if you have a very large site with many servers dedicated to each protocol supported. In this case, you might want to define a cluster for each protocol with a single port but with many servers.



In this example for the CBR component, two clusters are defined: **www.productworks.com** for port 80 (HTTP) and **www.testworks.com** for port 443 (SSL).

A third way of configuring CBR would be necessary if your site does content hosting for several companies or departments, each one coming into your site with a different URL. In this case, you might want to define a cluster for each company or department and then define any ports to which you want to receive connections at that URL, as shown in Figure 3.

2 clusters, each with 2 ports



In this example for the CBR component, two clusters are defined with port 80 (HTTP) and port 443 (SSL) for each of the sites at `www.productworks.com` and `www.testworks.com`.

Related tasks:

“Quick start configuration for Content Based Routing” on page 92

This quick start example shows how to configure three locally attached workstations using Content Based Routing (CBR) along with Caching Proxy, to load-balance web traffic between two web servers.

Setting up the Content Based Routing component

Set up the Content Based Routing (CBR) component in production environments and realistic test environments.

Before you begin

Before you set up the CBR component, you must be the root user (for AIX, HP-UX, Linux, or Solaris systems) or the Administrator (for Windows systems).

About this task

You need one IP address for each cluster of servers that is set up. A cluster address is an address that is associated with a host name (such as `www.company.com`). This IP address is used by a client to connect to the servers in a cluster. Specifically, this address is found in the URL request from the client. All requests that are made to the same cluster address are load balanced by CBR.

Before you use the CBR component, the system defaults for IPC (Inter-process Communication) must be modified. The maximum size of a shared memory

segment and the number of semaphore identifiers must be increased. To tune your system to support CBR, edit the `/etc/system` file on your system to add the following statements and then restart:

```
set shmsys:shminfo_shmmax=0x02000000
set semsys:seminfo_semmap=750
set semsys:seminfo_semmni=30
set semsys:seminfo_semmns=750
set semsys:seminfo_semmnu=30
set semsys:seminfo_semume=30
```

If you do not increase the shared memory segment to the values shown previously, **cbrcontrol executor start** command fails.

Procedure

1. Configure Caching Proxy to use CBR.

Important: To use CBR, you must install Caching Proxy. Caching Proxy is a service that starts automatically by default after installation. You must stop Caching Proxy before you start the CBR server function and modify the Caching Proxy service to start manually rather than automatically.

- a. Stop Caching Proxy by finding its process identifier by using the `ps -ef | grep ibmproxy` command and then ending the process by using the `kill process_id` command.
- b. Stop Caching Proxy from the Services panel.
- c. Modify the Caching Proxy configuration file (`ibmproxy.conf`), as follows:
 - Verify that the incoming `CacheByIncomingUrl` URL directive is off (default).
 - In the mapping rules section of the configuration file, add a mapping rule for every cluster that is similar to the following example:


```
Proxy /* http://cluster.domain.com/* cluster.domain.com
```

Tip: CBR sets the protocol, server, and target port later.

- d. Edit the four entries for the CBR plug-in, as follows:
 - `ServerInit`
 - `PostAuth`
 - `PostExit`
 - `ServerTerm`

Each entry must be on a single line. There are several instances of `ServerInit` in the `ibmproxy.conf` file, one for each plug-in. The entries for the CBR plug-in must be edited and uncommented.

Table 7. Additions to the configuration file, by operating system.

Operating system	Additions to CBR configuration file
AIX, Linux, and Solaris systems	<pre>ServerInit /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.so:ndServerInit PostAuth /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.so:ndPostAuth PostExit /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.so:ndPostExit ServerTerm /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.so:ndServerTerm</pre>
HP-UX	<pre>ServerInit /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.sl:ndServerInit PostAuth /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.sl:ndPostAuth PostExit /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.sl:ndPostExit ServerTerm /opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib/liblbcbcr.sl:ndServerTerm</pre>

Table 7. Additions to the configuration file, by operating system. (continued)

Operating system	Additions to CBR configuration file
Windows systems	<pre>ServerInit <install_root>\edge\ulb\cbr\servers\lib\liblbcbr.dll:ndServerInit PostAuth <install_root>\edge\ulb\cbr\servers\lib\liblbcbr.dll:ndPostAuth PostExit <install_root>\edge\ulb\cbr\servers\lib\liblbcbr.dll:ndPostExit ServerTerm <install_root>\edge\ulb\cbr\servers\lib\liblbcbr.dll:ndServerTerm</pre>

2. Start the server function by entering **cbrserver** on the command line. A default configuration file (`default.cfg`) automatically loads when you start **cbrserver**. If you decide to save the CBR configuration in the `default.cfg` file, then everything that is saved in this file automatically loads the next time **cbrserver** starts.
3. Start the executor function by entering **cbrcontrol executor start**. You might also change various executor settings now. Read the **dscontrol** executor topic to learn about the executor function.
4. Define a cluster and set cluster options. CBR balances requests that are sent for the cluster to the corresponding servers that are configured on the ports for that cluster.

The cluster is the symbolic name that is in the host portion of the URL and must match the name that is used in the Proxy statement of the `ibmproxy.conf` file.

Clusters that are defined in CBR must be defined to match the incoming request. A cluster must be defined by using the same host name or IP address that the incoming request contains. For example, if the request enters as the IP address, the cluster must be defined as the IP address. If there is more than one host name that resolves to a single IP address (and requests can arrive with any one of those host names) then all the host names must be defined as clusters.

Define a cluster by issuing the following command:

```
cbrcontrol cluster add cluster
```

Set cluster options by issuing the following command:

```
cbrcontrol cluster set cluster option value
```

5. Optional: Alias the network interface card (optional). If you are running Caching Proxy configured as a reverse proxy, when load balancing for multiple websites, you must add the cluster address for each website to at least one of the network interface cards of the Load Balancer server. Otherwise, you can omit this step.

Add the cluster address to the network interface by using the **ifconfig** command. Use the command for your operating system, as follows:

Table 8. . Commands to alias the network interface card

Operating system	Command
AIX	<code>ifconfig interface_name alias cluster_address netmask netmask</code>
HP-UX	<code>ifconfig interface_name cluster_address netmask netmask up</code>
Linux	<code>ifconfig interface_name cluster_address netmask netmask up</code>
Solaris 8, Solaris 9, and Solaris 10	<code>ifconfig interface_name addif cluster_address netmask netmask up</code>

Attention: For Linux and HP-UX systems, *interface_name* must have a unique number for each cluster address that is added, for example: eth0:1, eth0:2, and so on.

For Windows 2000, add the cluster address to the network interface, as follows:

- a. Click **Start > Settings > Control Panel**
- b. Double-click **Network and Dial-up Connections**
- c. Right-click **Local Area Connection**
- d. Click **Properties**.
- e. Select **Internet Protocol (TCP/IP)** and click **Properties**.
- f. Select **Use the following IP address** and click **Advanced**.
- g. Click **Add** and then type the IP address and subnet mask for the cluster.

For Windows 2003, add the cluster address to the network interface, as follows:

- a. Click **Start > Control Panel > Network Connections > Local Area Connection**
 - b. Click **Properties**.
 - c. Select **Internet Protocol (TCP/IP)** and click **Properties**.
 - d. Select **Use the following IP address** and click **Advanced**.
 - e. Click **Add** and then type the IP address and subnet mask for the cluster.
6. Define ports and set port options. The port number is the port on which the server applications are listening. For CBR with Caching Proxy running HTTP traffic, this is typically port 80.

Define a port to the cluster that you defined in the previous step by issuing the following command:

```
cbrcontrol port add cluster@port
```

Set port options by issuing the following command:

```
cbrcontrol port set cluster@port option value
```

Read the Commands reference for a complete list of the available to Load Balancer.

7. Define load balanced servers. The servers are the machines that run the applications that you want load balanced. The server is the symbolic name or dotted decimal address of the server machine. Define a server on the cluster and port by issuing the following command:

```
cbrcontrol server add cluster@port@server
```

You must define more than one server per port on a cluster to perform load balancing.

8. Add rules to your configuration. This step is the key step in configuring CBR with Caching Proxy. A rule defines how a URL request is distinguished and sent to one of the appropriate set of servers. The special rule type that is used by CBR is called a content rule. Define a content rule by issuing the following command:

```
cbrcontrol rule add cluster@port@rule type content pattern pattern
```

The value, *pattern*, is the regular expression that is compared to the URL in each client request.

Other rule types that are defined in Dispatcher can also be used in CBR.

9. Add servers to your rules. When a rule is matched by a client request, the rule's set of servers is queried for which server is best. The rule's server set is a subset of the servers that are defined in the port. Add servers to a rule's server set by issuing the following command:

```
cbrcontrol rule useserver cluster@port@rule server
```
10. Optional: Start the manager function. The manager function improves load balancing. Start the manager by issuing the following command:

```
cbrcontrol manager start
```
11. Optional: Start the advisor function. The advisors give the manager more information about the ability of the load balanced servers to respond to requests. An advisor is specific to a protocol. For example, start the HTTP advisor by issuing the following command:

```
cbrcontrol advisor start http port
```
12. Set cluster proportions, as required. If you start advisors, you might modify the proportion of importance that is given to advisor information that is included in the load balancing decisions. Set the cluster proportions by issuing the **cbrcontrol cluster set cluster proportions** command.
13. Start Caching Proxy.
 - a. Add to the environment variable.
 - Add to the LIBPATH environment variable:

```
/opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib
```
 - Add to the LD_LIBRARY_PATH environment variable:

```
/opt/IBM/WebSphere/Edge/ULB/cbr/servers/lib
```
 - Add to the PATH environment variable:

```
<install_root>ibm\edge\ulb\servers\lib
```
14. In the new environment, start Caching Proxy:
 - a. From the command prompt, issue **ibmproxy**
 - b. Start Caching Proxy from the Services panel: Click **Start > Settings (Windows 2000) > Control Panel > Administrative Tools > Services**.

Example

CBR configuration example

Configure CBR as follows:

1. Start CBR by issuing the **cbrserver** command.
2. Start the command line interface: issue the **cbrcontrol** command.
3. The **cbrcontrol** prompt displays. Issue the following commands (*cluster(c),port(p),rule(r),server(s)*):
 - `executor start`
 - `cluster add c`
 - `port add c@p`
 - `server add c@p@s`
 - `rule add c@p@r type content pattern uri=*`
 - `rule useserver c@p@r s`
4. Start Caching Proxy by issuing the **ibmproxy** command. (For Windows platforms, start Caching Proxy from the Services panel.)
5. Remove all proxy configurations from the browser.

6. Load `http://c/` into your browser, where *c* is the cluster that you configured previously.

- Server *s* is started
- The web page displays, as follows:

`http://s/`

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

“dscontrol executor” on page 233

Control the executor function with the `dscontrol executor` command.

Chapter 6. Administering Load Balancer

This section focuses on administering production environments and realistic test environments.

Enabling advisors to manage load balancing

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

Before you begin

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. They do this with a proactive client-like exchange with the servers. Advisors can be considered as lightweight clients of the application servers.

By writing your own advisors for the Load Balancer, you can customize how your server machines' load is determined.

For more information on how advisors work, read "Advisors" on page 113.

When using IPv6 protocols: If you are using IPv6 protocol on your machine and want to use advisors, you must modify the protocol file. To enable IPv6, insert the following line in the protocol file:

```
ipv6-icmp 58 IPv6-ICMP # IPv6 interface control message protocol
```

The protocol file is in the following directory:

- /etc/protocols
- C:\windows\system32\drivers\etc\

The product provides several protocol-specific advisors for the most popular protocols. However, it does not make sense to use all of the provided advisors with Load Balancer. Load Balancer also supports the concept of a "custom advisor" that allows users to write their own advisors.

Limitation on using advisors with bind-specific server applications:

- In order to use advisors on bind specific servers, start two instances of the server: One instance to bind on the *cluster@port* and the other instance to bind on the *server@port*. To determine if the server is bind specific, issue the `netstat -an` command and look for the *server@port*. If the server is not bind specific, the result from this command will be 0.0.0.0:80. If the server is bind specific, you will see an address such as 192.168.15.103:80.
- If using `arp publish` instead of the `ifconfig alias` command, Load Balancer will support the use of advisors when load-balancing servers with bind-specific server applications when they are binding to the cluster IP address.

About this task

You can start an advisor for a particular port across all clusters (group advisor). Or, you can choose to run different advisors on the same port, but on different clusters (cluster specific advisor).

Note: If Load Balancer is running on a computer with multiple network adapter cards, you cannot force the source IP address of the packet to a specific address when you want the advisor traffic to flow over a particular adapter.

Procedure

1. Start the advisor of your choice. For a list of possible advisors, refer to the list of advisors, or create a custom advisor.
 - **Cluster specific advisor:** To start an advisor on port 80 for clusterA, for example, specify both the cluster and port:

```
dscontrol advisor start ADV_name clusterA@80
```

This command will start an advisor on port 80 for clusterA. This advisor will advise on all servers attached to port 80 for clusterA.
 - **Group advisor:** To start an advisor on port 80 for all other clusters, simply specify the port:

```
dscontrol advisor start ADV_name 80
```

This command will start the advisor on port 80 for all clusters and sites that do not currently have a cluster or site specific advisor. Your advisor will advise on all servers attached to port 80.
 - a. Optional: If you are starting the HTTP or HTTPS advisor, you might want to define a unique client URL string to allow the advisor to monitor individual services in the server. For more information on this option, refer to “Getting service-specific advice with the advisor request or response option” on page 118.
2. Optional: Set the advisor interval. The advisor interval sets how often an advisor asks for status from the servers on the port it is monitoring and then reports the results to the manager. If the advisor interval is too low, it can mean poor performance as a result of the advisor constantly interrupting the servers. If the advisor interval is too high, it can mean that the manager’s decisions about weighting will not be based on accurate, up-to-date information.

Note: The advisor defaults should work efficiently for the great majority of possible scenarios. Be careful when entering values other than the defaults. For example, to set the interval to 3 seconds for the HTTP advisor for port 80, enter the following command:

```
dscontrol advisor interval http 80 3
```

It does not make sense to specify an advisor interval that is smaller than the manager interval. The default advisor interval is 7 seconds.

3. Optional: Set the advisor report timeout. To make sure that out-of-date information is not used by the manager in its load-balancing decisions, the manager will not use information from the advisor whose time stamp is older than the time set in the advisor report timeout. The advisor report timeout should be larger than the advisor polling interval. If the timeout is smaller, the manager will ignore reports that logically should be used. By default, advisor reports do not timeout — the default value is unlimited.

For example, to set the advisor report timeout to 30 seconds for the HTTP advisor for port 80, enter the following command:

```
dscontrol advisor timeout http 80 30
```

For more information on setting the advisor report timeout, see “dscontrol advisor” on page 227.

4. Optional: Set the advisors connect and receive timeout values. For Load Balancer, you can set the advisor’s timeout values at which it detects a particular port on the server (a service) is failed. The failed-server timeout values (connecttimeout and receivetimeout) determine how long an advisor waits before reporting that either a connect or receive has failed.

To obtain the fastest failed-server detection, set the advisor connect and receive timeouts to the smallest value (one second), and set the advisor and manager interval time to the smallest value (one second).

Note: If your environment experiences a moderate to high volume of traffic such that server response time increases, be careful not to set the connecttimeout and receivetimeout values too small, or the advisor may prematurely mark a busy server as failed.

For example, to set the connecttimeout and receivetimeout to 9 seconds for the HTTP advisor on port 80, type the following command:

```
dscontrol advisor connecttimeout http 80 9
```

```
dscontrol advisor receivetimeout http 80 9
```

The default for connect and receive timeout is 3 times the value specified for the advisor interval time.

5. Optional: Set the advisor retry value. Advisors have the ability to retry a connection before marking a server down. The advisor will not mark a server down until the server query has failed the number of retries plus 1. The retry value should be no larger than 3.

The following command sets a retry value of 2 for the LDAP advisor on port 389:

```
dscontrol advisor retry ldap 389 2
```

Related tasks:

“Getting advice with Metric Server” on page 121

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

Advisors

Advisors periodically open a TCP connection with each server and send a request message to the server. The content of the message is specific to the protocol running on the server. For example, the HTTP advisor sends an HTTP “HEAD” request to the server.

Advisors then listen for a response from the server. After getting the response, the advisor makes an assessment of the server. To calculate this “load” value, most advisors measure the time for the server to respond, and then use this value (in milliseconds) as the load.

Advisors then report the load value to the manager function, where it appears in the manager report in the “Port” column. The manager then calculates aggregate weight values from all its sources, per its proportions, and sets these weight values into the executor function. The Executor will then use these weights for load balancing new incoming client connections.

If the advisor determines that a server is alive and functioning properly, it will report a positive, non-zero load number to the Manager. If the advisor determines that a server is not active, it will return a special load value of negative one (-1). The Manager and the Executor will not forward any further connections to that server until that server has come back up.

Note: Before sending the initial request message, the advisor will ping the server. This is intended to provide quick status to determine if the machine is online. After the server responds to the ping, no more pings are sent. To disable the pings, add `-DLB_ADV_NB_PING` to the Load Balancer start script file.

Related tasks:

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

List of advisors

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

Table 9. List of advisors

Advisor Name	Description
caching proxy	The caching proxy advisor opens a connection, sends a caching proxy specific HTTP GET request, and interprets the response as a caching proxy load. Important: When using the Caching Proxy advisor, Caching Proxy needs to be running on all servers being load balanced. The machine on which the Load Balancer resides does not need to have Caching Proxy installed unless it is collocated on the same machine it is load balancing.
connect	The connect advisor does not exchange any protocol-specific data with the server. It simply measures the time it takes to open and close a TCP connection with the server. This advisor is useful for server applications which use TCP, but with a higher-level protocol for which an IBM-supplied or custom advisor is not available.
Custom advisors	Dispatcher provides the ability for a customer to write a custom (customizable) advisor. This enables support for proprietary protocols (on top of TCP) for which IBM has not developed a specific advisor. For more information, see “Creating a custom advisor” on page 125.

Table 9. List of advisors (continued)

Advisor Name	Description
db2	The DB2 advisor works in conjunction with the DB2 servers. Dispatcher has the built in capability of checking the health of DB2 servers without the need for customers to write their own custom advisors. The DB2 advisor communicates with the DB2 connection port only, not the Java connection port.
dns	The DNS advisor opens a connection, sends a pointer query for DNS, waits for a response, closes the connection and returns the elapsed time as a load.
ftp	The FTP advisor opens a connection, sends a SYST request, waits for a response, closes the connection, and returns the elapsed time as a load.
http	The HTTP advisor opens a connection, sends a HEAD request by default, waits for a response connection, and returns the elapsed time as a load. See “Getting service-specific advice with the advisor request or response option” on page 118 for more information on how to change the type of request sent by the HTTP advisor.
https	The HTTPS advisor is a Δ heavyweight Δ advisor for SSL connections. It performs a full SSL socket connection with the server. The HTTPS advisor opens an SSL connection, sends an HTTPS request, waits for a response, closes the connection, and returns the elapsed time as a load. (See also the SSL advisor, which is a Δ lightweight Δ advisor for SSL connections.) Note: The HTTPS advisor has no dependency upon server key or certificate content, but they must not be expired.
imap	The IMAP advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.
ldap	The LDAP advisor opens a connection, sends an anonymous BIND request, waits for a response, closes the connection, and returns the elapsed time as a load.

Table 9. List of advisors (continued)

Advisor Name	Description
ldapuri	<p>Note: The LDAP URI advisor allows you better gauge LDAP availability by processing a complete request to the LDAP server.</p> <p>The advisor:</p> <ol style="list-style-type: none"> 1. Opens a connection. 2. Sends a BIND request, which is based on the advisorrequest field that you define on the server object. 3. Waits for a response. 4. Closes the connection. 5. Returns the elapsed time as a load. <p>Read “Configuring the LDAP URI advisor” on page 119 for more information on configuring this advisor.</p>
nntp	<p>The NNTP advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.</p>
ping	<p>The ping advisor does not open a TCP connection with the servers, but instead reports whether the server responds to a ping. While the ping advisor may be used on any port, it is also designed for configurations using the wildcard port, over which multiple protocol traffic may be flowing. It is also useful for configurations using non-TCP protocols with their servers.</p>
pop3	<p>The POP3 advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.</p>
reach	<p>The reach advisor pings its target machines. This advisor is also designed for the Dispatcher's high availability components to determine reachability of its reach targets. Its results flow to high availability component and do not appear in the manager report. Unlike the other advisors, the reach advisor starts automatically by the manager function of the Dispatcher component.</p>
sip	<p>The SIP advisor opens a connection, sends an OPTIONS request, waits for a response, closes the connection, and returns the elapsed time as a load. The SIP advisor that is supported runs on TCP only and requires an application to be installed on a server that responds to an OPTIONS request.</p>

Table 9. List of advisors (continued)

Advisor Name	Description
smtp	The SMTP advisor opens a connection, waits for an initial message from the server, sends a quit, closes the connection, and returns the elapsed time as a load.
ssl	<p>The SSL advisor is a <code>lightweight</code> advisor for SSL connections. It does not establish a full SSL socket connection with the server. The SSL advisor opens a connection, sends an SSL CLIENT_HELLO request, waits for a response, closes the connection, and returns the elapsed time as a load. (See also the HTTPS advisor, which is a <code>heavyweight</code> advisor for SSL connections.)</p> <p>Note: Be aware of the following:</p> <ul style="list-style-type: none"> • The SSL advisor has no dependency upon key management or certificates. • You might want to use the TLS advisor if you find that the SSL advisor is marking servers down, you know that the server is still up, and you see messages stating "No ciphers found."
ssl2http	The ssl2http advisor starts and advises on the servers listed under port 443, but the advisor will open a socket to the <code>mapport</code> for HTTP requests. Only use the ssl2http advisor for CBR if the client-to-proxy protocol list is SSL and the proxy-to-server protocol is HTTP.
telnet	The Telnet advisor opens a connection, waits for an initial message from the server, closes the connection, and returns the elapsed time as a load.
tls	<p>Transport Layer Security (TLS) is the successor to the SSL protocol. TLS 1.0 was the follow up to SSLv3.</p> <p>You might want to use the TLS advisor if you find that the SSL advisor is marking servers down, you know that the server is still up, and you see messages stating "No ciphers found."</p>
was	The WAS (WebSphere Application Server) advisor works in conjunction with the WebSphere Application servers. Customizable sample files for this advisor are provided in the installation directory. For more information, see "Example: Implementing the WAS advisor" on page 142.

Table 9. List of advisors (continued)

Advisor Name	Description
wlm	The WLM (Workload Manager) advisor is designed to work in conjunction with servers on OS/390 mainframes running the MVS™ Workload Manager (WLM) component. For more information, see “The Workload Management Advisor” on page 124.

Related tasks:

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

“Configuring the LDAP URI advisor” on page 119

The LDAP URI advisor allows you better gauge Lightweight Directory Access Protocol (LDAP) availability by processing a complete request to the LDAP server. The LDAP URI advisor opens a connection to the LDAP server and sends a BIND request that is based on the `advisorrequest` field that you define on the server object. The advisor then waits for a response from the LDAP server and returns the elapsed time as a load.

Getting service-specific advice with the advisor request or response option

After you have started an HTTP or HTTPS advisor, you can define a unique client HTTP URL string, specific for the service that you want to query on the server. This allows the advisor to assess the health of the individual services within a server.

About this task

For each defined logical server under the HTTP port you can specify a unique client HTTP URL string, specific for the service that you want to query on the server. The HTTP or HTTPS advisor uses the `advisorrequest` string to query the health of the servers. The default value is "HEAD / HTTP/1.0."

The `advisorresponse` string is the response that the advisor scans for in the HTTP response. The advisor uses the `advisorresponse` string to compare to the real response that is received from the server. The default value is null.

Procedure

- Issue the server set command with the `advisorrequest` and `advisorresponse` parameters.
 - When issuing the command from the `dscontrol>>` shell prompt, you must place quotes around the string if a blank is contained within the string. For example:


```
server set cluster@port@server advisorrequest "head / http/1.0"
server set cluster@port@server advisorresponse "HTTP 200 OK"
```
 - When issuing the `dscontrol` command from the operating system prompt, you must precede the text with `△\△` and follow the text with `\`". For example:

```
dscontrol server set cluster@port@server advisorrequest "\"head / http/1.0\""
```

```
dscontrol server set cluster@port@server advisorresponse "\"HTTP 200 OK\""
```

When you create the request that the HTTP or HTTPS advisor sends to back-end servers to see if they are functioning, you type the start of the HTTP request and Load Balancer completes the end of the request with the following:

```
\r\nAccept:  
*/*\r\nUser-Agent:IBM_Load_Balancer_HTTP_Advisor\r\n\r\n
```

- Optional: If you want to add other HTTP header fields before Load Balancer appends this string to the end of the request, you can do so by including your own `\r\n` string in the request. The following is an example of what you might type to add the HTTP host header field to your request:

```
GET /pub/WWW/TheProject.html HTTP/1.0 \r\nHost: www.w3.org
```

Note: After starting an HTTP or HTTPS advisor for a specified HTTP port number, the advisor request and response value is enabled for servers under that HTTP port. See “dscontrol server” on page 249 for more information.

Related tasks:

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

Configuring the LDAP URI advisor

The LDAP URI advisor allows you better gauge Lightweight Directory Access Protocol (LDAP) availability by processing a complete request to the LDAP server. The LDAP URI advisor opens a connection to the LDAP server and sends a BIND request that is based on the advisorrequest field that you define on the server object. The advisor then waits for a response from the LDAP server and returns the elapsed time as a load.

About this task

In situations in which you cannot perform an anonymous bind request to an LDAP server, you can use the LDAP URI advisor to bind with an LDAP server that requires a user name and password. The LDAP URI advisor might provide a more precise measurement of workload, since the LDAP server will be required to process a complete request rather than perform only an anonymous bind.

Procedure

1. Set the advisorrequest field for the server that will use the LDAP URI advisor.
 - a. Set the advisorrequest field on the server object with the dscontrol server command. The string you use for the advisorrequest field must be compliant with the RFC2255 - The LDAP URL Format, but you need to format the string differently to work with Load Balancer. Specifically, RFC2255 does not allow for a password to be provided on the URL line, so Load Balancer has to take a modified string in the advisorrequest field, and convert that string to an RFC2255 compliant request. For example, a typical LDAP request might look like this:

```
ldapsearch -x -H ldap://bluepages.ibm.com:389 -D  
"uid=973524897,c=us,ou=bluepages,o=ibm.com" -w MYPASS -b "ou=bluepages,o=ibm.com" mail=user@us.ibm.com
```

To use this in Load Balancer's `advisorrequest` field, you should modify the string to look like this:

```
ldap://bluepages.ibm.com:389/ou=bluepages,o=ibm.com??sub?(mail=user@us.ibm.com)?bindname=uid=973524897,c=us,ou=bluepages,o=ibm.com?!bindpass=MYPASS
```

To use this `advisorrequest` string with the `dscontrol` server command, your command would look like this:

```
dscontrol server set cluster@server@port advisorrequest "ldap://bluepages.ibm.com:389/ou=bluepages,o=ibm.com??sub?(mail=user@us.ibm.com)?bindname=uid=973524897,c=us,
```

Note: Be aware of the following:

- Note the `?!bindpass=MYPASS` extension that is used in the examples. Replace `MYPASS` with the password that is used to authenticate the LDAP request.
 - Certain characters must be escaped using the `%` method that is described in RFC 1738. You must escape any of the following characters:
 - URL-illegal characters, like spaces.
 - URL special characters, as defined in section 2.2 of RFC 1738.
 - The reserved character `'?`, which is the ASCII 63 character, when it appears inside a `dn`, `filter`, or other element of an LDAP URL.
 - A comma character `,` if it occurs inside an extension value.
 - b. Optional: Set the `advisorresponse` field on the server object. If you set this field, you must set the value to a substring that is expected to be present in the response from the LDAP server.
2. Start the LDAP URI advisor. To start the LDAP URI advisor, use the `dscontrol` advisor command:

```
dscontrol advisor start ldapuri cluster@port
```

Note: Verify that you are using the LDAP URI advisor, and not the LDAP advisor. The LDAP advisor only supports anonymous bind requests to LDAP servers.

Related concepts:

“Advisors” on page 113

Related reference:

“Getting service-specific advice with the advisor request or response option” on page 118

After you have started an HTTP or HTTPS advisor, you can define a unique client HTTP URL string, specific for the service that you want to query on the server. This allows the advisor to assess the health of the individual services within a server.

“List of advisors” on page 114

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

“dscontrol advisor” on page 227

Use this command to control various features of the advisor function.

Related information:

“dscontrol server” on page 249

Configure servers and modify existing server configurations with the dscontrol server command.

Getting advice with Metric Server

Metric Server provides server load information to the Load Balancer in the form of system-specific metrics, reporting on the health of the servers.

Before you begin

The Metric Server agent must be installed and running on all servers that are being load balanced.

If you are using IPv6 protocol on your machine and want to use Metric Server, you must check to see if protocol 58 is defined to be ICMPv6 in the **protocol** file.

When using IPv6 protocols: If you are using IPv6 protocol on your machine and want to use advisors, you must modify the protocol file. To enable IPv6, insert the following line in the protocol file:

```
ipv6-icmp 58 IPv6-ICMP # IPv6 interface control message protocol
```

The protocol file is in the following directory:

- /etc/protocols
- C:\windows\system32\drivers\etc\

Metric Server Restriction: Like the Metric Server agent, the WLM agent reports on server systems as a whole, rather than on individual protocol-specific server daemons. Metric Server and WLM place their results into the system column of the manager report. As a consequence, running both the WLM advisor and Metric Server at the same time is not supported.

About this task

The Load Balancer manager queries the Metric Server agent residing on each of the servers, assigning weights to the load balancing process using the metrics gathered from the agents. The results are also placed into the manager report.

Note: When two or more metrics are gathered and normalized for each server into a single system load value, rounding errors may occur.

Procedure

1. Configure Metric Server on the Load Balancer machine.

Tip: You can use similar steps for configuring Metric Server for the other components of Load Balancer.

- a. Start **dsserver**. Start the executor, and add clusters, ports and servers to your configuration.
- b. Start the manager. Issue the command:

```
dscontrol manager start manager.log port
```

where *port* is the RMI port chosen for all the Metric Server agents to run on. The default RMI port that is set in the `metricserver.cmd` file is 10004.

- c. Add the system metric script to the cluster. Issue the command:

```
dscontrol metric add cluster@systemMetric
```

systemMetric is the name of the script (residing on the back-end server) which should run on each of the servers in the configuration under the specified cluster. Two scripts are provided for the customer - `cpuload` and `memload` - or you can create custom system metric scripts.

Tip: For Site Selector, `cpuload` and `memload` run automatically.

The script contains a command which should return a numeric value in the range of 0-100 or a value of -1 if the server is down. This numeric value should represent a load measurement, not an availability value.

Note: If the name of your System Metric script has an extension other than `.exe`, you must specify the full name of the file (for example, `mssystemscript.bat`). This is due to a Java limitation.

- d. Add to the configuration only servers that contain a Metric Server agent running on the port specified in the `metricserver.cmd` file. The port should match the port value specified in the manager start command.

Note: To ensure security:

- On the Load Balancer machine, create a key file (using the `lbkeys create` command).
- On the back-end server machine, copy the resulting key file, for the component you are using, to the `install_root/admin/keys` directory. Verify that the key file's permissions enable the file to be readable by the root.

2. Configure Metric Server on the server machines.

- a. Install the Metric Server package from the Load Balancer installation files.
- b. Check the metric server script in the `install_root/ms/bin` directory to verify that the desired RMI port is being used. The default RMI port is 10004.

Note: The RMI port value specified must be the same value as the RMI port value that was specified in the manager start command in Step 1b.

- c. Optional: You can write their own customized metric script files which define the command that the Metric Server will issue on the server machines. Ensure that any custom scripts are executable and located in the `install_root/ms/script` directory. Custom scripts must return a numeric load value in the range of 0-100.

Note: A custom metric script must be a valid program or script with a .bat or .cmd extension.

Specifically, for Linux and other UNIX-based systems, scripts must begin with the shell declaration, otherwise they may not properly run.

The following two scripts are provided for the customer in the *install_root*/ms/script directory:

- **cpuload:** returns the percentage of cpu in use ranging from 0-100
- **memload:** returns the percentage of memory in use ranging from 0-100.

- d. Start the metric server agent. On a command line of each server machine where Metric Server resides, type

```
metricserver start
```

Click **Start > Control Panel > Administrative Tools > Services**. Right-click **IBM Metric Server (ULB)** and select **Start**.

- e. Optional: Stop the metric server agent.

To stop the Metric Server agent, issue this command on every server machine where Metric Server resides:

```
metricserver stop
```

Click **Start > Control Panel > Administrative Tools > Services**. Right-click **IBM Metric Server (ULB)** and select **Stop**.

3. Optional: Change the log level in the Metric Server startup script. You can specify a log level range of 0 through 5, similar to the log level range in Load Balancer logs. This will generate an agent log in the *install_root*/ms/logs directory.
4. Optional: To have Metric Server run on an address other than the local host, you need to edit the metricserver file on the load balanced server machine.

Note: When gathering metrics across different domains, you must explicitly set the `java.rmi.server.hostname` in the server script (`dsserver`, etc) to the fully qualified domain name (FQDN) of the machine that is requesting the metrics. This is necessary because `InetAddress.getLocalHost.getHostName()` might not return the FQDN.

- a. After the occurrence of Java in the metricserver file, insert the following:

```
-Djava.rmi.server.hostname=OTHER_ADDRESS
```

- b. Before the `if` statements in the metricserver file, add the following line:

```
hostname OTHER_ADDRESS
```

- c. You will also need to alias the OTHER_ADDRESS on the Microsoft stack of the Metric Server machine. For example:

```
call netsh interface ip add address "Local Area Connection"  
addr=9.37.51.28 mask=255.255.240.0
```

5. Optional: Configure Metric Server for IPv4 only or IPv6 only. In a Load Balancer configuration that supports both IPv4 and IPv6 clusters, servers that run the Metric Server function can be configured as an IPv4 server only or as an IPv6 server only, but not both. To force Metric Server to use a particular IP protocol, specify the Java property `java.rmi.server.hostname` in the metricserver script.

Note: The host name specified in the Java property must be the physical IP address of the Metric Server.

- For Metric Server to communicate over the IPv6 address `2002:92a:8f7a:162:9:42:92:67`, specify the Java property after `$LB_CLASSPATH` in the metricserver startup script, in the *install_root*/bin directory, as follows:


```

install_root/java/jre/bin/java ..... $ULB_CLASSPATH
-Djava.rmi.server.hostname=2002:92a:8f7a:162:9:42:92:67
com.ibm.internet.nd.sma.SMA_Agent $LB_RMIPORT $LOG_LEVEL $LOG_SIZE $LOG_DIRECTORY $KEYS_DIRECTORY
$SCRIPT_DIRECTORY &

```

- For Metric Server to communicate over the IPv6 address 2002:92a:8f7a:162:9:42:92:67, you must edit the metricserver.cmd file, in the *install_root/bin* directory, as follows:

```

start
/min /wait %IBMLBPATH%\java\jre\bin\java
-Djava.rmi.server.hostname=2002:92a:8f7a:162:9:42:92:67
-Djava.net.preferIPv4Stack=false
-Djava.net.preferIPv6Stack=true -Xrs -cp
%LB_CLASSPATH% com.ibm.internet.nd.sma.SMA_Agent
%RMI_PORT% %LOG_LEVEL% %LOG_SIZE% %LOG_DIRECTORY% %KEYS_DIRECTORY%
%SCRIPT_DIRECTORY%
goto done

:stop
%IBMLBPATH%\java\jre\bin\java
-Djava.rmi.server.hostname=2002:92a:8f7a:162:9:42:92:67
-Djava.net.preferIPv4Stack=false
-Djava.net.preferIPv6Stack=true -cp %LB_CLASSPATH% com.ibm.internet.nd.sma.SMA_AgentStop %RMI_PORT%
:done

```

Related tasks:

Tuning

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

The Workload Management Advisor

WLM is code that runs on MVS mainframes. It can be queried to ask about the load on the MVS machine. When MVS Workload Management has been configured on your OS/390 system, Dispatcher can accept capacity information from WLM and use it in the load balancing process.

Using the WLM advisor, Dispatcher will periodically open connections through the WLM port on each server in the Dispatcher host table and accept the capacity integers returned. Because these integers represent the amount of capacity that is still available and Dispatcher expects values representing the loads on each machine, the capacity integers are inverted by the advisor and normalized into load values (that is, a large capacity integer but a small load value both represent a healthier server). The resulting loads are placed into the System column of the manager report.

There are several important differences between the WLM advisor and other Dispatcher advisors:

- Other advisors open connections to the servers using the same port on which flows normal client traffic. The WLM advisor opens connections to the servers using a port different from normal traffic. The WLM agent on each server machine must be configured to listen on the same port on which the Dispatcher WLM Advisor is started. The default WLM port is 10007.

- Other advisors only assess those servers defined in the Dispatcher `cluster@port@server` configuration for which the server's port matches the advisor's port. The WLM advisor advises upon every server in the Dispatcher configuration (regardless of the `cluster@port`). Therefore you must not define any non-WLM servers when using the WLM advisor.
- Other advisors place their load information into the manager report under its "Port" column. The WLM advisor places its load information into the manager report under its system column.
- It is possible to use both protocol-specific advisors along with the WLM advisor. The protocol-specific advisors will poll the servers on their normal traffic ports, and the WLM advisor will poll the system load using the WLM port.

Metric Server Restriction: Like the Metric Server agent, the WLM agent reports on server systems as a whole, rather than on individual protocol-specific server daemons. Metric Server and WLM place their results into the system column of the manager report. As a consequence, running both the WLM advisor and Metric Server at the same time is not supported.

Related tasks:

"Enabling advisors to manage load balancing" on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

Creating a custom advisor

A custom advisor is a small piece of Java code, provided as a class file, that is called by the Load Balancer base code to determine the load on a server. The base code provides all necessary administrative services, including starting and stopping an instance of the custom advisor, providing status and reports, recording history information in a log file, and reporting advisor results to the manager component.

About this task

Custom advisors are called after native, or standard, advisors have been searched. If the Load Balancer does not find a specified advisor among the list of standard advisors, it consults the list of custom advisors. When the Load Balancer base code calls a custom advisor, the following steps happen:

1. The Load Balancer base code opens a connection with the server machine.
2. If the socket opens, the base code calls the specified advisor's `GetLoad` function.
3. The advisor's `GetLoad` function performs the steps that the user has defined for evaluating the server's status, including waiting for a response from the server. The function terminates execution when the response is received.
4. The Load Balancer base code closes the socket with the server and reports the load information to the manager. Depending on whether the custom advisor operates in normal mode or in replace mode, the base code sometimes does additional calculations after the `GetLoad` function terminates.

Custom advisors can be designed to interact with the Load Balancer in either normal mode or replace mode. The choice for the mode of operation is specified in

the custom advisor file as a parameter in the constructor method. (Each advisor operates in only one of these modes, based on its design.)

- **Normal mode:** the custom advisor exchanges data with the server, and the base advisor code times the exchange and calculates the load value. The base code then reports this load value to the manager. The custom advisor returns the value zero to indicate success, or negative one to indicate an error.

To specify normal mode, set the replace flag in the constructor to false.

- **Replace mode:** the base code does not perform any timing measurements. The custom advisor code performs whatever operations are specified, based on its unique requirements, and then returns an actual load number. The base code accepts the load number and reports it, unaltered, to the manager. For best results, normalize your load numbers between 10 and 1000, with 10 representing a fast server and 1000 representing a slow server.

To specify replace mode, set the replace flag in the constructor to true.

Like all advisors, a custom advisor extends the functionality of the advisor base class, which is called `ADV_Base`. The advisor base performs most of the advisor's functions, such as reporting loads back to the manager for use in the manager's weight algorithm. The advisor base also performs socket connect and close operations and provides send and receive methods for use by the advisor. The advisor is used only for sending and receiving data on the specified port for the server that is being investigated. The TCP methods provided within the advisor base are timed to calculate load. A flag within the constructor of the advisor base overwrites the existing load with the new load returned from the advisor, if desired.

Note: Based on a value set in the constructor, the advisor base supplies the load to the weight algorithm at specified intervals. If the advisor has not completed processing and cannot return a valid load, the advisor base uses the previously reported load.

Procedure

1. Name your advisor. Custom advisor file names must follow the form `ADV_name.java`, where *name* is the name that you choose for your advisor.

Note:

- You must use the `ADV_` prefix for the advisor name.
 - You must name the custom advisor using lower-case alphabetic characters to eliminate case sensitivity when an operator types commands on a command line.
 - The custom advisor class must be located within the `install_root/lib/CustomAdvisors` subdirectory.
 - According to Java conventions, the name of the class defined within the file must match the name of the file.
2. Write your custom advisor. Read "Custom advisor methods and function calls" on page 127 for a list of methods and function calls to use in your advisor. Be aware that custom advisors need to have all the required routines. Advisors must have the following base class methods:
 - A constructor routine. The constructor calls the base class constructor.
 - An `ADV_AdvisorInitialize` method. This method provides a way to perform additional steps after the base class completes its initialization.

- A getLoad routine. The base advisor class performs the socket opening; the getLoad function only needs to issue the appropriate send and receive requests to complete the advising cycle.
3. Compile the advisor.
 - You must write custom advisors in the Java language and compile them with a Java compiler that is at the same level as the Load Balancer code. To check the version of Java on your system, run the following command from the `install_root/java/bin` directory:


```
java -fullversion
```

If the current directory is not part of your path, you will need to specify that Java should be run from the current directory to ensure you are getting the correct version information. In this case, run the following command from the `install_root/java/bin` directory:

```
./java -fullversion
```
 - The following files are referenced during compilation:
 - The custom advisor file.
 - The base classes file, `ibmnd.jar`, which is found in the `install_root/servers/lib` directory.
 - Your classpath environment variable must point to both the custom advisor file and the base classes file during the compilation. A compile command might have the following format, if your advisor is in the current directory:


```
install_path/java/bin/javac -classpath install_root/servers/lib/ibmnb.jar ADV_name.java
```
 - The output of the compilation is a class file, for example, `ADV_name.class`. Before starting the advisor, copy the class file to the `install_root/servers/lib/CustomAdvisors/` directory.

Note: You can compile custom advisors on one operating system and run on another operating system. For example, you can compile your advisor on a Windows system, copy the resulting class file, in binary format, to a Linux machine, and run the custom advisor there. For AIX, HP-UX, Linux, and Solaris operating systems, the syntax is similar.

4. Run your custom advisor. Custom advisors are called after native, or standard, advisors are searched. If Load Balancer does not find a specified advisor among the list of standard advisors, it consults the list of custom advisors.
 - a. If you have not already done so, copy the advisor's class file to the CustomAdvisors subdirectory on the Load Balancer machine. For example, for a custom advisor named `myping`, the file path is `install_root/servers/lib/CustomAdvisors/ADV_myping.class`.
 - b. Configure the Load Balancer, start its manager function, and issue the command to start your custom advisor. The custom advisor is specified by its name, excluding the `ADV_` prefix and the file extension:


```
dscontrol advisor start name.ext port
```

The port number specified in the command is the port on which the advisor will open a connection with the target server.

Related reference:

“Example: Sample advisor” on page 132
 This is a sample advisor file called `ADV_sample`.

Custom advisor methods and function calls

Use the following advisor methods and function calls in your custom advisors.

Be aware that custom advisors need to have all the required routines. Advisors must have the following base class methods:

- A constructor routine. The constructor calls the base class constructor.
- An `ADV_AdvisorInitialize` method. This method provides a way to perform additional steps after the base class completes its initialization.
- A `getLoad` routine. The base advisor class performs the socket opening; the `getLoad` function only needs to issue the appropriate send and receive requests to complete the advising cycle.

Constructor (provided by advisor base)

```
public <advisor_name> {  
    String sName;  
    String sVersion;  
    int iDefaultPort;  
    int iInterval;  
    String sDefaultLogFileName;  
    boolean replace  
}
```

sName

The name of the custom advisor

sVersion

The version of the custom advisor.

iDefaultPort

The port number on which to contact the server if no port number is specified in the call.

iInterval

The interval at which the advisor will query the servers.

sDefaultLogFileName

This parameter is required but not used. The only acceptable value is a null string, ""

replace

Whether or not this advisor functions in replace mode. Possible values are the following:

- *true* – Replace the load calculated by the advisor base code with the value reported by the custom advisor.
- *false* – Add the load value reported by the custom advisor to the load value calculated by the advisor base code.

ADV_AdvisorInitialize() method

```
void ADV_AdvisorInitialize()
```

This method is provided to perform any initialization that might be required for the custom advisor. This method is called after the advisor base module starts. In many cases, including the standard advisors, this method is not used and its code consists of a return statement only. This method can be used to call the `suppressBaseOpeningSocket()` on page 131 method, which is valid only from within this method.

In many cases, including the standard advisors, this method is not used and its code consists of a return statement only. You can use this method to call the `suppressBaseOpeningSocket` method, which is valid only from within the `ADV_AdvisorInitialize` method.

ADVLOG() method

The ADVLOG function allows a custom advisor to write a text message to the advisor base log file. The format follows:

```
void ADVLOG (int logLevel, String message)
```

This command has the following parameters:

logLevel

The status level at which the message is written to the log file. The advisor log file is organized in stages; the most urgent messages are given status level 0 and less urgent messages receive higher numbers. The most verbose type of message is given status level 5. These levels are used to control the types of messages that the user receives in real time (The dscontrol command is used to set verbosity). Catastrophic errors should always be logged at level 0.

message

The message to write to the log file. The value for this parameter is a standard Java string.

getAdvisorName function

The getAdvisorName function returns a Java string with the suffix portion of your custom advisor name. For example, for an advisor named ADV_cdload.java, this function returns the value cdload.

This function does not take parameters.

Note: It is not possible for this value to change during one instantiation of an advisor.

caller.getCurrentServerId()

The getCurrentServerId function returns a Java string which is a unique representation for the current server. Typically, this value changes each time you call your custom advisor, because the advisor base code queries all server machines in series.

This function takes no parameters.

caller.getCurrentClusterId()

The getCurrentClusterId function call returns a Java string which is a unique representation for the current cluster. Typically, this value changes each time you call your custom advisor, because the advisor base queries all clusters in series.

This function takes no parameters.

caller.getSocket()

The getSocket function call returns a Java socket which represents the socket opened to the current server for communication.

This function takes no parameters.

caller.getLatestLoad()

The `getLatestLoad` function allows a custom advisor to obtain the latest load value for a given server object. The load values are maintained in internal tables by the advisor base code and the manager daemon. This function call is useful if you want to make the behavior of one protocol or port dependent on the behavior of another. For example, you might use this function call in a custom advisor that disabled a particular application server if the Telnet server on that same machine was disabled.

The syntax is:

```
int caller.getLatestLoad (String clusterId, int port, String serverId)
```

The three arguments together define one server object.

This command has the following parameters:

clusterId

The cluster identifier of the server object for which to obtain the current load value. This argument must be a Java string.

port

The port number of the server object for which to obtain the current load value.

serverId

The server identifier of the server object for which to obtain the current load value. This argument must be a Java string. The return value is an integer.

- A positive return value represents the actual load value assigned for the object that was queried.
- The value -1 indicates that the server asked about is down.
- The value -2 indicates that the status of the server asked about is unknown.

caller.receive()

The `receive` function gets information from the socket connection. The syntax is:

```
caller.receive(StringBuffer *response)
```

This command has the following parameters:

response

This is a string buffer into which the retrieved data is placed. Additionally, the function returns an integer value with the following significance:

- 0 indicates data was sent successfully.
- A negative number indicates an error.

caller.send()

The `send` function uses the established socket connection to send a packet of data to the server, using the specified port. The syntax is as follows:

```
caller.send(String command)
```

This command has the following parameters:

command

This is a string containing the data to send to the server. The function returns an integer value with the following significance:

- 0 indicates data was sent successfully.
- A negative number indicates an error.

getLoad()

```
int getLoad( int iConnectTime; ADV_Thread *caller )
```

This function has the following parameters:

iConnectTime

The length of time, in milliseconds, that it took the connection to complete. This load measurement is performed by the advisor base code and passed to the custom advisor code, which can use or ignore the measurement when returning the load value. If the connection fails, this value is set to -1.

caller

The instance of the advisor base class where advisor base methods are provided. Function calls available to custom advisors: The methods, or functions, described in the following sections can be called from custom advisors. These methods are supported by the advisor base code. Some of these function calls can be made directly, for example, `function_name()`, but others require the prefix `caller`. `Caller` represents the base advisor instance that supports the custom advisor that is being executed.

getAdviseOnPort()

The `getAdviseOnPort` function returns the port number on which the calling custom advisor is running.

The return value is a Java integer (`int`), and the function does not take parameters.

Note: It is not possible for this value to change during one instantiation of an advisor.

getAdvisorName()

The `getAdvisorName` function returns a Java string with the suffix portion of your custom advisor's name. For example, for an advisor named `ADV_cdload.java`, this function returns the value `cdload`. This function takes no parameters. Note that it is not possible for this value to change during one instantiation of an advisor.

getInterval()

The `getInterval` function returns the advisor interval, that is, the number of seconds between advisor cycles. This value is equal to the default value set in the custom advisor's constructor, unless the value has been modified at run time by using the `dcontrol` command. The return value is a Java integer (`int`).

The function takes no parameters.

suppressBaseOpeningSocket()

The `suppressBaseOpeningSocket` function call allows a custom advisor to specify whether the base advisor code opens a TCP socket to the server on the custom advisor's behalf. If your advisor does not use direct communication with the server to determine its status, it might not be necessary to open this socket. This function call can be issued only once, and it must be issued from the "`ADV_AdvisorInitialize()` method" on page 128 routine.

The function takes no parameters.

Related tasks:

“Creating a custom advisor” on page 125

A custom advisor is a small piece of Java code, provided as a class file, that is called by the Load Balancer base code to determine the load on a server. The base code provides all necessary administrative services, including starting and stopping an instance of the custom advisor, providing status and reports, recording history information in a log file, and reporting advisor results to the manager component.

Related reference:

“Example: Sample advisor”

This is a sample advisor file called ADV_sample.

Example: Sample advisor

This is a sample advisor file called ADV_sample.

```
/ * *
* ADV_sample: The Load Balancer HTTP advisor
*
*
* This class defines a sample custom advisor for Load Balancer. Like all
* advisors, this custom advisor extends the function of the advisor base,
* called ADV_Base. It is the advisor base that actually performs most of
* the advisor's functions, such as reporting loads back to the Load Balancer
* for use in the Load Balancer's weight algorithm. The advisor base also
* performs socket connect and close operations and provides send and receive
* methods for use by the advisor. The advisor itself is used only for
* sending and receiving data to and from the port on the server being
* advised. The TCP methods within the advisor base are timed to calculate
* the load. A flag within the constructor in the ADV_base overwrites the
* existing load with the new load returned from the advisor if desired.
*
* Note: Based on a value set in the constructor, the advisor base supplies
* the load to the weight algorithm at specified intervals. If the actual
* advisor has not completed so that it can return a valid load, the advisor
* base uses the previous load.
*
* NAMING
*
* The naming convention is as follows:
*
* - The file must be located in the following Load Balancer directory:
*
*   ulb/servers/lib/CustomAdvisors/ (ulb\servers\lib\CustomAdvisors on Windows)
*
* - The Advisor name must be preceded with "ADV_". The advisor can be
* started with only the name, however; for instance, the "ADV_sample"
* advisor can be started with "sample".
*
* - The advisor name must be in lowercase.
*
* With these rules in mind, therefore, this sample is referred to as:
*
*   <base directory=""/>lib/CustomAdvisors/ADV_sample.class
*
*
* Advisors, as with the rest of Load Balancer, must be compiled with the
* prerequisite version of Java. To ensure access to Load Balancer classes, make
* sure that the ibmlb.jar file (located in the lib subdirectory of the base
* directory) is included in the system's CLASSPATH.
*
* Methods provided by ADV_Base:
*
* - ADV_Base (Constructor):
```



```

*
* - Pargs
*   - String sName = Name of the advisor
*   - String sVersion = Version of the advisor
*   - int iDefaultPort = Default port number to advise on
*   - int iInterval = Interval on which to advise on the servers
*   - String sDefaultName = Unused. Must be passed in as "".
*   - boolean replace = True - replace the load value being calculated
*                               by the advisor base
*                               False - add to the load value being calculated
*                                       by the advisor base
* - Return
*   - Constructors do not have return values.
*
* Because the advisor base is thread based, it has several other methods
* available for use by an advisor. These methods can be referenced using
* the CALLER parameter passed in getLoad().
*
* These methods are as follows:
*
* - send - Send a packet of information on the established socket connection
*         to the server on the specified port.
* - Pargs
*   - String sDataString - The data to be sent in the form of a string
* - Return
*   - int RC - Whether the data was successfully sent or not: zero indicates
*             data was sent; a negative integer indicates an error.
*
* - receive - Receive information from the socket connection.
* - Pargs
*   - StringBuffer sbDataBuffer - The data received during the receive call
* - Return
*   - int RC - Whether the data was successfully received or not; zero
*             indicates data was received; a negative integer indicates
*             an error.
*
* If the function provided by the advisor base is not sufficient,
* you can create the appropriate function within the advisor and
* the methods provided by the advisor base will then be ignored.
*
* An important question regarding the load returned is whether to apply
* it to the load being generated within the advisor base,
* or to replace it; there are valid instances of both situations.
*
* This sample is essentially the Load Balancer HTTP advisor. It functions
* very simply: a send request--an http head request--is issued. Once a
* response is received, the getLoad method terminates, flagging the advisor
* base to stop timing the request. The method is then complete. The
* information returned is not parsed; the load is based on the time
* required to perform the send and receive operations.
*/

```

```

package CustomAdvisors;
import com.ibm.internet.nd.advisors.*;
public class ADV_sample extends ADV_Base implements ADV_MethodInterface
{
    String COPYRIGHT =
        "(C) Copyright IBM Corporation 1997, All Rights Reserved.\n";
    static final String ADV_NAME = "Sample";
    static final int ADV_DEF_ADV_ON_PORT = 80;
    static final int ADV_DEF_INTERVAL = 7;

    // Note: Most server protocols require a carriage return ("\n") and line
    // feed ("\n") at the end of messages. If so, include them in
    // your string here.

    static final String ADV_SEND_REQUEST =

```

```

"HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
"IBM_Load_Balancer_HTTP_Advisor\r\n\r\n";

/**
 * Constructor.
 *
 * Parms: None; but the constructor for ADV_Base has several parameters
 *       that must be passed to it.
 */
public ADV_sample()
{
    super( ADV_NAME,
          "2.0.0.0-03.27.98",
          ADV_DEF_ADV_ON_PORT,
          ADV_DEF_INTERVAL,
          "", // not used false);
    super.setAdvisor( this );
}

/**
 * ADV_AdvisorInitialize
 *
 * Any Advisor-specific initialization that must take place after the
 * advisor base is started. This method is called only once and is
 * typically not used.
 */
public void ADV_AdvisorInitialize()
{
    return;
}

/**
 * getLoad()
 *
 * This method is called by the advisor base to complete the advisor's
 * operation, based on details specific to the protocol. In this sample
 * advisor, only a single send and receive are necessary; if more complex
 * logic is necessary, multiple sends and receives can be issued. For
 * example, a response might be received and parsed. Based on the
 * information learned thereby, another send and receive could be issued.
 *
 * Parameters:
 *
 * - iConnectTime - The current load as it refers to the length of time it
 *                  took to complete the connection to the server through
 *                  the specified port.
 *
 * - caller - A reference to the advisor base class where the Load
 *            Balancer-supplied methods are to perform simple TCP requests,
 *            mainly send and receive.
 *
 * Results:
 *
 * - The load - A value, expressed in milliseconds, that can either be added
 * to the existing load, or that can replace the existing load, as
 * determined by the constructor's "replace" flag.
 *
 * The larger the load, the longer it took the server to respond;
 * therefore, the lower the weight will become within the Load Balancer.
 *
 * If the value is negative, an error is assumed. An error from an
 * advisor indicates that the server the advisor is trying to reach is not
 * accessible and has been identified as being down. Load Balancer will
 * not attempt to load balance to a server that is down. Load Balancer will
 * resume load balancing to the server when a positive value is received.

```

```

*
*/
public int getLoad(int iConnectTime, ADV_Thread caller)
{
    int iRc;
    int iLoad = ADV_HOST_INACCESSIBLE; // -1

    // Send tcp request iRc = caller.send(ADV_SEND_REQUEST);
    if (iRc >= 0)
    {
        // Perform a receive
        StringBuffer sbReceiveData = new StringBuffer("");
        iRc = caller.receive(sbReceiveData);

        /**
         * In the normal advisor mode ("replace" flag is false), the load
         * returned is either 0 or 1 indicating the server is up or down.
         * If the receive is successful, a load of zero is returned
         * indicating that the load built within the base advisor is to be used.
         *
         * Otherwise ("replace" flag is true), return the desired load value.
         */

        if (iRc >= 0)
        {
            iLoad = 0;
        }
    }
    return iLoad;
}
} // End - ADV_sample

```

Example: Implementing standard advisors:

The following example demonstrates how to use a standard custom advisor.

This sample source code is similar to the standard Load Balancer HTTP advisor. It functions as follows:

1. A send request, a "HEAD/HTTP" command, is issued.
2. A response is received. The information is not parsed, but the response causes the getLoad method to terminate.
3. The getLoad method returns 0 to indicate success or -1 to indicate a failure.

This advisor operates in normal mode, so the load measurement is based on the elapsed time in milliseconds required to perform the socket open, send, receive, and close operations.

```

package CustomAdvisors;
import com.ibm.internet.lb.advisors.*;
public class ADV_sample extends ADV_Base implements ADV_MethodInterface {
    static final String ADV_NAME = "Sample";
    static final int ADV_DEF_ADV_ON_PORT = 80;
    static final int ADV_DEF_INTERVAL = 7;
    static final String ADV_SEND_REQUEST =
        "HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
        "IBM_Load_Balancer_HTTP_Advisor\r\n\r\n";

    //-----
    // Constructor
    public ADV_sample() {
        super(ADV_NAME, "3.0.0.0-03.31.00",
            ADV_DEF_ADV_ON_PORT, ADV_DEF_INTERVAL, "",
            false);
        super.setAdvisor( this );
    }
}

```

```

    }

    //-----
    // ADV_AdvisorInitialize
    public void ADV_AdvisorInitialize() {
        return; // usually an empty routine
    }

    //-----
    // getLoad

    public int getLoad(int iConnectTime, ADV_Thread caller) {
        int iRc;
        int iLoad = ADV_HOST_INACCESSIBLE; // initialize to inaccessible

        iRc = caller.send(ADV_SEND_REQUEST); // send the HTTP request to
        // the server
        if (0 <= iRc) { // if the send is successful
            StringBuffer sbReceiveData = new StringBuffer(""); // allocate a buffer
            // for the response
            iRc = caller.receive(sbReceiveData); // receive the result

            // parse the result here if you need to

            if (0 <= iRc) { // if the receive is successful
                iLoad = 0; // return 0 for success
            } // (advisor's load value is ignored by
            // base in normal mode)
        }
        return iLoad;
    }
}

```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called ADV_sample.

Example: Implementing a side stream advisor:

The following example demonstrates how a side stream advisor can be implemented. This sample illustrates suppressing the standard socket opened by the advisor base. Instead, this advisor opens a side stream Java socket to query a server. This procedure can be useful for servers that use a different port from normal client traffic to listen for an advisor query.

In this example, a server is listening on port 11999 and when queried returns a load value with a hexadecimal int "4". This sample runs in replace mode, that is, the last parameter of the advisor constructor is set to true and the advisor base code uses the returned load value rather than the elapsed time.

Note the call to `suppressBaseOpeningSocket()` in the initialization routine. Suppressing the base socket when no data will be sent is not required. For example, you might want to open the socket to ensure that the advisor can contact the server. Examine the needs of your application carefully before making this choice.

```

package CustomAdvisors;
import java.io.*;
import java.net.*;
import java.util.*;
import java.util.Date;
import com.ibm.internet.lb.advisors.*;
import com.ibm.internet.lb.common.*;
import com.ibm.internet.lb.server.SRV_ConfigServer;

```

```

public class ADV_sidea extends ADV_Base implements ADV_MethodInterface {
    static final String ADV_NAME = "sidea";
    static final int ADV_DEF_ADV_ON_PORT = 12345;
    static final int ADV_DEF_INTERVAL = 7;

    // create an array of bytes with the load request message
    static final byte[] abHealth = {(byte)0x00, (byte)0x00, (byte)0x00,
                                     (byte)0x04};

    public ADV_sidea() {
        super(ADV_NAME, "3.0.0.0-03.31.00", ADV_DEF_ADV_ON_PORT,
              ADV_DEF_INTERVAL, "",
              true); // replace mode parameter is true
        super.setAdvisor( this );
    }

    //-----
    // ADV_AdvisorInitialize
    public void ADV_AdvisorInitialize()
    {
        suppressBaseOpeningSocket(); // tell base code not to open the
                                     // standard socket
        return;
    }

    //-----
    // getLoad
    public int getLoad(int iConnectTime, ADV_Thread caller) {
        int iRc;
        int iLoad = ADV_HOST_INACCESSIBLE; // -1
        int iControlPort = 11999; // port on which to communicate
                                  // with the server
        String sServer = caller.getCurrentServerId(); // address of server to query
        try {
            socket soServer = new Socket(sServer, iControlPort); // open socket to
                                                                  // server
            DataInputStream disServer = new DataInputStream(
                soServer.getInputStream());
            DataOutputStream dosServer = new DataOutputStream(
                soServer.getOutputStream());
            int iRecvTimeout = 10000; // set timeout (in milliseconds)
                                      // for receiving data
            soServer.setSoTimeout(iRecvTimeout);
            dosServer.writeInt(4); // send a message to the server
            dosServer.flush();
            iLoad = disServer.readByte(); // receive the response from the server
        } catch (exception e) {
            system.out.println("Caught exception " + e);
        }
        return iLoad; // return the load reported from the server
    }
}

```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called ADV_sample.

Example: Implementing a two-port advisor:

The following example shows how to implement a two-port advisor. This custom advisor sample demonstrates the capability to detect failure for one port of a server based upon both its own status and on the status of a different server daemon that is running on another port on the same server machine.

For example, if the HTTP daemon on port 80 stops responding, you might also want to stop routing traffic to the SSL daemon on port 443.

This advisor is more aggressive than standard advisors, because it considers any server that does not send a response to have stopped functioning, and marks it as down. Standard advisors consider unresponsive servers to be very slow. This advisor marks a server as down for both the HTTP port and the SSL port based on a lack of response from either port.

To use this custom advisor, the administrator starts two instances of the advisor: one on the HTTP port, and one on the SSL port. The advisor instantiates two static global hash tables, one for HTTP and one for SSL. Each advisor tries to communicate with its server daemon and stores the results of this event in its hash table. The value that each advisor returns to the base advisor class depends on both the ability to communicate with its own server daemon and the ability of the partner advisor to communicate with its daemon.

The following custom methods are used.

- `ADV_nte()` is a simple container object to hold information about a server. These objects are stored in the hash table as table elements. Each object has a time stamp that is used to determine whether the element is current.
- `putNte()` and `getNte()` are synchronized methods that ensure that the two advisor instances access the hash table in a controlled fashion.
- `getLoadHTTP` is a method that queries the responsiveness of an HTTP server. It is a low-level routine and does not gather or use information about SSL.
- `getLoadSSL()` is a method that queries the responsiveness of an SSL server. It is a low-level routine and does not gather or use information about HTTP.
- `getLoad()` is the entry point routine for this custom advisor. It can handle both protocols and can store and fetch information from the hash table. This is the routine that links the two ports.

The following error conditions are detected:

- Unresponsive server machine - The base advisor classes periodically send a ping signal to the server address. If the address is not reachable, the base advisor classes marks the server down. Neither of the two instances of the custom advisor is called, and both servers on that machine are marked down.
- One daemon on a server machine becomes unresponsive, but the other is working - When the base code attempts to open a socket with the server, the connection is refused, and the base advisor for this protocol marks the server as down. The custom advisor code for that protocol is not called. Although the custom advisor for the other protocol continues communicating with its server, it learns from the hash table that the other custom advisor cannot communicate with its server daemon. Therefore, the second protocol's advisor also marks its server as down.
- One daemon does not send a response, but the other daemon does - The custom advisor for the unresponsive protocol detects the failure to communicate, marks the server as down, and stores the data in the hash table. The custom advisor for the other port learns that information from the hash table and marks its server as down.

This sample is written to link ports 80 for HTTP and 443 for SSL, but it can be tailored to any combination of ports:

```

package CustomAdvisors;
import java.io.*;
import java.net.*;
import java.util.*;
import java.util.Date;
import com.ibm.internet.lb.advisors.*;
import com.ibm.internet.lb.common.*;
import com.ibm.internet.lb.manager.*;
import com.ibm.internet.lb.server.SRV_ConfigServer;

//-----
// Define the table element for the hash tables used in this custom advisor

class ADV_nte implements Cloneable {
    private String sCluster;
    private int iPort;
    private String sServer;
    private int iLoad;
    private Date dTimestamp;

//-----
// constructor

    public ADV_nte(String sClusterIn, int iPortIn, String sServerIn,
        int iLoadIn) {
        sCluster = sClusterIn;
        iPort = iPortIn;
        sServer = sServerIn;
        iLoad = iLoadIn;
        dTimestamp = new Date();
    }

//-----
// check whether this element is current or expired
    public boolean isCurrent(ADV_twop oThis) {
        boolean bCurrent;
        int iLifetimeMs = 3 * 1000 * oThis.getInterval();    // set lifetime as
                                                            // 3 advisor cycles

        Date dNow = new Date();
        Date dExpires = new Date(dTimestamp.getTime() + iLifetimeMs);

        if (dNow.after(dExpires)) {
            bCurrent = false;
        } else {
            bCurrent = true;
        } return bCurrent;
    }

//-----
// value accessor(s)

    public int getLoadValue() { return iLoad; }

//-----
// clone (avoids corruption between threads)

    public synchronized Object Clone() {
        try {
            return super.clone();
        } catch (cloneNotSupportedException e) {
            return null;
        }
    }
}

//-----

```

```

// define the custom advisor

public class ADV_twop extends ADV_Base
    implements ADV_MethodInterface, ADV_AdvisorVersionInterface {
    static final int ADV_TWOP_PORT_HTTP = 80;
    static final int ADV_TWOP_PORT_SSL = 443;

    //-----
    // define tables to hold port-specific history information

    static Hashtable htTwopHTTP = new Hashtable();
    static Hashtable htTwopSSL = new Hashtable();
    static final String ADV_TWOP_NAME = "twop";
    static final int ADV_TWOP_DEF_ADV_ON_PORT = 80;
    static final int ADV_TWOP_DEF_INTERVAL = 7;
    static final String ADV_HTTP_REQUEST_STRING =
        "HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
        "IBM_LB_Custom_Advisor\r\n\r\n";

    //-----
    // create byte array with SSL client hello message

    public static final byte[] abClientHello = {
        (byte)0x80, (byte)0x1c,
        (byte)0x01, // client hello
        (byte)0x03, (byte)0x00, // SSL version
        (byte)0x00, (byte)0x03, // cipher spec len (bytes)
        (byte)0x00, (byte)0x00, // session ID len (bytes)
        (byte)0x00, (byte)0x10, // challenge data len (bytes)
        (byte)0x00, (byte)0x00, (byte)0x03, // cipher spec
        (byte)0x1A, (byte)0xFC, (byte)0xE5, (byte)0x20, // challenge data
        (byte)0xFD, (byte)0x3A, (byte)0x3C, (byte)0x18,
        (byte)0xAB, (byte)0x67, (byte)0xB0, (byte)0x52,
        (byte)0xB1, (byte)0x1D, (byte)0x55, (byte)0x44, (byte)0x0D, (byte)0x0A };

    //-----
    // constructor

    public ADV_twop() {
        super(ADV_TWOP_NAME, VERSION, ADV_TWOP_DEF_ADV_ON_PORT,
            ADV_TWOP_DEF_INTERVAL, "",
            false); // false = load balancer times the response
        setAdvisor ( this );
    }

    //-----
    // ADV_AdvisorInitialize

    public void ADV_AdvisorInitialize() {
        return; }

    //-----
    // synchronized PUT and GET access routines for the hash tables

    synchronized ADV_nte getNte(Hashtable ht, String sName, String sHashKey) {
        ADV_nte nte = (ADV_nte)(ht.get(sHashKey));
        if (null != nte) {
            nte = (ADV_nte)nte.clone();
        }
        return nte;
    }
    synchronized void putNte(Hashtable ht, String sName, String sHashKey,
        ADV_nte nte) { ht.put(sHashKey,nte); return;
    }

    //-----

```



```

// getLoadHTTP - determine HTTP load based on server response

int getLoadHTTP(int iConnectTime, ADV_Thread caller) {
    int iLoad = ADV_HOST_INACCESSIBLE;
    int iRc = caller.send(ADV_HTTP_REQUEST_STRING); // send request message
                                                    // to server
    if (0 <= iRc) { // did the request return a failure?
        StringBuffer sbReceiveData = new StringBuffer("") // allocate a buffer
                                                            // for the response
        iRc = caller.receive(sbReceiveData); // get response from server

        if (0 <= iRc) { // did the receive return a failure?
            if (0 < sbReceiveData.length()) { // is data there?
                iLoad = SUCCESS; // ignore retrieved data and
                                // return success code
            }
        }
    }
    return iLoad;
}

//-----
// getLoadSSL() - determine SSL load based on server response

int getLoadSSL(int iConnectTime, ASV_Thread caller) {
    int iLoad = ADV_HOST_INACCESSIBLE;
    int iRc;

    CMNByteArrayWrapper cbawClientHello = new CMNByteArrayWrapper(
                                                abClientHello);
    Socket socket = caller.getSocket();

    try {
        socket.getOutputStream().write(abClientHello); // Perform a receive.
        socket.getInputStream().read(); // If receive is successful,
                                    // return load of 0. We are not
                                    // concerned with data's contents,
                                    // and the load is calculated by
                                    // the ADV_Thread thread.

        iLoad = 0;
    } catch (IOException e) { // Upon error, iLoad will default to it.
    }
    return iLoad;
}

//-----
// getLoad - merge results from the HTTP and SSL methods

public int getLoad(int iConnectTime, ADV_Thread caller) {
    int iLoadHTTP;
    int iLoadSSL;
    int iLoad;
    int iRc;

    String sCluster = caller.getCurrentClusterId(); // current cluster address
    int iPort = getAdviseOnPort();
    String sServer = caller.getCurrentServerId();
    String sHashKey = sCluster + ":" + sServer; // hash table key

    if (ADV_TWOP_PORT_HTTP == iPort) { // handle an HTTP server
        iLoadHTTP = getLoadHTTP(iConnectTime, caller); // get the load for HTTP

        ADV_nte nteHTTP = newADV_nte(sCluster, iPort, sServer, iLoadHTTP);
        putNte(htTwopHTTP, "HTTP", sHashKey, nteHTTP); // save HTTP load
                                                        // information
    }
}

```

```

ADV_nte nteSSL = getNte(htTwopSSL, "SSL", sHashKey); // get SSL
                                                    // information
if (null != nteSSL) {
    if (true == nteSSL.isCurrent(this)) {          // check the time stamp
        if (ADV_HOST_INACCESSIBLE != nteSSL.getLoadValue()) { // is SSL
                                                                // working?
            iLoad = iLoadHTTP;
        } else {                                     // SSL is not working, so mark the HTTP server down
            iLoad= ADV_HOST_INACCESSIBLE;
        }
    } else {                                        // SSL information is expired, so mark the
                                                // HTTP server down
        iLoad = ADV_HOST_INACCESSIBLE;
    }
} else {                                           // no load information about SSL, report
                                                // getLoadHTTP() results
    iLoad = iLoadHTTP;
}
}
else if (ADV_TWOP_PORT_SSL == iPort) {            // handle an SSL server
    iLoadSSL = getLoadSSL(iConnectTime, caller); // get load for SSL

    ADV_nte nteSSL = new ADV_nte(sCluster, iPort, sServer, iLoadSSL);
    putNte(htTwopSSL, "SSL", sHashKey, nteSSL); // save SSL load info.

    ADV_nte nteHTTP = getNte(htTwopHTTP, "SSL", sHashKey); // get HTTP
                                                            // information
    if (null != nteHTTP) {
        if (true == nteHTTP.isCurrent(this)) { // check the timestamp
            if (ADV_HOST_INACCESSIBLE != nteHTTP.getLoadValue()) { // is HTTP
                                                                    // working?
                iLoad = iLoadSSL;
            } else { // HTTP server is not working, so mark SSL down
                iLoad = ADV_HOST_INACCESSIBLE;
            }
        } else { // expired information from HTTP, so mark SSL down
            iLoad = ADV_HOST_INACCESSIBLE;
        }
    } else { // no load information about HTTP, report
                                                // getLoadSSL() results
        iLoad = iLoadSSL;
    }
}
}

//-----
// error handler

else {
    iLoad = ADV_HOST_INACCESSIBLE;
}
return iLoad;
}
}

```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called ADV_sample.

Example: Implementing the WAS advisor:

The following examples show how custom advisors can be implemented.

A sample custom advisor for WebSphere Application Server is included in the *install_root/servers/samples/CustomAdvisors/* directory. The full code is not duplicated in this document. Ensure that the following will be implemented:

- ADV_was.java is the advisor source code file that is compiled and run on the Load Balancer machine.
- LBAdvisor.java.servlet is the servlet source code that must be renamed to LBAdvisor.java, compiled, and run on the WebSphere Application Server machine.

The complete advisor is only slightly more complex than the sample. It adds a specialized parsing routine that is more compact than the StringTokenizer example shown in the topic “Example: Using data returned from advisors” on page 144.

The more complex part of the sample code is in the Java servlet. Among other methods, the servlet contains two methods required by the servlet specification: `init()` and `service()`, and one method, `run()`, that is required by the `Java.lang.thread` class.

- `init()` is called once by the servlet engine at initialization time. This method creates a thread named `_checker` that runs independently of calls from the advisor and sleeps for a period of time before resuming its processing loop.
- `service()` is called by the servlet engine each time the servlet is invoked. In this case, the method is called by the advisor. The `service()` method sends a stream of ASCII characters to an output stream.
- `run()` contains the core of the code execution. It is called by the `start()` method that is called from within the `init()` method.

The relevant fragments of the servlet code appear below:

```
...
public void init(ServletConfig config) throws ServletException {
    super.init(config);
    ...
    _checker = new Thread(this);
    _checker.start();
}

public void run() {
    setStatus(GOOD);

    while (true) {
        if (!getKeepRunning())
            return;
        setStatus(figureLoad());
        setLastUpdate(new java.util.Date());

        try {
            _checker.sleep(_interval * 1000);
        } catch (Exception ignore) { ; }
    }
}

public void service(HttpServletRequest req, HttpServletResponse res)
    throws ServletException, IOException {
    ServletOutputStream out = null;
    try {
        out = res.getOutputStream();
    } catch (Exception e) { ... }
    ...
    res.setContentType("text/x-application-LBAdvisor");
    out.println(getStatusString());
    out.println(getLastUpdate().toString());
    out.flush(); return;
}
...
```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called ADV_sample.

Example: Using data returned from advisors:

Whether you use a standard call to an existing part of the application server or add a new piece of code to be the server-side counterpart of your custom advisor, you possibly want to examine the load values returned and change server behavior.

The Java StringTokenizer class, and its associated methods, make this investigation easy to do. The content of a typical HTTP command might be

```
GET /index.html HTTP/1.0 90
```

A typical response to this command might be the following:

```
HTTP/1.1 200 OK
Date: Mon, 20 November 2000 14:09:57 GMT
Server: Apache/1.3.12 (Linux and UNIX)
Content-Location: index.html.en
Vary: negotiate
TCN: choice
Last-Modified: Fri, 20 Oct 2000 15:58:35 GMT
ETag: "14f3e5-1a8-39f06bab;39f06a02"
Accept-Ranges: bytes
Content-Length: 424
Connection: close
Content-Type: text/html
Content-Language: en

<!DOCTYPE HTML PUBLIC "-//w3c//DTD HTML 3.2 Final//EN">
<HTML><HEAD><TITLE>Test Page</TITLE></HEAD>
<BODY><H1>Apache server</H1>
<HR>
<P><P>This Web server is running Apache 1.3.12.
</P>
<P><IMG SRC="apache_pb.gif" ALT="">
</P></P>
</HR>
</BODY></HTML>
```

The items of interest are contained in the first line, specifically the HTTP return code. The HTTP specification classifies return codes that can be summarized as follows:

- 2xx return codes are successes
- 3xx return codes are redirections
- 4xx return codes are client errors
- 5xx return codes are server errors

If you know precisely what codes the server can possibly return, your code might not need to be as detailed as this example. However, keep in mind that limiting the return codes you detect might limit the future flexibility of your program.

The following example is a stand-alone Java program that contains a minimal HTTP client. The example invokes a simple, general-purpose parser for examining HTTP responses.

```

import java.io.*;
import java.util.*;
import java.net.*;

public class ParseTest {
    static final int iPort = 80;
    static final String sServer = "www.ibm.com";
    static final String sQuery = "GET /index.html HTTP/1.0\r\n\r\n";
    static final String sHTTP10 = "HTTP/1.0";
    static final String sHTTP11 = "HTTP/1.1";

    public static void main(String[] Arg) {
        String sHTTPVersion = null;
        String sHTTPReturnCode = null;
        String sResponse = null; int iRc = 0;
        BufferedReader brIn = null;
        PrintWriter psOut = null;
        Socket soServer= null;
        StringBuffer sbText = new
        StringBuffer(40);

        try {
            soServer = new Socket(sServer, iPort);
            brIn = new BufferedReader(new InputStreamReader(
                soServer.getInputStream()));
            psOut = new PrintWriter(soServer.getOutputStream());
            psOut.println(sQuery);
            psOut.flush();
            sResponse = brIn.readLine();
            try {
                soServer.close();
            } catch (Exception sc) {};
        } catch (Exception swr) {};

        StringTokenizer st = new StringTokenizer(sResponse, " ");
        if (true == st.hasMoreTokens()) {
            sHTTPVersion = st.nextToken();
            if (sHTTPVersion.equals(sHTTP10) || sHTTPVersion.equals(sHTTP11)) {
                System.out.println("HTTP Version: " + sHTTPVersion);
            } else {
                System.out.println("Invalid HTTP Version: " + sHTTPVersion);
            }
        } else {
            System.out.println("Nothing was returned");
            return;
        }

        if (true == st.hasMoreTokens()) {
            sHTTPReturnCode = st.nextToken();
            try {
                iRc = Integer.parseInt(sHTTPReturnCode);
            } catch (NumberFormatException ne) {};

            switch (iRc) {
            case(200):
                System.out.println("HTTP Response code: OK, " + iRc);
                break;
            case(400): case(401): case(402): case(403): case(404):
                System.out.println("HTTP Response code: Client Error, " + iRc);
                break;
            case(500): case(501): case(502): case(503):
                System.out.println("HTTP Response code: Server Error, " + iRc);
                break;
            default:
                System.out.println("HTTP Response code: Unknown, " + iRc);
                break;
            }
        }
    }
}

```

```

    }

    if (true == st.hasMoreTokens()) {
        while (true == st.hasMoreTokens()) {
            sbText.append(st.nextToken());
            sbText.append(" ");
        }
        System.out.println("HTTP Response phrase: " + sbText.toString());
    }
}
}
}

```

Related reference:

“Example: Sample advisor” on page 132
 This is a sample advisor file called ADV_sample.

Configuring high availability

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the "health" of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

About this task

When you configure high availability, each of the two machines is assigned a specific role, either primary or backup. The primary machine sends connection data to the backup machine on an ongoing basis. While the primary is active (load balancing), the backup is in a standby state, continually updated and ready to take over, if necessary.

The communication sessions between the two machines are referred to as heartbeats. The heartbeats allow each machine to monitor the health of the other. If the backup machine detects that the active machine has failed, it will take over and begin load balancing. At that point the statuses of the two machines are reversed: the backup machine becomes active and the primary becomes standby.

Note: In the high availability configuration, both primary and backup machines must be on the same subnet with identical configuration.

For the complete syntax see “dscontrol highavailability” on page 236. For a more complete discussion of many of the tasks below, see “Configuring the Dispatcher machine” on page 69.

Tips for configuring high availability:

1. To configure a single Dispatcher machine to route packets without a backup, do not issue any of the high availability commands at startup.
2. To convert two Dispatcher machines configured for high availability to one machine running alone, stop the executor on one of the machines, then delete the high availability features (the heartbeats, reach, and backup) on the other.
3. **Linux for s/390:** In both of the two cases above, you must alias the network interface card with cluster addresses, as required.
4. When running two Dispatcher machines in a high availability configuration, unexpected results can occur if you set any of the parameters for the executor, cluster, port, or server (for example, port stickytime) to different values on the two machines.

Procedure

1. If you are running Linux for s/390 operating systems in layer 3 mode, create alias script files on each of the two Dispatcher machines. See “Scripts to run with high availability” on page 151 for more information on the available scripts. The scripts you create should contain commands to:

- Configure the cluster ip address on the interface
- Add an iptables rule to drop incoming packets destined to the cluster address

See the topic “Configuring the Dispatcher machine” on page 69 for more information.

2. Start the server on both Dispatcher server machines.
3. Start the executor on both machines.
4. Ensure that the non-forwarding address (NFA) of each Dispatcher machine is configured, and is a valid IP address for the subnet of the Dispatcher machines.
5. Add the heartbeat information on both machines:

```
dscontrol highavailability heartbeat add source_address destination_address
```

Source_address and *destination_address* are the IP addresses (either DNS names or IP addresses) of the Dispatcher machines. The values will be reversed on each machine. For example:

```
Primary - highavailability heartbeat add 9.67.111.3 9.67.186.8
Backup - highavailability heartbeat add 9.67.186.8 9.67.111.3
```

At least one heartbeat pair must have the NFAs of the pair as the source and destination address. If possible, at least one of the heartbeat pairs should be across a separate subnet than the regular cluster traffic. Keeping the heartbeat traffic distinct will help prevent false takeovers during very heavy network loads and also improve complete recovery times after a failover.

- a. Optional: Set the number of seconds that the executor uses to timeout high availability heartbeats. The default is 2 seconds. For example:

```
dscontrol executor set hatimeout 3
```

6. On both machines, configure the list of IP addresses that the Dispatcher must be able to reach in order to ensure full service, using the reach add command. Reach targets are recommended but not required. See “Detecting server failures with heartbeats and reach targets” on page 149 for more information. For example:

```
dscontrol highavailability reach add 9.67.125.18
```

7. Add the backup information to each machine:

- a. For the primary machine:

```
dscontrol highavailability backup add primary [auto | manual] port
```

- b. For the backup machine:

```
dscontrol highavailability backup add backup [auto | manual] port
```

Note: Select an unused port on your machines as the port. The port number entered will be used as a key to ensure the correct host is receiving the packet.

8. Check the high availability status on each machine:

```
dscontrol highavailability status
```

The machines should each have the correct role (backup or primary) and states. The primary should be active; the backup should be in standby mode. The recovery strategies must be the same.

9. Set up the cluster, port, and server information on both machines.
10. Optional: Enable replication of connection and affinity records.

Note: This feature provides the ability to replicate connection and affinity records between high availability partners. When the records are replicated, connection and affinity states are preserved so that the connections can continue even after takeover has taken place.

- Enable replication for connection and affinity records:
`dscontrol port set cluster@port repstrategy both`
- Enable replication for only connection records:
`dscontrol port set cluster@port repstrategy connection`
- Enable replication for only affinity records:
`dscontrol port set cluster@port repstrategy affinity`
- To disable replication:
`dscontrol port set cluster@port repstrategy none`

For more information on this command, refer to the topic on the `dscontrol port` command.

11. Optional: Suppress the unreachable packets for the ICMP port that are generated by the operating system in response to heartbeat packets
`# iptables -t filter -A INPUT -p udp --destination-port <port> -j DROP`

`<port>` is the port number for the heartbeat. For more information on this issue, read the technote [ICMP Port unreachable sent when HA packet received](#).

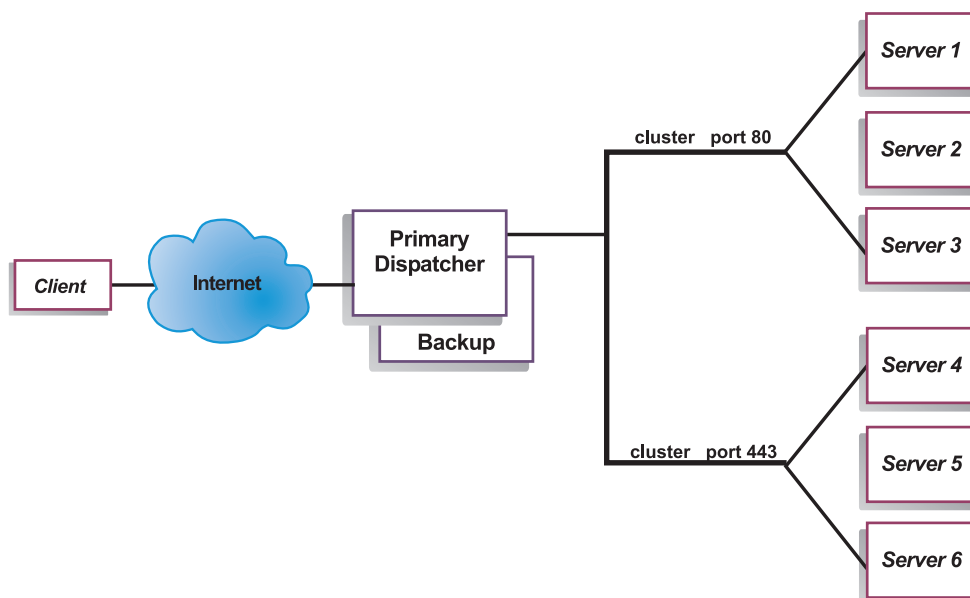
12. Start the manager and advisors on both machines.

Related reference:

“`dscontrol port`” on page 244

Configure ports and port settings with the `dscontrol port` command.

How high availability works



To improve Dispatcher availability, the Dispatcher high availability functions as follows:

1. Two Dispatchers with connectivity to the same clients, and the same cluster of servers, as well as connectivity between the Dispatchers. Both Dispatchers must run on the same type of operating system and platform.
2. A “heartbeat” mechanism between the two Dispatchers detects a Dispatcher failure. At least one heartbeat pair must have the NFAs of the pair as the source and destination address. If possible, at least one of the heartbeat pairs should be across a separate subnet than the regular cluster traffic. Keeping the heartbeat traffic distinct will help prevent false takeovers during very heavy network loads and also improve complete recovery times after a failover.
3. A list of reach targets, addresses that both Dispatcher machines must be able to contact in order to load balance traffic normally. For more information, see “Detecting server failures with heartbeats and reach targets.”
4. Synchronization of the Dispatcher information
5. Logic to elect the active Dispatcher which is in charge of a given cluster of servers, and the standby Dispatcher which continuously gets synchronized for that cluster of servers.
6. A mechanism to perform IP takeover, when the logic or an operator decides to switch active and standby.

Planning for high availability

When configuring for high availability, consider that the Load Balancer machine is supported with the following limitations or special considerations:

- If you are using IPv6 protocol on your machine and want to use high availability, you must check to see if protocol 58 is defined to be ICMPv6 in the protocol file.
- In the high availability configuration, both primary and backup machines must be on the same subnet with identical configuration.
- The heartbeat pairs (which is the mechanism between the primary and standby Dispatchers to detect Dispatcher failure) must be both IPv4 format or both IPv6 format.
- In a high availability or a stand-alone environment, you must not alias the cluster address against the network adaptor.
- The HighAvailChange script can be moved from the *install_root/servers/samples* directory to the *install_root/servers/bin* directory to log high availability state changes for the Dispatcher machine, but this script does not need to be changed.

Related tasks:

“Configuring high availability” on page 146

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the “health” of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

Detecting server failures with heartbeats and reach targets

Configure heartbeats and reach targets to detect server failures and control when failovers can occur.

About this task

Besides the basic criteria of failure detection (the loss of connectivity between active and standby Dispatchers, detected through the heartbeat messages), there is another failure detection mechanism named reachability criteria. When you configure the Dispatcher you can provide a list of hosts that each of the Dispatchers should be able to reach in order to work correctly. The two high availability partners continually communicate with each other through heartbeats, and they update one another on how many reach targets either one of them can ping. If the standby pings more reach targets than the active, a failover occurs.

Heartbeats: Heartbeats are sent by the active Dispatcher and are expected to be received by the standby Dispatcher every half second. If the standby Dispatcher fails to receive a heartbeat within 2 seconds, a failover begins. All heartbeats must break for a takeover from the standby Dispatcher to occur. In other words, when two heartbeat pairs are configured, both heartbeats must break. To stabilize a high availability environment and to avoid failover, add more than one heartbeat pair.

Reach target considerations: For reach targets, you should choose at least one host for each subnet your Dispatcher machine uses. The hosts could be routers, IP servers or other types of hosts. Host reachability is obtained by the reach advisor, which pings the host. Failover takes place either if the heartbeat messages cannot go through, or if the reachability criteria are met better by the standby Dispatcher than by the primary Dispatcher. To make the decision based on all available information, the active Dispatcher regularly sends the standby Dispatcher its reachability capabilities. The standby Dispatcher then compares those capabilities with its own and decides whether to switch.

Note: When you configure the reach target, the reach advisor must also be started. The reach advisor starts automatically when you start the manager function. For more information on the reach advisor, see “List of advisors” on page 114.

Procedure

- Use the `dscontrol highavailability` command to add or delete a reach target to a server:

```
dscontrol highavailability reach add|delete address mask
```

- Use the `dscontrol highavailability` command with the heartbeat option to add a heartbeat:

```
dscontrol highavailability heartbeat add srcaddress dstaddress
```

To delete a heartbeat, use the following:

```
dscontrol highavailability heartbeat delete address
```

Related tasks:

“Scripts to run with high availability” on page 151

High Availability recovery strategy for failed servers

The recovery strategy dictates how Load Balancer behaves when one Dispatcher machine fails and there is another configured as a backup.

Two Dispatcher machines are configured: the primary machine, and a second machine called the backup. At startup, the primary machine sends all the connection data to the backup machine until that machine is synchronized. The

primary machine becomes active, that is, it begins load balancing. The backup machine, meanwhile, monitors the status of the primary machine, and is said to be in standby state.

If the backup Load Balancer machine detects that the primary machine has failed, it performs a takeover load balancing functions and becomes the active machine. After the primary machine has once again become operational, the machines respond according to how the recovery strategy has been configured by the user.

There are two kinds of strategy:

- **Automatic** - The primary machine resumes routing packets as soon as it becomes operational again.
- **Manual** - intervention is required to return the primary machine to active state and reset the backup machine to standby. The manual recovery strategy allows you to force the routing of packets to a particular machine, using the takeover command. Manual recovery is useful when maintenance is being performed on the other machine

Note: The strategy parameter must be the same for both machines.

Related tasks:

“Configuring high availability” on page 146

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the "health" of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

“Detecting server failures with heartbeats and reach targets” on page 149

Configure heartbeats and reach targets to detect server failures and control when failovers can occur.

“Scripts to run with high availability”

Scripts to run with high availability

The following sample script can be used:

highAvailChange

The HighAvailChange script runs whenever the high availability state changes within the Load Balancer. You can create this script to use state change information, for instance, to alert an Administrator or simply record the event.

On Linux for s/390: Load Balancer issues a gratuitous ARP to move IP addresses from one Load Balancer to another. This mechanism is therefore tied to the underlying network type. When running Linux for s/390, Load Balancer can natively do high availability takeovers, complete with IP address moves, only on those interfaces which can issue a gratuitous ARP and configure the address on the local interface. This mechanism will not work properly on point-to-point interfaces such as IUCV and CTC and will not work properly in certain configurations of qeth/QDIO.

Related tasks:

“Configuring high availability” on page 146

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the "health" of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

“Detecting server failures with heartbeats and reach targets” on page 149

Configure heartbeats and reach targets to detect server failures and control when failovers can occur.

Related reference:

“High Availability recovery strategy for failed servers” on page 150

The recovery strategy dictates how Load Balancer behaves when one Dispatcher machine fails and there is another configured as a backup.

Using NAT forwarding to forward traffic

Using Network Address Translation (NAT) Dispatcher capability removes the limitation for the backend servers to be on a locally attached network. With the NAT forwarding method, Dispatcher load balances the incoming request to the server. The server returns the response to Dispatcher. The Dispatcher machine then returns the response to the client.

Before you begin

The following limitations apply to NAT forwarding:

- Dispatcher's implementation of NAT is a simple implementation of this feature. It analyzes and operates only the contents of TCP/IP packet headers. It does not analyze the contents of the data portion of the packets. For Dispatcher, NAT does not work with application protocols, such as FTP, which embed the addresses or port numbers in the data portion of the messages. This is a known limitation of header-based NAT.
- Dispatcher's NAT cannot work with the affinity selection algorithm. It works with connection and connection + affinity selection algorithms only.
- Dispatcher does not support Network Address Port Translation (NAPT) intrinsically as part of its implementation. Hence, `mapport` is obsolete and it is recommended to use the OS implementations of NAPT techniques directly on the backend servers.

About this task

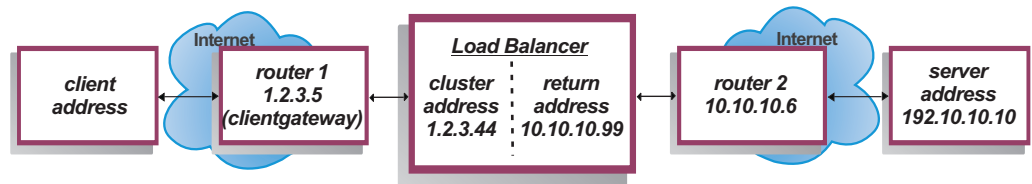
You need three IP addresses for the Dispatcher machine – NFA, cluster, and return address. The return address is a unique address or host name that you configure on the Dispatcher machine. Dispatcher uses the return address as its source address when load balancing the client's request to the server. Using the return address ensures that the server returns the packet to the Dispatcher machine, rather than sending the packet directly to the client (Dispatcher then forwards the IP packet to the client).

You must specify the return address value when you add the server. You cannot modify the return address unless you remove the server and then add it again. The return address cannot be the same as the cluster, server, or NFA address. When you use the NAT forwarding method, you must define a return address for communication between Load Balancer and the backend servers. The number of

connections that Load Balancer can keep active with the backend server is limited by the return addresses and the server combination.

In addition, you must configure a client-gateway, which is the router to send the response back to the client. Specify the router address to reach the backend server. If the backend server is in same subnet as the dispatcher machine, then the router IP address must be same as the backend server.

Important: All addresses must belong to same IP family.



Procedure

1. Set the client gateway parameter on the **dscontrol executor set** command for the Load Balancer Server. Client gateway is an IP address that is used as the router address through which traffic in the return direction is forwarded from Load Balancer to clients. This value must be set to a non-zero IP address before you can use NAT. For example, `dscontrol executor set clientgateway 1.2.3.5`
2. Add a NAT enabled server on the Load Balancer Server. For example, `dscontrol server addclusterip@clusterport@servername address server_ip_address nat yes returnaddress return_ip_address router router_ip_address`. Use the `dscontrol server` command:

```
dscontrol server add 1.2.3.44@22@1bopt1 address 192.10.10.10 nat yes returnaddress 10.10.10.99 router 10.10.10.6
```

Attention: No configuration is required on the managed server, unless it is for port translation.

Attention: NAT is used when the backend application must receive the data on a port that is different from the port on which data is sent by the client. As an alternate to **mapport**, use port translation techniques that are available in the operating system on which the application is deployed, for example: `iptables` on Linux, `netsh (portproxy)` on Windows, `ipfilter` on AIX, Solaris, and HP-UX.

Related tasks:

Chapter 5, “Configuring Load Balancer,” on page 61

This section focuses on configuring production environments and realistic test environments.

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

Related reference:

“dscontrol executor” on page 233

Control the executor function with the dscontrol executor command.

“dscontrol logstatus” on page 239

Use this command to display the log settings for a server.

“dscontrol manager” on page 239

You can control the manager function with the dscontrol manager command.

“dscontrol metric” on page 243

You can configure system metrics with the dscontrol metric command.

“dscontrol server” on page 249

Configure servers and modify existing server configurations with the dscontrol server command.

Use encapsulation forwarding to forward traffic across network segments

Use encapsulation forwarding when the back-end server is not located on the same network segment or if you are using virtualization technology and need to forward packets that are otherwise unable to be forwarded.

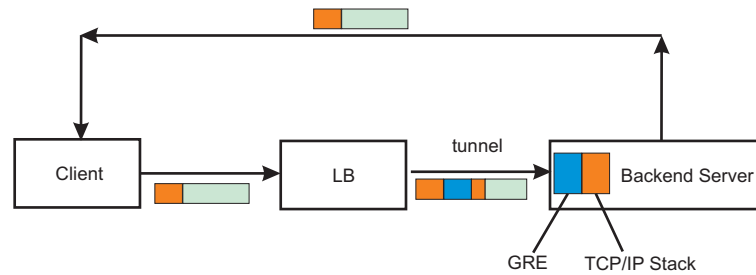
About this task

In a typical configuration, Load Balancer receives a packet, P, and forwards it as packet P', where only the time-to-live (TTL) has been decremented. When you enable encapsulation, Load Balancer receives a packet P, and forwards it as E(P'), where the encapsulated packet E contains P'. The outer packet E has a unique IP header, which permits Load Balancer to forward packets across routers and across some types of virtualization technology that you could not otherwise forward packets across.

Encapsulation forwarding:

- Is implemented like MAC forwarding:
 - Packets from server to client do not go through the load balancer
 - Alias the loopback device to cluster address on back-end server
- Requires that you configure an IP/IP or GRE tunnel only on the back-end server.
- Does not require you to add routes while configuring the tunnel.

Load Balancer will act as the tunnel on the other end.



Additionally, this functionality allows you to forward packets to Solaris zones or AIX workload partitions that are on the same host, since Load Balancer can use the existing stack configuration instead of bypassing it entirely.

Procedure

1. On the Load Balancer machine, add a server with encapsulation enabled. When this server is selected to forward the packet, it is encapsulated. Use the `dscontrol server` command:

```
dscontrol server set encap_source_IP encapforward [yes/no] encaptype [ipip/gre] encapcond [auto/always]
```

For example, you can type the following at the prompt:

```
dscontrol server set 1.2.3.4@80@1.2.3.5 encapforward yes encaptype ipip encapcond always
```

2. Configure the IPIP or GRE tunnel on the back-end server for network traffic. For example, you can type the following:

```
ifconfig gre0 tunnel 9.184.119.242 9.184.118.200 # The IP address of the server and Load Balancer
ifconfig gre0 inet 9.184.114.25 # Some IP address on this subnet
### loopback...
ifconfig lo0 alias 9.184.114.24 netmask 255.255.255.255
```

- To set up a GRE tunnel, use the following:

```
sysctl -w net.ipv4.conf.all arp_ignore=3 net.ipv4.conf.all arp_announce=2

# for gre
modprobe ipgre
ip link set gre0 up
ip addr add <cluster_ip> scope host dev gre0
sysctl -w net.ipv4.conf.all rp_filter=0
sysctl -w net.ipv4.conf.gre0 rp_filter=0
```

To set up an IPIP tunnel, use the following:

```
sysctl -w net.ipv4.conf.all arp_ignore=3 net.ipv4.conf.all arp_announce=2
modprobe ipip
ip link set tunl0 up
ip addr add <cluster_ip> scope host dev tunl0
sysctl -w net.ipv4.conf.all rp_filter=0
sysctl -w net.ipv4.conf.tunl0 rp_filter=0
```

```
./sbin/ifconfig ip.tun0 plumb 9.184.114.25 netmask 255.255.255.255 up # Some free IP address on this subnet
./sbin/ifconfig ip.tun0 9.184.114.25 9.184.114.222 up # Some free IP address on this subnet
./sbin/ifconfig ip.tun0 up ./sbin/ifconfig ip.tun0 tsrc 9.184.112.183 tdst 9.184.118.203 # The IP address of this machine and that of the server.
### loopback...
ifconfig lo0:1 plumb 9.184.114.24 netmask 255.0.0.0 up
```

Note: You might need to disable reverse path filtering on the backend servers so encapsulation forwarding will function properly. Use the following command to disable reverse path filtering on linux systems:

```
sysctl -w net.ipv4.conf.all rp_filter=0
```

- Tunneling is not supported on Windows operating systems.

Related tasks:

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Quiesce servers or services for maintenance windows

To remove a server or service from the Load Balancer configuration for any reason (updates, upgrades, service, etc.), you can use the `dscontrol`(or `cbrcontrol`) manager quiesce command. You cannot use quiesce in the Site Selector component.

About this task

The quiesce subcommand allows existing connections to complete without being severed and disallows any new connections to a server or service. You can quiesce servers on a scheduled time to perform upgrades or general maintenance. The `daily` option specifies to quiesce the server at a time that you specify.

Note: When you want to bring down a particular service on a server, like HTTP, you can now quiesce individual daemons. If you have added `server1` to port 80 and port 22 of a cluster, for example, the usual quiesce command will stop forwarding traffic to `server1` on port 80 and port 22. This new option allows you to stop traffic only on port 80 if you need to stop that particular service.

Tip: You can use similar steps for quiescing the server for the CBR component.

Procedure

- Quiesce a server or individual service daemon immediately.

1. Quiesce a server. Use the following command:

```
dscontrol manager quiesce server
```

The following is an example of using the option to quiesce server 9.40.25.67:

```
dscontrol manager quiesce 9.40.25.67
```

2. Quiesce a service. Use the following command:

```
dscontrol manager quiesce cluster@port@server
```

The following is an example of using the option to quiesce the HTTP service on server 9.40.25.67, which is on cluster1:

```
dscontrol manager quiesce cluster1@80@9.40.25.67
```

If you use this command in conjunction with a command to quiesce a server, the commands will act together. For example:

- If you quiesce `server1` with the `dscontrol manager quiesce server` command, you can unquiesce a specific service by using the following command:

```
dscontrol manager unquiesce cluster@port@server
```

This command will unquiesce that particular port for service, but all other ports will remain quiesced.

- If you quiesce an individual service by using the `dscontrol manager quiesce cluster@port@server` command, you will unquiesce all services on the server if you use the following command:


```
dscontrol manager unquiesce server
```

This will ignore settings that you have configured for individual services.

- Quiesce a server on a daily schedule. Use the following command:

```
dscontrol manager quiesce server daily start_hour end_hour
```

Note:

- Start hour and end hour are values from 0 to 24. For example, (1 2) indicates to quiesce the server from 1:00 AM to 2:00 AM. (12 13) indicates to quiesce the server from 12:00 PM to 1:00 PM, which is 1 hour.
- Specify (-1 -1) to disable the daily quiesce for a particular server.

The following is an example of using the daily option to quiesce server 9.40.25.67 from 2:00 AM to 5:00 AM:

```
dscontrol manager quiesce 9.40.25.67 daily 2 5
```

Related concepts:

Chapter 6, “Administering Load Balancer,” on page 111

This section focuses on administering production environments and realistic test environments.

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Restricting incoming traffic with ipchains and iptables

Built into the Linux kernel is a firewall facility called ipchains. When Load Balancer and ipchains run concurrently, Load Balancer sees packets first, followed by ipchains. This allows the use of ipchains to harden a Linux Load Balancer machine, which could be, for example, a Load Balancer machine that is used to load balance firewalls.

About this task

In general, an appropriate ipchains strategy for the Load Balancer machines is to disallow all traffic, except that which is to or from the back-end servers, the partner high availability Load Balancer, any reach targets, or any configuration hosts.

It is not recommended to activate iptables when running Load Balancer on Linux kernel version 2.4.10.x. Activation on this Linux kernel version can result in performance degradation over time.

Procedure

- To activate iptables or ipchains, configure them to be completely restricted, so no inbound or outbound traffic permitted. The packet-forwarding portion of Load Balancer continues to function normally.

Some additional traffic must be permitted for all of Load Balancer to function properly. Some examples of this communication are:

- Advisors communicate between the Load Balancer machine and the back-end servers.
- Load Balancer pings back-end servers, reach targets, and high availability partner Load Balancer machines.

- User interfaces (graphical user interface, command line, and wizards) use RMI.
- Back-end servers must respond to pings from the Load Balancer machine.
- To deactivate iptables:
 1. List the modules which are using ip_tables and ip_conntrack. Issue the following command:


```
lsmod
```
 2. Remove them by issuing the following commands:


```
rmmod ip_tables
rmmod ip_conntrack
```

When you reboot the machine these modules will be added again, so you need to repeat these steps each time you reboot.

Related tasks:

“Quick start configuration” on page 61

This quick start example shows how to configure three locally attached workstations to load-balance web traffic between two web servers.

“Installing the product on distributed operating systems” on page 33

You use IBM Installation Manager to install this version of Websphere Load Balancer for IPv4 and IPv6.

Logging with Load Balancer

Load Balancer posts entries to a server log, a manager log, a metric monitor log (logging communications with Metric Server agents), and a log for each advisor you use.

About this task

You might want to configure logging for a variety of reasons, like troubleshooting or keeping records. Load Balancer can log activities to varying degrees of detail and keep the log files on your system. You can customize the following logging features:

- Setting the logging level to define the expansiveness of the messages written to the log. At level 0, errors are logged and Load Balancer also logs headers and records of events that happen only once (for example, a message about an advisor starting to be written to the manager log). Level 1 includes ongoing information, and so on, with level 5 including every message produced to aid in debugging a problem if necessary. The default for the manager, advisor, server, or subagent logs is 1.
- Setting the maximum size of a log. When the log file reaches the specified size, it is saved as <component>.log.full, and subsequent logs are written in the currently running file. In next cycle, the existing <component>.log.full is deleted Load Balancer will create a new one.

Important: You cannot set the log size to a value that is smaller than the current one. The higher you set the log level, the more carefully you should choose the log size. At level 0, it is probably safe to leave the log size to the default of 1MB; however, when logging at level 3 and above, you should limit the size without making it too small to be useful.

Log entries are time stamped so you can tell the order in which they were written.

Tip: You can use similar steps for logging for the other components of Load Balancer.

Procedure

- Display the current settings for the server log. Use the `dscontrol logstatus` command:
`dscontrol logstatus`
- Configure the logging level or maximum log size for a server log. Use the `dscontrol set` command:
`dscontrol set loglevel level logsize size`
where:
 - *level* is 0-5.
 - *size* is unlimited or a file size in bytes.
- Configure the logging level or maximum log size for a manager log. Use the `dscontrol manager` command:
`dscontrol manager loglevel level`
`dscontrol manager logsize size`
where:
 - *level* is 0-5.
 - *size* is unlimited or a file size in bytes.
- Configure the logging level or maximum log size for the metric monitor log that logs communication with Metric Server agents. Use the `dscontrol manager metric set` command:
`dscontrol manager metric set loglevel level`
`dscontrol manager metric set logsize size`
- Configure the logging level or maximum log size for an advisor log. Use the `dscontrol advisor` command:
`dscontrol advisor loglevel name cluster@port level`
`dscontrol advisor logsize name cluster@port size`
where:
 - *cluster@port* the cluster is the address in IP address format or symbolic name. The port is the number of the port that the advisor is monitoring. The cluster value is optional on the advisor commands, but the port value is required. If the cluster value is not specified, then the advisor will start running on the port for all clusters. If you specify a cluster, then the advisor will start running on the port, but only for the cluster you have specified. Read the Enabling advisors to manage load balancing topic for more information.

Related tasks:

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

Related reference:

“dscontrol logstatus” on page 239

Use this command to display the log settings for a server.

“dscontrol metric” on page 243

You can configure system metrics with the `dscontrol metric` command.

“dscontrol manager” on page 239

You can control the manager function with the `dscontrol manager` command.

Logging server statistics with binary logging

The binary logging feature allows server information to be stored in binary files. These files can then be processed to analyze the server information that has been gathered over time.

About this task

The following information is stored in the binary log for each server defined in the configuration.

- cluster address
- port number
- serverID
- server address
- server weight
- server total connections
- server active connections
- server port load
- server system load

Some of this information is retrieved from the executor as part of the manager cycle. Therefore the manager must be running in order for the information to be logged to the binary logs.

A sample Java program and command file have been provided in the `install_root/servers/samples/BinaryLog` directory. This sample shows how to retrieve all the information from the log files and print it to the screen. It can be customized to do any type of analysis you want with the data. An example using the supplied script and program to get a report of the Load Balancer's server information from 8:00 AM to 5:00 PM on May 1, 2001:

```
dslogreport 2001/05/01 8:00 2001/05/01 17:00
```

Tip: You can use similar steps for configuring binary logging for the CBR component. Binary logging is not available for the Site selector component.

Procedure

Use `dscontrol binlog` command set to configure binary logging

- Start binary logging:

```
dscontrol binlog start
```

The `start` option starts logging server information to binary logs in the `logs` directory. One log is created at the start of every hour with the date and time as the name of the file.

- Stop binary logging:

```
dscontrol binlog stop
```

The `stop` option stops logging server information to the binary logs. The log service is stopped by default.

- Set the interval value to control how often information is written to the logs.

```
dscontrol binlog interval seconds
```

The manager will send server information to the log server every manager interval. The information is written to the logs only if the specified log interval seconds have elapsed since the last record was written to the log. By default, the log interval is set to 60 seconds. There is some interaction between the settings of the manager interval and the log interval. Since the log server is provided with information no faster than manager interval seconds setting the log interval less than the manager interval effectively sets it to the same as the manager interval.

This logging technique allows you to capture server information at any granularity. You can capture all changes to server information that are seen by the manager for calculating server weights. However, this amount of information is probably not required to analyze server usage and trends. Logging server information every 60 seconds gives you snapshots of server information over time. Setting the log interval very low can generate huge amounts of data.

- Set the retention option to control how long log files are kept.

```
dscontrol binlog retention hours
```

Log files older than the retention *hours* specified are deleted by the log server. This will only occur if the log server is being called by the manager, so stopping the manager will cause old log files not to be deleted.

- View the current status for binary logging:

```
dscontrol binlog status
```

The `status` option returns the current settings of the log service. These settings are whether the service is started, what the interval is, and what the retention hours are.

Related tasks:

“Logging with Load Balancer” on page 158

Load Balancer posts entries to a server log, a manager log, a metric monitor log (logging communications with Metric Server agents), and a log for each advisor you use.

Related reference:

“dscontrol binlog” on page 231

You can control the settings and operation of the binary log file with the `dscontrol binlog` command.

Support for ICMP forwarding and messaging

Load Balancer now supports forwarding and processing ICMP messages to improve the robustness of connection protocols and permit Load Balancer to receive ICMP fragmentation messages.

Load Balancer will forward an ICMP message based on the following guidelines:

- For ICMP packets that contain headers with IP and TCP/UDP fragments, Load Balancer will forward to packets to the correct back-end server.
- For ICMP packets that do not contain TCP/UDP fragments, Load Balancer will forward the packets in a round robin manner.
- For ICMP messages that are for an IPGRE or an IPIP message that Load Balancer generated, Load Balancer will limit the outbound size appropriately for future packets.
- If Load Balancer forwards an IP packet, but the time to live (TTL) for the packet becomes zero when the TTL is decremented, Load Balancer will send an "ICMP Time Exceeded" message to the client.
- When Load Balancer cannot forward a packet, it will generate an ICMP message and send the appropriate message back to the client:
 - The outbound interface MTU is too small, or you need to use encapsulation.
 - Load Balancer cleaned up a connection record. For example, the cluster and port designation match, but the server is not present.

Related reference:

Chapter 6, “Administering Load Balancer,” on page 111

This section focuses on administering production environments and realistic test environments.

Configure rules to manage traffic to busy or unavailable servers

Use rules-based load balancing to fine tune when and why packets are sent to which servers. Load Balancer reviews any rules you add from first priority to last priority, stopping on the first rule that it finds to be true, then load balancing the traffic between any servers associated with the rule. It already balances the load based on the destination and port, but using rules expands your ability to distribute connections.

About this task

You should use rules-based load balancing with Dispatcher and Site Selector when you want to use a subset of your servers. Always use rules for the CBR component.

In most cases when configuring rules, you should configure a default always true rule in order to catch any request that is passed by other higher priority rules. This default can be a "Sorry, the site is currently down, try again later" response when all other servers fail for the client request.

All rules have a name, type, priority, and might have a begin range and end range, along with a set of servers. Rules are evaluated in priority order. A rule with a priority of 1 (lower number) is evaluated before a rule with a priority of 2 (higher number). The first rule that is satisfied will be used. When a rule has been satisfied, no further rules are evaluated. For a rule to be satisfied, it must meet two conditions:

1. The predicate of the rule must be true. That is, the value it is evaluating must be between the begin and end ranges, or the content must match the regular expression that is specified in the rule's pattern. For rules of type "true," the predicate is always satisfied, regardless of the begin and end ranges. If a rule has no servers that are associated with it, the rule only needs to meet this first condition to be satisfied. In this case, Load Balancer will drop the connection request.
2. If there are servers associated with the rule, at least one server must have a weight greater than 0 to forward packets so Load Balancer will have a server to which connections can be forwarded.

If a connection request does not satisfy any rules, Load Balancer will select a server from the full set of servers available on the port.

Choose from the following types of rules:

- Dispatcher
 - Active connections total
 - Always true
- CBR
 - Client IP address
 - Time of day
 - Connections per second
 - Active connections total
 - Always true
 - Content of a request
- Site Selector
 - Client IP address
 - Time of day
 - Metric all
 - Metric average
 - Always true

Procedure

- Configure a rule that is based on the total active connections. You may want to use rules based on active connections total on a port if your servers get overloaded and start throwing packets away. Certain Web servers will continue to accept connections even though they do not have enough threads to respond to the request. As a result, the client requests time out and the customer coming to your Web site is not served. You can use rules based on active connections to

balance capacity within a pool of servers. For example, you know from experience that your servers will stop serving after they have accepted 250 connections.

Note: The manager must be running for the rules to work.

Create a rule using the **dscontrol** or **cbrcontrol** rule command. You would then add to the rule your current servers plus some additional servers, which will otherwise be used for other processing. For example:

```
dscontrol rule add 130.40.52.153@80@pool2 type active beginrange 250 endrange 500
```

- Create a rule that always evaluates as true. Such a rule will always be selected, unless all the servers associated with it are down. Therefore, this rule should ordinarily be at a lower priority than other rules. You can even have multiple "always true" rules, each with a set of servers that are associated with it. Load Balancer will choose a rule based on the first rule that is true and has an available server.

For example, assume you have six servers. You want two of them to handle your traffic under all circumstances, unless they are both down. If the first two servers are down, you want a second set of servers to handle the traffic. If all four of these servers are down, then you will use the final two servers to handle the traffic. You could set up three "always true" rules, then the first set of servers will always be chosen as long as at least one is up. If both servers are down, one from the second set is chosen, and so forth.

As another example, you might want an "always true" rule to ensure that if incoming clients do not match any of the rules you have set, they will not be served. Then you would not add any servers to the rule, causing the clients packets to be dropped with no response. You can define more than one "always true" rule, and thereafter adjust which one gets run by changing their priority levels. Create a rule using the **dscontrol** rule command:

```
dscontrol rule add 130.40.52.153@80@jamais type true priority 100
```

You do not need to set a **beginrange** or **endrange** values when you create an always true rule.

- Add one or more servers to a rule set. You can use the **dscontrol rule useserver** command to add one or more servers to a rule set that is already defined. For example:

```
dscontrol rule useserver 130.40.52.153@80@jamais server1
```

```
dscontrol rule useserver 130.40.52.153@80@jamais server1+server2+server3
```

Additional configuration rules that you can add with CBR and Site Selector components

- Rules based on the client IP address. You might use rules based on the client IP address to screen the customers and allocate resources based on where they are coming from. For example:

```
cbrcontrol rule add 9.67.131.153@80@ni type ip beginrange 9.0.0.0 endrange 9.255.255.255
```

The **ni** rule screens out any connection from unwanted clients. You would then add to the rule the servers that you want accessible, or if you do not add any servers to the rule, requests coming from the **9.x.x.x** addresses are not served by any of your servers.

- Rules based on the time of day. You might use rules based on the time of day for capacity planning reasons. For example:

```
cbrcontrol rule add 9.67.131.153@80@ni type time beginrange 3 endrange 4
```


When defining the beginrange and endrange of time intervals, each value must be an integer representing only the hour portion of the time; portions of an hour are not specified. To specify a single hour, for example, the hour between 3:00 and 4:00 am, specify a beginrange of 3 and an endrange of 3. This range signifies all the minutes beginning with 3:00 and ending with 3:59. Specifying a beginrange of 3 and an endrange of 4 covers the two-hour period from 3:00 through 4:59.

- Rules based on the connections per second (CBR only).

Tip: The manager must be running for this scenario to work.

You might use rules based on connections per second to share some of your servers with other applications. For example, you can set two rules:

- If connections per second on port 80 is between 0 and 2000, then use these 2 servers
- If connections per second on port 80 is greater than 2000, then use these 10 servers

Setting rule evaluate option upserveronrule in conjunction with the connection type rule: When using the connections type rule and setting the upserveronrule option, if some of the servers in the server set are down, then you can ensure that the remaining servers will not overload. Read “Server evaluation option for rules” on page for more information.

- Metric all rule (Site Selector only). For the metric all rule, choose a system metric (cpuload, memload, or your own customized system metric script), and Site Selector compares the system metric value (returned by the Metric Server agent residing in each load-balanced server) with the begin and end range that you specify in the rule. Set the current system metric value for all the servers in the server within the range for the rule to run.

Tip: The system metric script you choose must reside on each of the load-balanced servers.

The following example describes how to add a metric all rule to your configuration:

```
sscontrol rule add dnsload.com:allrule1 type metricall metricname cpuload beginrange 0 endrange 100
```

- Metric average rule (Site Selector only). For the metric average rule, you choose a system metric (cpuload, memload, or your own customized system metric script), and Site Selector compares the system metric value (returned by the Metric Server agent residing in each load-balanced server) with the begin and end range that you specify in the rule. The average of the current system metric values for all the servers in the server set must be within the range for the rule to run.

Tip: The system metric script you choose must reside on each of the load-balanced servers.

The following example describes how to add a metric average rule to your configuration:

```
sscontrol rule add dnsload.com:avgrule1 type metricavg metricname cpuload beginrange 0 endrange 100
```

Related concepts:

Chapter 7, "Tuning Load Balancer," on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Related reference:

"dscontrol rule" on page 247

Control the executor function with the `dscontrol rule` command.

Optimize connections with client-to-server affinity

Use the affinity feature to map a client IP address to a back-end server.

Optimize connections with client-to-server affinity for the Dispatcher component

The Dispatcher affinity feature maps a client IP address to a back-end server. Affinity is established once a packet's destination IP address matches the cluster, the destination port matches the Load Balancer port, and the source IP address matches.

About this task

When affinity is established, subsequent packets are sent to the same back-end server. When affinity is broken, due to a server down or a server removal, all affinity and thus connections to that server are broken. Also, there is no "connection" information reported in the command line or GUI clients. Only the number of active affinity records are used.

This approach has the advantages of providing a hard affinity and of being more efficient for Load Balancer. The affinity method that is used decreases memory and CPU utilization as compared to connection forwarding.

gotcha: The removal of an affinity record also breaks connections, so when you migrate from Load Balancer for IPv4 to Load Balancer for IPv4 and IPv6, the value you have set for maximum staletimeout setting should be used as the value for stickytime setting for Load Balancer for IPv4 and IPv6.

Choose a selection algorithm based on how Load Balancer should route connections. You can choose the following selection algorithms:

affinity

Specifies that when a subsequent packet arrives from the same client as a previous packet, the request will be directed to the same server as the previous packet. Over time, the client will finish sending packets, and the affinity record will go away.

Each affinity record lives for the stickytime value, which is in seconds:

- If a subsequent connection is received within the interval of the stickytime value, the affinity record is still valid and the request will go to the same server.
- If a subsequent connection is not received within the interval for the stickytime value, the record is purged. Load Balancer will select a new server for a connection that is received after the stickytime value has expired.

You can use the `quiesce` command to take a server offline. The server is not taken down until after the `staletimeout` value expires. This is an example of the `quiesce` command:

```
dscontrol manager quiesce server
```

connection

Specifies that when a new TCP connection is received from a client, Load Balancer will choose the server at that moment in time, and forwards the packets to it. If a subsequent connection comes in from the same client, Load Balancer treats this connection as an unrelated, new connection and choose an appropriate server at that moment.

connection+affinity

Specifies that when a packet arrives from an existing connection, it is forwarded to the server which was chosen for that connection. If the packet does not belong to an existing connection, but the packet arrives from the same client, then Load Balancer selects the same server that was chosen for the previous connection.

The behaviour for the `connection+affinity` algorithm is similar to the `affinity` algorithm, but the timeout value for `affinity` is calculated from start of every connection. That means that if the time between the start of two connections from the same client is less than the `stickytime` value, then both connections will be forwarded to same server.

Procedure

- Optional: Enable affinity by adding a port and setting the selection algorithm and `stickytime` at the port level to some number of seconds using the `dscontrol port` command.

Note: There are now three selections you can choose from for selection algorithm:

- **affinity:** specifies that the server selection is based on client affinity.
- **connection:** specifies that the server selection is based on weighted round-robin selection (default).
- **conn+affin:** specifies that server selection is based on the relationship to an existing connection. For new connections, the server selection is based on affinity.

For example, use the following command to set the selection algorithm to `affinity` and `stickytime` to 60 seconds:

```
dscontrol port add cluster@port selectionalgorithm affinity stickytime 60
```

Note: When you enable affinity, it cannot be disabled unless you completely remove the port and add it again without affinity configured.

- Optional: Enable cross-port affinity. Cross port affinity is the affinity feature that has been expanded to cover multiple ports. For example, if a client request is first received on one port and the next request is received on another port, cross port affinity allows Load Balancer to send the client request to the same server. To use this feature, the ports must:
 - Share the same cluster address.
 - Share the same servers.
 - Use the same selection algorithm, which must be `affinity` or `conn+affin`.
 - Have the same `stickytime` value, which is not zero. After cross port affinity has been established, you have the flexibility to modify the `stickytime` value

for the port. However, it is recommended that you change the stickytime values for all shared ports to the same value, otherwise results might occur that are not expected.

More than one port can link to the same crossport. When subsequent connections arrive from the same client on the same port or a shared port, the same server will be accessed. Configure the crossport value with the dscontrol port add command. The following is an example of configuring multiple ports with:

- a selection algorithm of conn+affin
- a stickytime value of 60 seconds
- cross port affinity set to port 10

Enter the following:

```
dscontrol port add cluster@20 selectionalgorithm conn+affin stickytime 60 crossport 10
dscontrol port add cluster@30 selectionalgorithm conn+affin stickytime 60 crossport 10
dscontrol port add cluster@40 selectionalgorithm conn+affin stickytime 60 crossport 10
```

Note: You can only specify a value for crossport with the dscontrol port add command and cannot be modified afterwards. You cannot use the dscontrol port set command to configure the crossport value.

See dscontrol port for detailed information on command syntax for the crossport option.

- Optional: Enable port affinity override for a specific server. You can override the stickiness of a port for a specific server with this feature, but you must be using the conn+affinity selection algorithm. The port affinity override feature can be useful in certain situations, like if you are using rules. For example, assume that you:
 - Use a rule to limit the amount of connections to each application server
 - Want to establish an overflow server with an always true rule. This rule responds to incoming connections with a message like "Try again later" for that application.
 - Do not want the client to be sticky to that server, because the port has a stickytime value of 25 minutes.

With port affinity override, you can change the overflow server to override the affinity that would normally be associated with that port. The next time the client requests the cluster, it is load balanced to the best available application server and not the overflow server.

To enable port affinity override, enter the following command:

```
dscontrol server set cluster@port@server sticky no
```

Related tasks:

“Quiesce servers or services for maintenance windows” on page 156
To remove a server or service from the Load Balancer configuration for any reason (updates, upgrades, service, etc.), you can use the dscontrol(or cbrcontrol) manager quiesce command. You cannot use quiesce in the Site Selector component.

Related reference:

“dscontrol port” on page 244

Configure ports and port settings with the dscontrol port command.

“dscontrol server” on page 249

Configure servers and modify existing server configurations with the dscontrol server command.

Optimize connections with client-to-server affinity for the Content Based Router component

The Content Based Router (CBR) affinity feature maps a client IP address to a back-end server. Affinity is established once a packet's destination IP address matches the cluster, the destination port matches the Load Balancer port, and the source IP address matches.

About this task

When affinity is established, subsequent packets are sent to the same back-end server. When affinity is broken, due to a server down or a server removal, all affinity and thus connections to that server are broken. Also, there is no connection information reported in the command line or GUI clients. Only the number of active affinity records are used.

This approach has the advantages of providing a hard affinity and of being more efficient for Load Balancer. The affinity method that is used decreases memory and CPU utilization as compared to connection forwarding.

gotcha: The removal of an affinity record also breaks connections, so when you migrate from Load Balancer for IPv4 to Load Balancer for IPv4 and IPv6, the value you have set for maximum stalemtimeout setting should be used as the value for stickytime setting for Load Balancer for IPv4 and IPv6.

Specify affinity for the CBR component as active cookie, passive cookie and URI, as follows:

Active cookie

Enables load-balancing Web traffic with affinity to the same server based upon cookies generated by Load Balancer. When a rule has been enabled for active cookie affinity, new client requests are load-balanced using standard CBR algorithms, while succeeding requests from the same client are sent to the initially chosen server. The chosen server is stored as a cookie in the response to the client. As long as the client's future requests contains the cookie, and each request arrives within the stickytime interval, the client will maintain affinity with the initial server.

Making a rule sticky would normally be used for CGI or servlets that store client state on the server. The state is identified by a cookie ID (these are server cookies). Client state is only on the selected server, so the client needs the cookie from that server to maintain that state between requests.

Passive cookie

Enables load-balancing Web traffic with affinity to the same server based

upon self-identifying cookies generated by the servers. When passive cookie affinity is enabled, Load Balancer will choose the server based on the cookie name in the HTTP header of the client request.

The first time Load Balancer finds a server whose cookie value contains the client's cookie name, Load Balancer chooses that server for the request. If the cookie name in the client request is not found or does not match any of the content within the servers' cookie values, the server is chosen using existing server selection or the weighted round-robin technique.

- URI** URI affinity allows you to load-balance web traffic to Caching Proxy servers which allow unique content to be cached on each individual server. As a result, you will effectively increase the capacity of your site's cache by eliminating redundant caching of content on multiple machines. This is useful in scenarios wherein your Web site supports a moderate volume of client traffic to very diverse content and you prefer to have a larger cache spread across multiple servers.

The default for the affinity option is none.

Procedure

- Configure active cookie affinity. For example, set the affinity parameter to `activecookie`

```
Cbrcontrol rule set cluster@port@rule stickytime 60
Cbrcontrol rule set cluster@port@rule affinity activecookie
```

The active cookie affinity option, for the rule command, can only be set to `activecookie` if `port stickytime` is zero (not enabled). When active cookie affinity is active on a rule then you cannot enable `stickytime` on the port.

Active cookie affinity has a default expiration of the current server time, plus the `stickytime` interval, plus twenty-four hours. If your clients (those sending requests to your CBR machine) have inaccurate times on their system (for example, they are more than one day ahead of the server time), then those clients' systems will ignore the cookies from CBR because the system will assume that the cookies have already expired. To set a longer expiration time, modify the `cbrserver` script. In the script file, edit the `javaw` line, adding the following parameter after `LB_SERVER_KEYS:-DCOOKIEEXPIREINTERVAL=X` where `X` is the number of days to add to the expiration time.

- Configure passive cookie affinity:
 - Set the affinity parameter to `passivecookie` on the `cbrcontrol` rule `[add | set]` command. Also, the `cookieName` parameter must be set to the name of the cookie that Load Balancer should look for in the client HTTP header request.
 - Set the `cookieValue` parameter, for each server in the rule's server set, on the `cbrcontrol server [add | set]` command.

You can set the passive cookie affinity option for the rule command only to `passivecookie` if `port stickytime` is zero (not enabled). When passive cookie affinity is active on a rule, you cannot enable `stickytime` on the port.

- Configure URI affinity:
 - Set the affinity parameter to `uri` on the `cbrcontrol` rule `[add | set]` command. Set the URI affinity option `URI` only if `port stickytime` is zero (not enabled). When URI affinity is active on a rule you cannot enable `stickytime` on the port.

Related tasks:

“Quiesce servers or services for maintenance windows” on page 156

To remove a server or service from the Load Balancer configuration for any reason (updates, upgrades, service, etc.), you can use the dscontrol(or cbrcontrol) manager quiesce command. You cannot use quiesce in the Site Selector component.

Related reference:

“dscontrol port” on page 244

Configure ports and port settings with the dscontrol port command.

“dscontrol server” on page 249

Configure servers and modify existing server configurations with the dscontrol server command.

Optimize connections with client-to-server affinity for the Site Selector component

The Site Selector affinity feature maps a client IP address to a back-end server. Affinity is established once a packet's destination IP address matches the cluster, the destination port matches the Load Balancer port, and the source IP address matches.

About this task

When affinity is established, subsequent packets are sent to the same back-end server. When affinity is broken, due to a server down or a server removal, all affinity and thus connections to that server are broken. Also, there is no "connection" information reported in the command line or GUI clients. Only the number of active affinity records are used.

This approach has the advantages of providing a hard affinity and of being more efficient for Load Balancer. The affinity method that is used decreases memory and CPU utilization as compared to connection forwarding.

gotcha: The removal of an affinity record also breaks connections, so when you migrate from Load Balancer for IPv4 to Load Balancer for IPv4 and IPv6, the value you have set for maximum stalemtimeout setting should be used as the value for stickytime setting for Load Balancer for IPv4 and IPv6.

Choose a selection algorithm based on how Load Balancer should route connections. You can choose the following selection algorithms:

affinity

Specifies that when a subsequent packet arrives from the same client as a previous packet, the request will be directed to the same server as the previous packet. Over time, the client will finish sending packets, and the affinity record will go away.

Each affinity record lives for the stickytime value, which is in seconds:

- If a subsequent connection is received within the interval of the stickytime value, the affinity record is still valid and the request will go to the same server.
- If a subsequent connection is not received within the interval for the stickytime value, the record is purged. Load Balancer will select a new server for a connection that is received after the stickytime value has expired.

You can use the `quiesce` command to take a server offline. The server is not taken down until after the `staletimeout` value expires. This is an example of the `quiesce` command:

```
dscontrol manager quiesce server
```

connection

Specifies that when a new TCP connection is received from a client, Load Balancer will choose the server at that moment in time, and forwards the packets to it. If a subsequent connection comes in from the same client, Load Balancer treats this connection as an unrelated, new connection and choose an appropriate server at that moment.

connection+affinity

Specifies that when a packet arrives from an existing connection, it is forwarded to the server which was chosen for that connection. If the packet does not belong to an existing connection, but the packet arrives from the same client, then Load Balancer selects the same server that was chosen for the previous connection.

The behaviour for the `connection+affinity` algorithm is similar to the `affinity` algorithm, but the timeout value for `affinity` is calculated from start of every connection. That means that if the time between the start of two connections from the same client is less than the `stickytime` value, then both connections will be forwarded to same server.

Procedure

- Optional: Enable affinity by adding a port and setting the selection algorithm and `stickytime` at the port level to some number of seconds using the `dscontrol port` command.

Note: There are now three selections you can choose from for selection algorithm:

- **affinity:** specifies that the server selection is based on client affinity.
- **connection:** specifies that the server selection is based on weighted round-robin selection (default).
- **conn+affin:** specifies that server selection is based on the relationship to an existing connection. For new connections, the server selection is based on affinity.

For example, use the following command to set the selection algorithm to `affinity` and `stickytime` to 60 seconds:

```
dscontrol port add cluster@port selectionalgorithm affinity stickytime 60
```

Note: When you enable affinity, it cannot be disabled unless you completely remove the port and add it again without affinity configured.

- Optional: Enable cross-port affinity. Cross port affinity is the affinity feature that has been expanded to cover multiple ports. For example, if a client request is first received on one port and the next request is received on another port, cross port affinity allows Load Balancer to send the client request to the same server. To use this feature, the ports must:
 - Share the same cluster address.
 - Share the same servers.
 - Use the same selection algorithm, which must be `affinity` or `conn+affin`.
 - Have the same `stickytime` value, which is not zero. After cross port affinity has been established, you have the flexibility to modify the `stickytime` value

for the port. However, it is recommended that you change the stickytime values for all shared ports to the same value, otherwise results might occur that are not expected.

More than one port can link to the same crossport. When subsequent connections arrive from the same client on the same port or a shared port, the same server will be accessed. Configure the crossport value with the dscontrol port add command. The following is an example of configuring multiple ports with:

- a selection algorithm of conn+affin
- a stickytime value of 60 seconds
- cross port affinity set to port 10

Enter the following:

```
dscontrol port add cluster@20 selectionalgorithm conn+affin stickytime 60 crossport 10
dscontrol port add cluster@30 selectionalgorithm conn+affin stickytime 60 crossport 10
dscontrol port add cluster@40 selectionalgorithm conn+affin stickytime 60 crossport 10
```

Note: You can only specify a value for crossport with the dscontrol port add command and cannot be modified afterwards. You cannot use the dscontrol port set command to configure the crossport value.

See dscontrol port for detailed information on command syntax for the crossport option.

- Optional: Enable port affinity override for a specific server. You can override the stickiness of a port for a specific server with this feature, but you must be using the conn+affinity selection algorithm. The port affinity override feature can be useful in certain situations, like if you are using rules. For example, assume that you:
 - Use a rule to limit the amount of connections to each application server
 - Want to establish an overflow server with an always true rule. This rule responds to incoming connections with a message like "Try again later" for that application.
 - Do not want the client to be sticky to that server, because the port has a stickytime value of 25 minutes.

With port affinity override, you can change the overflow server to override the affinity that would normally be associated with that port. The next time the client requests the cluster, it is load balanced to the best available application server and not the overflow server.

To enable port affinity override, enter the following command:

```
dscontrol server set cluster@port@server sticky no
```

Related tasks:

“Quiesce servers or services for maintenance windows” on page 156

To remove a server or service from the Load Balancer configuration for any reason (updates, upgrades, service, etc.), you can use the `dscontrol`(or `cbrcontrol`) manager quiesce command. You cannot use quiesce in the Site Selector component.

Related reference:

“`dscontrol port`” on page 244

Configure ports and port settings with the `dscontrol port` command.

“`dscontrol server`” on page 249

Configure servers and modify existing server configurations with the `dscontrol server` command.

Sample scripts to generate alerts and record server failure

Load Balancer provides user exits that trigger scripts that you can customize. You can create the scripts to perform automated actions, such as alerting an Administrator when servers are marked down by the manager or simply record the event of the failure.

Sample scripts, which you can customize, are in the `install_root/servers/samples` directory. In order to run the files, you must move them to the `install_root/servers/bin` directory and remove the `△sample△` file extension. The following sample scripts are provided:

- **serverDown** — a server is marked down by the manager.
- **serverUp** — a server is marked back up by the manager.
- **managerAlert** — all servers are marked down for a particular port.
- **managerClear** — at least one server is now up, after all were marked down for a particular port.

If all servers on a cluster are marked down (either by the user or by the advisors), the `managerAlert` (if configured) starts, and Load Balancer attempts to route traffic to the servers using a round-robin technique. The `serverDown` script does not start when the last server in the cluster is detected as offline. By design, Load Balancer attempts to continue to route the traffic in case a server comes back online and responds to the request. If Load Balancer instead dropped all traffic, the client would receive no response. When Load Balancer detects that the first server of a cluster is back online, the `managerClear` script (if configured) starts, but the `serverUp` script (if configured) is not run until an additional server is brought back online.

Here are some considerations for using the `serverUp` and `serverDown` scripts:

- If you define the manager cycle to be less than 25% of the advisor time, false reports of servers up or down can result. By default, the manager runs every 2 seconds, but the advisor runs every 7 seconds. Therefore, the manager expects new advisor information within 4 cycles. However, removing this restriction (that is, defining the manager cycle to be greater than 25% of the advisor time) significantly decreases performance because multiple advisors can advise on a single server.
- When a server goes down, the `serverDown` script starts. However, if you issue a `serverUp` command, it is assumed that the server is up until the manager obtains new information from the advisor cycle. If the server is still down, the `serverDown` script runs again.

Chapter 7. Tuning Load Balancer

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Attention: If you are not using the Dispatcher component, substitute **dscontrol** with the following commands:

- Content Based Routing: **cbrcontrol**
- Site Selector: **sscontrol**

The manager report

The manager function of Load Balancer calculates a weight for each server. These weights are used to determine how many connections a server should receive as compared with the other servers in the same cluster and port configuration. Understanding the manager report is critical to understanding how the network traffic is distributed.

The manager report contains a list of each cluster, port, and server that is defined to that cluster:port combination. Each server shows two weights, now and new, and four columns that are used to calculate the weight:

- Active connections (ACTV) or CPU (in the case of Site Selector component based on input from the Metric Server)
- New connections (NEWC) or Memory (in the case of Site Selector component based on input from the Metric Server)
- Port Load (PORT)
- System load (SYS)

Each of the four columns is assigned a percentage that is used to calculate the weight for the server. The percentages are set with the cluster set proportion command. By default, only the active connections and new connections are considered when calculating the weight of the server in the case of Dispatcher and CBR. These values are based on information that is generated and stored internally in the executor. For Site Selector, the manager obtains the first two values (CPU and memory) from Metric Server.

When an advisor is started, the proportion for the port load is set to 1% so that the port load is used in the weight calculation. Similarly, when a metric is added the proportion for the system load is set to 1%. The manager function returns the following values for each server:

Active connections (ACTV)

Active connections are TCP connections that are closed at the start of the manager cycle.

New connections (NEWC)

New connections represent the increase in total connections from the start of the manager cycle to the start of the last manager cycle.

Port Load (PORT)

The port load is the value that is obtained from an advisor that is defined

on this cluster:port combination. If an advisor is not started, the port load is always zero. When an advisor is defined, the port load typically represents the number of milliseconds for the advisor to receive a response from the server.

When the port load is shown as -1, the advisor did not receive a successful response to its query. Increase the log level and log size for the advisor to investigate why the server did not respond. If the server never responds to the connection request, complete the following steps:

1. Ensure that you can successfully ping the server from the Load Balancer machine.
2. Verify that the server application is started and listening on the port that is defined. The server should be listening on the wildcard address (0.0.0.0), or both the cluster IP address and the real server IP address to successfully respond to the advisor requests.

If the server responds to the connection, then it might be responding to the query in a manner that is different from what Load Balancer is expecting to see. Check the advisorresponse string that is defined to ensure it matches what the server has transmitted. This scenario applies to both http and https advisors.

System load (SYS)

The system load represents the value that is returned from the metric server. If metrics have not been added for this cluster:port combination, the system load is zero (0). When metrics are defined, the system load is a value in between -1 and 100, which represents the status of the server. 100 is very busy and zero (0) is idle.

If the system load shows -1, Load Balancer cannot communicate with the metric server on the back-end machine. Ensure Load Balancer keys are properly distributed to the server, that the server can be pinged from the Load Balancer, and that metric server is started on the machine. If the problem persists, complete the following steps:

1. Edit the script for the metric server on the back-end machine and increase both the log level and log size.
2. Restart the metric server.
3. Increase the the log size and the log level for the metric monitor at the Load Balancer.
4. Examine the log files on both the Load Balancer machine and the back-end machine to determine why communication is failing.

CPU and Memory (Site Selector component)

The CPU and Memory represent the values that are returned from the metric server . Cpload and memload metrics are added automatically for the Site Selector configuration. CPU and Memory are values in between -1 and 100, which represent the status of the server. 100 is very busy and zero (0) is idle. If these metrics show -1, Load Balancer cannot communicate with the metric server on the back-end machine. Ensure Load Balancer keys are properly distributed to the server, that the server can be pinged from the Load Balancer, and that metric server is started on the machine. If the problem persists, complete the following steps:

1. Edit the script for the metric server on the back-end machine and increase both the log level and log size.
2. Restart the metric server.

3. Increase the log size and the log level for the metric monitor at the Load Balancer.
4. Examine the log files on both the Load Balancer machine and the back-end machine to determine why communication is failing.

The number of active connections and new connections are determined based upon the number of connections that the executor has forwarded within the last cycle of the last manager function. The manager cycle is two seconds, by default.

Configuring server weights

Under normal circumstances, Load Balancer uses all of the values that have proportions that are not zero to calculate the new weight. For example, if the proportions are 40 40 20 0, the active connections and new connections are 40% of the weight calculation each and the port load is 20%.

As an example, assume the manager function returns the following values:

	ACTV	NEWC	PORT	SYS
Server1	50	200	25	0
Server2	25	100	50	0

The initial weight calculations will be:

- Server1 = $.40(50) + .40(200) + .2(25) = 20 + 80 + 5 = 105$
- Server2 = $.40(25) + .40(100) + .2(50) = 10 + 40 + 10 = 60$

The initial weights are scaled to be proportional to the weightbound for the cluster:port. By default, the weightbound is 10. Thus, in the previous example, the final weights, which are rounded to the nearest whole number, are:

- Server1 = $(105/165) * 10 = 6$
- Server2 = $(60/165) * 10 = 4$

The calculated weight is shown as the NEW weight in the manager report. The weight is only pushed to the executor function if it exceeds the sensitivity level that is configured for the cluster:port combination. The NOW weight represents the weight that is obtained from the executor at the start of this manager cycle.

If the port load or the system load is -1, and the respective proportion for the port or system column is greater than 0, the calculated weight is zero (0). Zero (0) indicates that the server is not active and new requests are not sent to the server.

If you quiesce a server, you will see that the weight is also shown as zero (0), but the port load is positive if the server is still online. If a quiesced server goes offline, the port load is -1.

If a user issues a "server down" function on a server to prevent Load Balancer from sending requests to that server, the weight is -1 regardless of the value for the port load and system load.

Related tasks:

Tuning

Optimizing the manager interval

To optimize overall performance, the manager is restricted in how often it can interact with the executor. You can make changes to this interval by entering the dscontrol manager interval and dscontrol manager refresh commands.

About this task

The manager interval specifies how often the manager will update the server weights that the executor uses in routing connections. If the manager interval is too low, it can mean poor performance as a result of the manager constantly interrupting the executor. If the manager interval is too high, it can mean that the executor's request routing will not be based on accurate, up-to-date information.

The manager refresh cycle specifies how often the manager will ask the executor for status information. The refresh cycle is based on the interval time.

Procedure

1. Set the manager interval time, in seconds. For example, to set the manager interval to 1 second, enter the following command:

```
dscontrol manager interval 1
```
2. Set the manager refresh time, in seconds. For example, to set the manager refresh cycle to 3, enter the following command:

```
dscontrol manager refresh 3
```

This will cause the manager to wait for 3 intervals before asking the executor for status.

Related tasks:

Chapter 7, "Tuning Load Balancer," on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Tuning the proportion of importance given to status information

The manager uses ratios to determine the importance of status information coming from advisors and Load Balancer. You can change the default ratios that the manager uses to weight this information.

About this task

The manager can use some or all of the following external factors in its weighting decisions:

- Active connections: The number of active connections on each load balanced server machine (as tracked by the executor), or
- CPU (in the case of the Site Selector component, based on input from the Metric Server)
- New connections: The number of new connections on each load balanced server machine (as tracked by the executor), or
- Memory (in the case of the Site Selector component, based on input from the Metric Server)
- Port-specific: The input from advisors listening on the port.
- System metric: The input from the system monitoring tools, such as Metric Server or WLM.

Along with the current weight for each server and some other information required for its calculations, the manager gets the first two values (active and new

connections) from the executor. These values are based on information that is generated and stored internally in the executor.

You can change the relative proportion of importance of the four values on a per cluster basis. Think of the proportions as percentages; the sum of the relative proportions must equal 100%. The default ratio is 50/50/0/0, which ignores the advisor and system information. In your environment, you may need to try different proportions to find the combination that gives the best performance.

Note:

- When adding an advisor (other than WLM), if the port proportion is zero, then the manager increases this value to 1. Because the sum of the relative proportions must total 100, the highest value is then decreased by 1.
- When adding the WLM advisor, if the system metric proportion is zero, then the manager increases this value to 1. Because the sum of the relative proportions must total 100, the highest value is then decreased by 1.

The number of active connections is dependent upon the number of clients as well as the length of time necessary to use the services that are being provided by the load balanced server machines. If the client connections are quick (such as small Web pages served using HTTP GET), then the number of active connections are fairly low. If the client connections are slower (such as a database query), then the number of active connections are higher.

You should avoid setting active and new connections proportions values too low. You will disable load balancing and smoothing unless you have these first two values set to at least 20 each.

Procedure

To set the proportion of importance that is given to the different factors, use the “dscontrol cluster” on page 231 command. For example:

```
dscontrol cluster set cluster proportions value
```

See the topic on the “dscontrol cluster” on page 231 command for more information.

Managing traffic with server weights

Weights are applied to all servers on a port. For any particular port, the requests are distributed between servers based on their weights relative to each other. For example, if one server is set to a weight of 10, and the other to 5, the server set to 10 should get twice as many requests as the server set to 5.

About this task

Weights are set by the manager function based upon internal counters in the executor, feedback from the advisors, and feedback from a system-monitoring program, such as Metric Server. If you want to set weights manually while running the manager, specify the fixedweight option on the dscontrol server command. For a description of the fixedweight option, see “dscontrol manager” on page 239.

The maximum weight boundary affects how much difference there can be between the number of requests each server will get. If you set the maximum weightbound to 1, then all the servers can have a weight of 1, 0 if quiesced, or -1 if marked

down. As you increase this number, the difference in how servers can be weighted is increased. At a maximum weightbound of 2, one server could get twice as many requests as another. At a maximum weightbound of 10, one server could get 10 times as many requests as another. The default maximum weightbound is 20.

If an advisor finds that a server has gone down, it tells the manager, which sets the weight for the server to zero. As a result, the executor will not send any additional connections to that server as long as that weight remains zero. If there were any active connections to that server before the weight changed, they will be left to complete normally.

If all the servers are down, the manager sets the weights to half the weightbound.

Procedure

- To specify the maximum weight boundary that any server can have, use the following command:

```
dscontrol port set port weightbound weight
```

- Configure fixed weights for servers.

1. Turn on the fixedweight option. For more information, see “dscontrol server” on page 249.

```
dscontrol server set cluster@port@server fixedweight yes
```

The server weight value remains fixed while the manager is running until you issue another dscontrol server command with fixedweight set to no.

2. After fixedweight is set to yes, use the dscontrol server set weight command to set the weight to the value you desire. For example:

```
dscontrol server set cluster@port@server weight value
```

Related concepts:

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Related tasks:

“Reset connections to down servers”

The TCP reset feature sends a TCP reset to clients that are connected to servers that have weights of 0. A server's weight can be 0 if it is configured to 0 or if an advisor marks the server down.

Related reference:

“dscontrol manager” on page 239

You can control the manager function with the dscontrol manager command.

“dscontrol port” on page 244

Configure ports and port settings with the dscontrol port command.

Reset connections to down servers

The TCP reset feature sends a TCP reset to clients that are connected to servers that have weights of 0. A server's weight can be 0 if it is configured to 0 or if an advisor marks the server down.

About this task

This feature is useful for long-lived connections where it hastens the client's ability to renegotiate a failed connection. A TCP reset will cause the connection to be immediately closed. When you activate the TCP reset feature, Load Balancer will send a TCP reset to the client when the client has a connection to a server with a weight of 0.

Procedure

Activate the TCP reset to reset connections to down servers.

Note: To use the TCP reset feature, you must set the `clientgateway` for IPv4 addresses, or `clientgateway6` for IPv6 addresses, to a router address. Use the `dscontrol` executor command. `Clientgateway` has to be set for a port that is on IPv4 cluster, and `clientgateway6` has to be set on a port that is on IPv6 cluster. To activate TCP reset, use the following command:

```
dscontrol port add|set cluster@port@server reset yes
```

The default value for the reset feature is no.

Related reference:

“`dscontrol manager`” on page 239

You can control the manager function with the `dscontrol manager` command.

“`dscontrol port`” on page 244

Configure ports and port settings with the `dscontrol port` command.

Update the network configuration automatically

Load Balancer requires information about all interfaces and IP addresses configured on a machine, and this information can be obtained manually or configured to be automatically updated.

About this task

Attention: This feature is not available with the CBR and Site Selector components.

For Load Balancer to work properly, it must have current information for the network interfaces and connections on the Load Balancer machine. There are two methods for obtaining these updates:

Manual mode

This is the default mode, and you are required to update Load Balancer when there are network configuration changes.

Auto mode

Specifies that Load Balancer detects the configuration for the Network Identification Card (NIC) automatically and updates the settings every 10 seconds.

Procedure

- Enable automatic updating. Use the following command:

```
dscontrol executor nicconfig auto
```
- Enable manual updating, and update the configuration.
 1. Enable manual updates. Use the following command:

```
dscontrol executor nicconfig manual
```

2. Update the network configuration. Use the following command:

```
dscontrol executor nicconfig update
```

Related tasks:

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Related reference:

“dscontrol executor” on page 233

Control the executor function with the dscontrol executor command.

Optimizing the sensitivity threshold

To work at top speed, updates to the weights for the servers are only made if the weights have changed significantly. Constantly updating the weights when there is little or no change in the server status could create unnecessary overhead.

About this task

When the percentage weight change for the total weight for all servers on a port is greater than the sensitivity threshold, the manager updates the weights used by the executor to distribute connections. Consider, for example, that the total weight changes from 100 to 105. The change is 5%. With the default sensitivity threshold of 5, the manager will not update the weights used by the executor, because the percentage change is not above the threshold. If, however, the total weight changes from 100 to 106, the manager will update the weights.

Note: In most cases, you will not need to change this value.

Procedure

To set the manager’s sensitivity threshold to a value other than the default (for example, 6), enter the following command:

```
dscontrol manager sensitivity 6
```

Related information:

Chapter 7, “Tuning Load Balancer,” on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Optimizing the smoothing index

The smoothing index limits the amount that a server’s weight can change, effectively smoothing the change in the distribution of requests.

About this task

The manager calculates the server weights dynamically. As a result, an updated weight can be very different from the previous one. Under most circumstances, this will not be a problem. Occasionally, however, it may cause an oscillating effect in the way the requests are load balanced. For example, one server can end up

receiving most of the requests due to a high weight. The manager will see that the server has a high number of active connections and that the server is responding slowly. It will then shift the weight over to the free servers and the same effect will occur there too, creating an inefficient use of resources.

To alleviate this problem, the manager uses a smoothing index. A higher smoothing index will cause the server weights to change less drastically. A lower index will cause the server weights to change more drastically. The default value for the smoothing index is 1.5. At 1.5, the server weights can be rather dynamic. An index of 4 or 5 will cause the weights to be more stable.

Note: In most cases, you will not need to change this value.

Procedure

Set the smoothing index, in seconds. For example, to set the smoothing index to 4 enter the following command:

```
dscontrol manager smoothing 4
```

Related concepts:

Chapter 7, "Tuning Load Balancer," on page 175

How well a Web site performs while receiving heavy user traffic is an essential factor in the overall success of an organization. This topic highlights a few main ways you can improve performance through a combination of product features and application development considerations.

Controlling connection records with the `staletimeout` value

Connections are considered stale when there has been no activity on that connection for the number of seconds specified in `stale timeout`. When the number of seconds has been exceeded with no activity, Load Balancer will remove that connection record from its tables, and subsequent traffic for that connection is discarded. The `staletimeout` command controls the way Load Balancer handles idle connections and the associated connection records.

About this task

Use the `staletimeout` command to control the period during which Load Balancer should keep connections in the "Established" state and accept traffic when no active traffic has been seen in the Dispatcher tables.

A client sends a FIN packet after it has sent all its packets so that the server will know that the transaction is finished. When Dispatcher receives the FIN packet, it marks the transaction from active state to FIN state. When a transaction is marked FIN, the memory that is reserved for the connection can be cleared.

Procedure

To change the `staletimeout` value, use the `dscontrol` executor `set` command. Type the following at a command prompt:

```
dscontrol executor set staletimeout time
```

where the value for *time* is in seconds.

Note: Some services might have `staletimeout` values of their own.

Note: For example, LDAP (Lightweight Directory Access Protocol) has a configuration parameter called `idletimeout`. When `idletimeout` seconds have been exceeded, an idle client connection will be forcibly closed. `idletimeout` may also be set to 0, which means that the connection will never be forcibly closed.

Connectivity problems can occur when Load Balancer's stale timeout value is smaller than the service's timeout value. In the case of LDAP, for example, if the Load Balancer `staletimeout` value defaults to 6400 seconds, and there is no activity on the connection for 6400 seconds, Load Balancer will remove the connection record from its tables. Furthermore, if the `idletimeout` value at the LDAP server is larger than 6400 seconds (or set to 0), the client might still believe that it has a connection to the server. When the client sends packets, the packets will be discarded by Load Balancer. This causes the LDAP client to hang when a request is made to the server.

To avoid this problem, set the LDAP `idletimeout` to a nonzero value that is the same or smaller than the Load Balancer `staletimeout` value.

Related reference:

Chapter 6, "Administering Load Balancer," on page 111

This section focuses on administering production environments and realistic test environments.

Chapter 8. Troubleshooting Load Balancer components

Use the information that is provided to help you solve problems that can occur in Load Balancer.

Troubleshooting Content Based Router

Use the information that is provided to help you solve problems that can occur in Content Based Routing.

Click a link in the table to go to a full description and possible solution for the problem that you are experiencing.

Table 10. CBR Troubleshooting table

Symptom	Possible cause
CBR not running correctly	Conflicting port numbers
The cbrcontrol or lbadm command fails with 'Server not responding' or 'unable to access RMI server' message	Commands fail due to socksified stack. Or commands fail due to not starting cbrserver
Requests are not being load balanced	Caching Proxy was started before the executor was started
On Solaris, the cbrcontrol executor start command fails with 'Error: Executor was not started.' message	Command fails because the system IPC defaults may need to be modified, or link to library is incorrect.
URL rule does not work	Syntactical or configuration error
Unexpected GUI behavior when using Windows systems paired with Matrox AGP video card	Problem occurs when using Matrox AGP video cards while running the Load Balancer GUI
GUI hangs (or unexpected behavior) when trying to load a large configuration file.	Java does not have access to enough memory to handle such a large change to the GUI
On Windows platform, corrupted Latin-1 national characters appear in command prompt	Change font properties of command prompt window
On HP-UX platform, the following message occurs: java.lang.OutOfMemoryError unable to create new native thread	Some HP-UX installations by default allow 64 threads per process. This is insufficient.
On Windows platform, advisors and reach targets mark all servers down	Task offloading is not disabled or may need to enable icmp.
On Windows platform, problem resolving IP address to host name when more than one address is configured to an adapter	The IP address you want as your hostname must appear first in the registry.
On Solaris systems, Load Balancer processes end when you exit the terminal session window from which they started	Use the nohup command to prevent the processes that you started from receiving a hangup signal when you exit the terminal session.

CBR will not run

This problem can occur when another application is using one of the ports used by CBR. For more information, go to [Checking CBR port numbers](#).

cbrcontrol or lbadm command fails

1. The **cbrcontrol** command returns: Error: Server not responding. Or, the **lbadm** command returns: Error: unable to access RMI server. These errors result when your machine has a socksified stack. To correct this problem, edit the `socks.cnf` file to contain the following lines:

```
EXCLUDE-MODULE java
EXCLUDE-MODULE javaw
```

2. The administration consoles for Load Balancer interfaces (command line, graphical user interface, and wizards) communicate with `cbrserver` using remote method invocation (RMI). The default communication uses three ports; each port is set in the **cbrserver** start script:

- 11099 to receive commands from **cbrcontrol**
- 10004 to send metric queries to Metric Server
- 11199 for the RMI server port

This can cause problems when one of the administration consoles runs on the same machine as a firewall or through a firewall. For example, when Load Balancer runs on the same machine as a firewall, and you issue **cbrcontrol** commands, you might see errors such as Error: Server not responding.

To avoid this problem, edit the **cbrserver** script file to set the port used by RMI for the firewall (or other application). Change the line: `LB_RMISERVERPORT=11199` to `LB_RMISERVERPORT=yourPort`. Where *yourPort* is a different port.

When complete, restart **cbrserver** and open traffic for ports 11099, 10004, 11199, and 11100, or for the chosen port for the host address from which the administration console runs.

3. These errors can also occur if you have not already started **cbrserver**.

Requests not being load balanced

Caching Proxy and CBR have been started, but requests are not being load balanced. This error can occur if you start Caching Proxy before starting the executor. If this happens, the `stderr` log for Caching Proxy will contain the following error message: "ndServerInit: Could not attach to executor." To avoid this problem, start the executor before starting Caching Proxy.

On Solaris systems, cbrcontrol executor start command fails

On Solaris systems, the **cbrcontrol executor start** command returns: "Error: Executor was not started." This error occurs if you do not configure the IPC (Inter-process Communication) for the system so that the maximum size of a shared memory segment and semaphore IDs are bigger than the operating system's default. In order to increase the size of the shared memory segment and semaphore IDs, you must edit the `/etc/system` file. For more information on how to configure this file, see the section on modifying the system defaults for IPCs (Inter-process Communication).

Syntactical or configuration error

If the URL rule does not work, this can be a result of either a syntactical or configuration error. For this problem check the following:

- Verify the rule is configured correctly. See Appendix B. Content rule (pattern) syntax, for details.
- Issue a **cbrcontrol** rule report for this rule, and check the 'Times Fired' column to see if it has incremented according to the number of requests made. If it has incremented correctly, recheck the server configuration.
- If the rule is not being fired, add an 'always true' rule. Issue a **cbrcontrol** rule report on the 'always true' rule to verify that it is getting fired.

On Windows platform, unexpected GUI behavior when using Matrox AGP video cards

On Windows platform when using a Matrox AGP card, unexpected behavior can occur in the Load Balancer GUI. When clicking the mouse, a block of space slightly larger than the mouse pointer can become corrupted causing possible highlighting reversal or images to shift out of place on the screen. Older Matrox cards have not shown this behavior. There is no known fix when using Matrox AGP cards.

Disconnect from host occurs when resize Netscape browser window while using Web administration

If you are using remote Web administration to configure Load Balancer, do not resize (Minimize, Maximize, Restore Down, and so on) the Netscape browser window in which the Load Balancer GUI appears. Because Netscape reloads a page every time browser windows are resized, this will cause a disconnect from host. You will need to reconnect to host each time you resize the window. If you are performing remote Web administration on a Windows platform, use Internet Explorer.

On Windows platform, corrupted Latin-1 national characters appear in command prompt window

In a command prompt window on the Windows operating system, some national characters of the Latin-1 family might appear corrupted. For example, the letter "a" with a tilde may display as a pi symbol. To fix this, you must change the font properties of the command prompt window. To change the font, do the following:

1. Click the icon in the upper left corner of the command prompt window
2. Select properties, then click the font tab
3. The default font is raster fonts; change this to Lucida Console and click OK

On HP-UX, Java out of memory, thread error occurs

Some HP-UX 11i installations are pre-configured to allow only 64 threads per process. However, some Load Balancer configurations require more than this amount. For HP-UX systems, set the threads per process to at least 256. To increase this value, use the sam utility to set the max_thread_proc kernel parameter. If heavy use is expected, you might need to increase max_thread_proc beyond 256.

Increase the max_thread_proc parameter, as follows:

1. Type sam from the command line
2. Select **Kernel Configuration > Configurable Parameters**
3. Select **max_thread_proc**
4. Press the spacebar to highlight **max_thread_proc**
5. Press Tab one time, then press the right-arrow key until you select **Actions**

6. Press Enter to display the Actions menu, then press **M** to select Modify Configurable Parameter. (If you do not see this option, highlight **max_thread_proc**)
7. Press Tab until you select the Formula/Value field
8. Type a value of 256 or greater.
9. Click **OK**
10. Press Tab one time, then select **Actions**
11. Press **K** for Process New Kernel.
12. Select **Yes**
13. Restart your system

On Windows systems, advisors and reach targets mark all servers down

When configuring your adapter on a Load Balancer machine, you must ensure that the following two settings are correct for the advisor to work:

- Disable Task Offloading, which is most commonly used on 3Com adapter cards, as follows:
 1. Go to **Start > Settings > Control Panel > Network and Dial-up Connections**, then select the adapter.
 2. In the pop-up window, click **Properties**.
 3. Click **Configure**, then select the **Advanced** tab.
 4. In the property pane, select the **Task Offload** property, then select **disable** in the value field.
- Enable Protocol 1 (ICMP) for IP protocols if you are enabling TCP/IP filtering. If ICMP is not enabled, the ping test to the backend server will not succeed. Check if ICMP is enabled, as follows:
 1. Go to **Start > Settings > Control Panel > Network and Dial-up Connections**, then select the adapter.
 2. In the pop-up window, click **Properties**.
 3. From the components pane, select **Internet Protocol (TCP/IP)**, then click **Properties**.
 4. Click **Advanced**, then select the Options tab.
 5. Select TCP/IP filtering in the options pane, then click **Properties**.
 6. If you have selected Enable TCP/IP Filtering and permit only for IP protocols, you must add IP Protocol 1 in addition to the existing TCP and UDP ports that you enabled.

On Windows systems, resolving IP address to host name when more than one address is configured to an adapter

On Windows platform, when configuring an adapter with more than one IP address, configure the IP address that you want affiliated to the host name first in the registry.

Because Load Balancer is dependent on `InetAddress.getLocalHost()` in many instances (for example, `lbkeys create`), multiple IP addresses aliased to a single adapter might cause problems. To avoid this problem, list the IP address to which you want your host name to resolve first in the registry.

To address this issue, reorder the adapters in the Advanced Settings for the Control Panel's Network Connections option. For example:

1. Open the Control Panel.
2. Open the Network Connections option.
3. From the menu bar, select **Advanced > Advanced Settings...**
4. Reorder the adapters that are listed in the Advanced Settings panel.

Related concepts:

"Content Based Routing (CBR) component" on page 16

CBR works with Caching Proxy to proxy client requests to specified HTTP or HTTPS (SSL) servers. Use CBR to manipulate caching details for faster web document retrieval with low network bandwidth requirements. CBR and Caching Proxy examines HTTP requests by using specified rule types.

Checking CBR port numbers

If you are experiencing problems running CBR, it might be that one of your applications is using a port number that CBR normally uses.

Be aware that CBR uses the following port numbers:

- 11099 to receive commands from **cbrcontrol**
- 10004 to send metric queries to Metric Server

If another application is using one of the port numbers for CBR, you can either change the CBR port numbers or change the application port number.

Change the CBR port numbers, as follows:

- Change the port that receives the commands by modifying the *LB_RMIPORT* variable at the top of the **cbrserver** file to the port that you want CBR to receive commands.
- Change the port that receives the metric reports from the Metric Server by modifying the *RMI_PORT* variable in the metric server file to the port that you want CBR to communicate with Metric Server. Provide the *metric_port* argument when the manager starts. Read the description of the manager start command syntax for more information.

Note: You can find **cbrserver** in the `<install_root>Edge/ULB/cbr/servers/bin` directory.

Related concepts:

"Troubleshooting Content Based Router" on page 185

Use the information that is provided to help you solve problems that can occur in Content Based Routing.

Related reference:

"dscontrol manager" on page 239

You can control the manager function with the `dscontrol manager` command.

Troubleshooting Dispatcher

Use the information that is provided to help you solve problems that can occur in Content Based Routing.

Use the table to see a full description and possible solution for the problem that you are experiencing.

Table 11. Troubleshooting table for Load Balancer

Symptom	Possible Cause
Dispatcher not running correctly	Conflicting port numbers
Connections from client machines not being served or connections timing out	<ul style="list-style-type: none"> • Wrong routing configuration • Server does not have loopback device aliased to the cluster address • Extra route not deleted • Port not defined for each cluster
Dispatcher, Microsoft IIS, and SSL are not working or will not continue	Unable to send encrypted data across protocols
The dscontrol or lbadmin command fails with 'Server not responding' or 'unable to access RMI server' message	<ol style="list-style-type: none"> 1. Commands fail due to socksified stack. Or commands fail due to not starting dsserver 2. RMI ports are not set correctly 3. Host file has incorrect local host
Advisors not working correctly	Advisors are not running
"Cannot Find the File..." error message, when running Netscape as default browser to view online help (Windows platform)	Incorrect setting for HTML file association
Graphical user interface does not start correctly	Insufficient paging space
Graphical user interface does not display correctly.	Resolution is incorrect.
Help panels sometimes disappear behind other windows	Java limitation
GUI hangs (or unexpected behavior) when trying to load a large configuration file.	Java does not have access to enough memory to handle such a large change to the GUI
Korean Load Balancer interface displays overlapping or undesirable fonts on AIX and Linux systems	Default fonts must be changed
Unexpected GUI behavior when using Windows platform paired with Matrox AGP video card	Problem occurs when using Matrox AGP video cards while running the Load Balancer GUI
Slow response time when running commands on the Dispatcher machine	Slow response time can be due to machine overloading from a high volume of client traffic
SSL or HTTPS advisor not registering server loads	Problem occurs because the SSL server application not configured with the cluster IP address
On Windows platform, corrupted Latin-1 national characters appear in command prompt	Change font properties of command prompt window
On Windows platform, advisors and reach targets mark all servers down	Task offloading is not disabled or may need to enable ICMP.
On Windows platform, advisors not working in a high availability setup after a network outage	When the system detects a network outage, it clears its Address Resolution Protocol (ARP) cache

Table 11. Troubleshooting table for Load Balancer (continued)

Symptom	Possible Cause
On Linux systems, "IP address add" command and multiple cluster loopback aliases are incompatible	When aliasing more than one address on the loopback device, should use ifconfig command, not ip address add
On Solaris systems, Load Balancer processes end when you exit the terminal session window from which they started	Use the nohup command to prevent the processes that you started from receiving a hangup signal when you exit the terminal session.
Slow down occurs when loading Load Balancer configurations	The delay might be due to Domain Name System (DNS) calls that are made to resolve and verify the server address.
On Windows systems, the following error message appears: There is an IP address conflict with another system on the network	If high availability is configured, cluster addresses may be configured on both machines for a brief period which causes this error message to appear.
On Windows systems, "Server not responding" error occurs when issuing a dscontrol or lbadm command	When more than one IP address exists on a Windows system and the host file does not specify the address to associate with the hostname.
Dispatcher MAC forwarding configuration limitations with zSeries® and S/390® platforms	On Linux, there are limitations when using zSeries or S/390 servers that have Open System Adapter (OSA) cards. Possible workarounds are provided.
On Linux systems, iptables can interfere with the routing of packets	Linux iptables can interfere with load balancing of traffic and must be disabled on the Load Balancer machine.
On Solaris systems, when you try to configure an IPv6 server on the Dispatcher machine, the message "unable to add server" appears	This can be caused by the way the Solaris operating system handles the ping request for an IPv6 address.
Upgrading the Java fileset provided with the Load Balancer installations	If a problem is found with the Java file set, you should report the problem to IBM Service so that you can receive an upgrade for the Java file set that was provided with the Load Balancer installation.
Client requests fail when forwarding to HP-UX back-end servers	After setting up Load Balancer for IPv6 on the HP-UX operating system, client requests to the cluster address fail. This error is a result of the interaction between the neighbor discovery function for the operating system and the Load Balancer.
Load Balancer for IPv4 and IPv6 conflicts with IP security (IPsec)	<p>If you are using the Load Balancer for IPv4 and IPv6 with IP security (IPsec) enabled, output packets might be incorrect and dispatcher configuration information might display incorrectly in the command line interface and administrative console for WebSphere Application Server.</p> <p>Load Balancer reports that it is forwarding connections, but clients do not receive responses.</p>

Table 11. Troubleshooting table for Load Balancer (continued)

Symptom	Possible Cause
The serverUp script might run when you issue commands for Load Balancer that affect the status of servers	You might experience problems if you run a command that affects the status of a server, such as the dscontrol manager unquiesce and dscontrol manager quiesce commands, after a manager cycle has already retrieved the weights of the servers. If you run these commands, it might overwrite the values that are saved during the manager cycle and cause the serverUp script to run unexpectedly.

Dispatcher not running correctly

This problem can occur when another application is using one of the ports used by the Load Balancer. For more information, read the Configuring the Load Balancer machine topic.

Load Balancer requests are not being balanced

This problem has symptoms such as connections from client machines not being served or connections timing out. Check the following to diagnose this problem:

1. Have you configured the nonforwarding address, clusters, ports, and servers for routing? Check the configuration file.
2. Does the loopback device on each server have the alias set to the cluster address?
Use `netstat -ni` to check.
3. Is the extra route deleted?
Use `netstat -nr` to check.
4. Use the **dscontrol cluster status** command to check the information for each cluster you have defined. Make sure you have a port defined for each cluster.
5. Use the **dscontrol server report ::** command to make sure that your servers are neither down nor set to a weight of zero.

Dispatcher, Microsoft IIS, and SSL do not work (Windows platform)

When using Dispatcher, Microsoft IIS, and SSL, if they do not work together, there may be a problem with enabling SSL security. For more information about generating a key pair, acquiring a certificate, installing a certificate with a key pair, and configuring a directory to require SSL, see the *Microsoft Information and Peer Web Services* documentation.

dscontrol or lbadm command fails

1. The **dscontrol** command returns: Error: Server not responding. Or, the **lbadm** command returns: Error: unable to access RMI server. These errors can result when your machine has a socksified stack. To correct this problem, edit the `socks.cnf` file to contain the following lines:

```
EXCLUDE-MODULE java
EXCLUDE-MODULE javaw
```
2. The administration consoles for Load Balancer interfaces (command line, graphical user interface, and wizards) communicate with **dserver** using

remote method invocation (RMI). The default communication uses three ports; each port is set in the `dsserver` start script:

- 10099 to receive commands from **dscontrol**
- 10004 to send metric queries to Metric Server
- 10199 for the RMI server port

This can cause problems when one of the administration consoles runs on the same machine as a firewall or through a firewall. For example, when Load Balancer runs on the same machine as a firewall, and you issue **dscontrol** commands, you might see errors such as Error: Server not responding.

To avoid this problem, edit the `dsserver` script file to set the port used by RMI for the firewall (or other application). Change the line:

`LB_RMISERVERPORT=10199` to `LB_RMISERVERPORT=yourPort`. Where *yourPort* is a different port.

When complete, restart `dsserver` and open traffic for ports 10099, 10004, 10199, and 10100, or for the chosen port for the host address from which the administration console will be run.

3. These errors can also occur if you have not already started **dsserver**.
4. If there are multiple adapters on the machine, you must designate which adapter that `dsserver` is to use by adding the following in the `dsserver` script: `java.rmi.server.hostname=<host_name or IPaddress>`

For example: `java -Djava.rmi.server.hostname="10.1.1.1"`

Advisors not working correctly

An ICMP ping is issued to the servers before the advisor request. If a firewall exists between Load Balancer and the servers, ensure that pings are supported across the firewall. If this setup poses a security risk to your network, modify the `java` statement in `dsserver` to turn off all pings to the servers by adding the `java` property:

```
LB_ADV_NO_PING="true"
java -DLB_ADV_NO_PING="true"
```

"Cannot find the file..." error message when trying to view online Help (Windows platform)

For Windows platforms, when using Netscape as your default browser, the following error message may result: "Cannot find the file '<filename>.html' (or one of its components). Make sure the path and filename are correct and that all required libraries are available."

The problem is due to an incorrect setting for HTML file association. The solution is the following:

1. Click **My Computer**, click **Tools**, select **Folder Options**, and click **File Types** tab
2. Select "Netscape Hypertext Document"
3. Click **Advanced** button, select **open**, click **Edit** button
4. Enter `NSShell` in the Application: field (not the Application Used to Perform Action: field), and click **OK**

Graphical user interface (GUI) does not start correctly

The graphical user interface (GUI), which is **ladmin**, requires a sufficient amount of paging space to function correctly. If insufficient paging space is available, the

GUI might not start up completely. If this occurs, check your paging space and increase it if necessary.

Graphical user interface (GUI) does not display correctly

If you experience a problem with the appearance of the Load Balancer GUI, check the setting for the operating system's desktop resolution. The GUI is best viewed at a resolution of 1024x768 pixels.

On Windows platform, help windows sometimes disappear behind other open windows

On Windows platform, when you first open help windows, they sometimes disappear into the background behind existing windows. If this occurs, click on the window to bring it forward again.

GUI hangs (or unexpected behavior) when trying to load a large configuration file

When using **lbadmin** or Web administration (**lbwebaccess**) to load a large configuration file (roughly 200 or more **add** commands), the GUI may hang or display unexpected behavior, such as responding to screen changes at an extremely slow rate of speed.

This occurs because Java does not have access to enough memory to handle such a large configuration.

There is an option on the runtime environment that can be specified to increase the memory allocation pool available to Java.

The option is `-Xmxn` where `n` is the maximum size, in bytes, for the memory allocation pool. `n` must be a multiple of 1024 and must be greater than 2MB. The value `n` may be followed by `k` or `K` to indicate kilobytes, or `m` or `M` to indicate megabytes. For example, `-Xmx128M` and `-Xmx81920k` are both valid. The default value is 64M. Solaris 8 has a maximum value of 4000M.

For example, to add this option, edit the **lbadmin** script file, modifying "javaw" to "javaw -Xmxn" as follows. For AIX systems, modify "java" to "java -Xmxn".

- AIX systems

```
java -Xmx256m -cp $LB_CLASSPATH $LB_INSTALL_PATH $LB_CLIENT_KEYS  
com.ibm.internet.nd.framework.FWK_Main 1>/dev/null 2>&1 &
```

- HP-UX systems

```
java -Xmx256m -cp $LB_CLASSPATH $LB_INSTALL_PATH $LB_CLIENT_KEYS  
com.ibm.internet.nd.framework.FWK_Main 1>/dev/null 2>&1 &
```

- Linux systems

```
javaw -Xmx256m -cp $LB_CLASSPATH $LB_INSTALL_PATH $LB_CLIENT_KEYS  
com.ibm.internet.nd.framework.FWK_Main 1>/dev/null 2>&1 &
```

- Solaris systems

```
java -Xmx256m -cp $LB_CLASSPATH $LB_INSTALL_PATH $LB_CLIENT_KEYS  
com.ibm.internet.nd.framework.FWK_Main 1>/dev/null 2>&1 &
```

- Windows systems

```
START javaw -Xmx256m -cp %LB_CLASSPATH% %LB_INSTALL_PATH%  
%LB_CLIENT_KEYS% com.ibm.internet.nd.framework.FWK_Main
```

There is no recommended value for n , but it should be greater than the default option. A good place to start would be with twice the default value.

Korean Load Balancer interface displays overlapping or undesirable fonts on AIX and Linux systems

To correct overlapping or undesirable fonts in the Korean Load Balancer interface:

- On AIX systems
 1. Stop all Java processes on the AIX system.
 2. Open the `font.properties.ko` file in an editor. This file is located in `home/jre/lib` where `home` is the Java home.
 3. Search for the following string:

```
-Monotype-TimesNewRomanWT-medium-r-normal  
--*-%d-75-75-*--ksc5601.1987-0
```
 4. Replace all instances of the string with:

```
-Monotype-SansMonoWT-medium-r-normal  
--*-%d-75-75-*--ksc5601.1987-0
```
 5. Save the file.
- On Linux systems
 1. Stop all Java processes on the system.
 2. Open the `font.properties.ko` file in an editor. This file is located in `home/jre/lib` where `home` is the Java home.
 3. Search for the following string (with no spaces):

```
-monotype-  
timesnewromanwt-medium-r-normal---*-%d-75-75-p*-microsoft-symbol
```
 4. Replace all instances of the string with:

```
-monotype-sansmonowt-medium-r-normal---*-%d-75-75-p*-microsoft-symbol
```
 5. Save the file.

On Windows platform, unexpected GUI behavior when using Matrox AGP video cards

On Windows platform when using a Matrox AGP card, unexpected behavior can occur in the Load Balancer GUI. When clicking the mouse, a block of space slightly larger than the mouse pointer can become corrupted causing possible highlighting reversal or images to shift out of place on the screen. Older Matrox cards have not shown this behavior. There is no known fix when using Matrox AGP cards.

Slow response time running commands on Dispatcher machine

If you are running the Dispatcher component for load balancing, it is possible to overload the computer with client traffic. The Load Balancer kernel module has the highest priority, and if it is constantly handling client packets, the rest of the system may become unresponsive. Running commands in user space may take a very long time to complete, or may never complete.

If this happens, you should begin to restructure your setup to avoid overloading the Load Balancer machine with traffic. Alternatives include spreading the load across several Load Balancer machines, or replacing the machine with a stronger and faster computer.

When trying to decide if the slow response time on the machine is due to high client traffic, consider whether this occurs during client peak traffic times.

Misconfigured systems that cause routing loops can also cause the same symptoms. But before changing the Load Balancer setup, determine whether the symptoms may be due to high client load.

SSL or HTTPS advisor not registering server loads

Load Balancer will send packets to the servers using the cluster address that is aliased on the loopback. Some server applications (such as SSL) require that configuration information, such as certificates, are based on the IP address. The IP address must be the cluster address which is configured on the loopback in order to match the contents of the incoming packets. If the IP address of the cluster is not used when configuring the server application, then the client request will not get properly forwarded to the server.

On Windows systems, corrupted Latin-1 national characters appear in command prompt window

In a command prompt window on the Windows operating system, some national characters of the Latin-1 family might appear corrupted. For example, the letter "a" with a tilde may display as a pi symbol. To fix this, you must change the font properties of the command prompt window. Change the font, as follows:

1. Click the icon in the upper left corner of the command prompt window
2. Select Properties, then click the Font tab
3. The default font is Raster fonts; change this to Lucida Console and click OK

On Windows systems, advisors and reach targets mark all servers down

When configuring your adapter on a Load Balancer machine, you must ensure that the following two settings are correct for the advisor to work:

- Disable Task Offloading.
 - To disable Task offloading: Go to **Start > Settings > Control Panel > Network and Dial-up Connections**, then select the adapter.
 - In the pop-up window, click **Properties**.
 - Click **Configure**, then select the Advanced tab.
 - In the property pane, select the **Task Offload** property, then select disable in the value field.
- Enable Protocol 1 (ICMP) for IP protocols if you are enabling TCP/IP filtering. If ICMP is not enabled, the ping test to the back-end server will not succeed. To check whether ICMP is enabled:
 - Go to **Start > Settings > Control Panel > Network and Dial-up Connections**, then select the adapter.
 - In the pop-up window, click **Properties**.
 - From the components pane, select **Internet Protocol (TCP/IP)**, then click **Properties**.
 - Click **Advanced**, then select the Options tab.
 - Select TCP/IP filtering in the options pane, then click Properties.
 - If you have selected **Enable TCP/IP Filtering** and **permit only** for IP protocols, you must add IP Protocol 1. This must be added in addition to the existing TCP and UDP ports that you enabled.

On Windows systems, after network outage, advisors not working in a high availability setup

By default, when the Windows operating system detects a network outage, it clears its address resolution protocol (ARP) cache, including all static entries. After the network is available, the ARP cache is repopulated by ARP requests sent on the network.

With a high availability configuration, both servers take over primary operations when a loss of network connectivity affects one or both. When the ARP request is sent to repopulate the ARP cache, both servers respond, which causes the ARP cache to mark the entry as not valid. Therefore, the advisors are not able to create a socket to the backup servers.

Preventing the Windows operating system from clearing the ARP cache when there is a loss of connectivity solves this problem. Microsoft has published an article that explains how to accomplish this task. This article is on the Microsoft Web site, located in the Microsoft Knowledge Base, article number 239924: <http://support.microsoft.com/default.aspx?scid=kb;en-us;239924>.

The following is a summary of the steps, described in the Microsoft article, to prevent the system from clearing the ARP cache:

1. Use the Registry editor (regedit or regedit32) to open the registry.
2. View the following key in the registry:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters
3. Add the following registry value: Value Name: DisableDHCPMediaSense Value Type: REG_DWORD.
4. After the key is added, edit the value and set it to 1.
5. Reboot the machine for the change to take effect.

Attention: This affects the ARP cache regardless of the DHCP setting.

On Linux systems, do not use "IP address add" command when aliasing multiple clusters on the loopback device

Certain considerations must be taken when using Linux kernel 2.4.x servers. If the server has a cluster address configured on the loopback device using the **ip address add** command, only one cluster address can be aliased.

When aliasing multiple clusters to the loopback device use the **ifconfig** command, for example:

```
ifconfig lo:num clusterAddress netmask 255.255.255.255 up
```

Additionally, there are incompatibilities between the **ifconfig** method of configuring interfaces and the **ip** method of configuring interfaces. Best practice suggests that a site choose one method and use that method exclusively.

On Solaris systems, Load Balancer processes end when you exit the terminal window from which they started

On Solaris systems, after starting Load Balancer scripts (such as `dsserver` or `lbadmin`) from a terminal window, if you exit from that window, the Load Balancer process also exits.

To resolve this problem, start the Load Balancer scripts with the **nohup** command. For example: **nohup dsserver**. This command prevents the processes started from the terminal session from receiving a hangup signal from the terminal when it exits, allowing the processes to continue even after the terminal session has ended. Use the **nohup** command in front of any Load Balancer scripts that you want to continue to process beyond the end of a terminal session.

Delay occurs while loading a Load Balancer configuration

Loading a Load Balancer configuration might take a long time due to Domain Name System (DNS) calls that are made to resolve and verify the server address.

If the DNS of the Load Balancer machine is configured incorrectly, or if DNS in general takes a long time, this will cause a slow down in loading the configuration due to the Java processes that are sending DNS requests on the network.

A workaround for this is to add your server addresses and hostnames to your local `/etc/hosts` file.

On Windows systems, an IP address conflict error message appears

If high availability is configured, the cluster addresses may be configured on both machines for a brief period and cause the following error message to occur: There is an IP address conflict with another system on the network. In this case, you can safely ignore the message. It is possible for a cluster address to be briefly configured on both high availability machines at the same time, especially during startup of either machine, or when a takeover has been initiated.

On Windows systems, "Server not responding" error occurs when issuing dscontrol or lbadm

To resolve this problem, update the `c:\Windows\system32\drivers\etc\hosts` file with your machine host name and the IP address that you want to associate with the host name.

If you are using **dscontrol**, you can specify the connection address using the following command:

```
dscontrol host@<ip_address or host_name> <command>
```

gotcha: The IP address cannot be a cluster address.

On Linux, Dispatcher configuration limitations when using zSeries or S/390 servers that have Open System Adapter (OSA) cards

In general, servers in the Load Balancer configuration must all be on the same network segment regardless of the platform. Active network devices such as router, bridges, and firewalls interfere with Load Balancer. This is because Load Balancer functions as a specialized router, modifying only the link-layer headers to its next and final hop. Any network topology in which the next hop is not the final hop is not valid for Load Balancer.

Note: Tunnels, such as channel-to-channel (CTC) or inter-user communication vehicle (IUCV), are often supported. However, Load Balancer must forward across the tunnel directly to the final destination, it cannot be a network-to-network tunnel.

There is a limitation for zSeries and S/390 servers that share the OSA card, because this adapter operates differently than most network cards. The OSA card has its own virtual link layer implementation, which has nothing to do with ethernet, that is presented to the Linux and z/OS hosts behind it. Effectively, each OSA card looks just like ethernet-to-ethernet hosts (and not to the OSA hosts), and hosts that use it will respond to it as if it is ethernet.

The OSA card also performs some functions that relate to the IP layer directly. Responding to ARP (address resolution protocol) requests is one example of a function that it performs. Another is that shared OSA routes IP packets based on destination IP address, instead of on ethernet address as a layer 2 switch. Effectively, the OSA card is a bridged network segment unto itself.

Load Balancer that runs on an S/390 Linux or zSeries Linux host can forward to hosts on the same OSA or to hosts on the ethernet. All the hosts on the same shared OSA are effectively on the same segment.

Load Balancer can *forward out* of a shared OSA because of the nature of the OSA bridge. The bridge knows the OSA port that owns the cluster IP. The bridge knows the MAC address of hosts directly connected to the ethernet segment. Therefore, Load Balancer can MAC-forward across one OSA bridge.

However, Load Balancer cannot forward into a shared OSA. This includes the Load Balancer on an S/390 Linux when the back-end server is on a different OSA card than the Load Balancer. The OSA for the back-end server advertises the OSA MAC address for the server IP, but when a packet arrives with the ethernet destination address of the server's OSA and the IP of the cluster, the server's OSA card does not know which of its hosts, if any, should receive that packet. The same principles that permit OSA-to-ethernet MAC-forwarding to work out of one shared OSA do not hold when trying to forward into a shared OSA.

Workaround:

In Load Balancer configurations that use zSeries or S/390 servers that have OSA cards, there are two approaches you can take to work around the problem that has been described.

1. Using platform features

If the servers in the Load Balancer configuration are on the same zSeries or S/390 platform type, you can define point-to-point (CTC or IUCV) connections between Load Balancer and each server. Set up the endpoints with private IP addresses. The point-to-point connection is used for Load Balancer-to-server traffic only. Then add the servers with the IP address of the server endpoint of the tunnel. With this configuration, the cluster traffic comes through the Load Balancer OSA card and is forwarded across the point-to-point connection where the server responds through its own default route. The response uses the server's OSA card to leave, which might or might not be the same card.

2. Using Load Balancer's encapsulation feature.

If the servers in the Load Balancer configuration are not on the same zSeries or S/390 platform type, or if it is not possible to define a point-to-point connection

between Load Balancer and each server, it is recommended that you use Load Balancer's encapsulation feature, which is a protocol that permits Load Balancer to forward across routers.

When using encapsulation, the client->cluster IP packet is received by Load Balancer, encapsulated, and sent to the server. At the server, the original client->cluster IP packet is excapsulated, and the server responds directly to the client. The advantage with using GRE is that Load Balancer sees only the client-to-server traffic, not the server-to-client traffic. The disadvantage is that it lowers the maximum segment size (MSS) of the TCP connection due to encapsulation overhead.

Refer to the topic "Use encapsulation forwarding to forward traffic across network segments" on page 154 for more information on how to configure Load Balancer to forward with encapsulation.

Linux iptables can interfere with the routing of packets

Linux iptables can interfere with load balancing of traffic and must be disabled on the Dispatcher machine.

Issue the following command to determine if iptables are loaded:

```
lsmod | grep ip_tables
```

The output from the preceding command might be similar to this:

```
ip_tables          22400  3
iptables_mangle,iptable_nat,iptable_filter
```

Issue the following command for each iptable listed in the output to display the rules for the tables:

```
iptables -t <short_name> -L
```

For example:

```
iptables -t mangle -L
iptables -t nat -L
iptables -t filter -L
```

If iptable_nat is loaded, it must be unloaded. Because iptable_nat has a dependency on iptable_conntrack, iptable_conntrack also must be removed. Issue the following command to unload these two iptables:

```
rmmmod iptable_nat iptable_conntrack
```

Unable to add an IPv6 server to the Load Balancer configuration on Solaris systems

On Solaris systems, when you try to configure an IPv6 server on a Websphere Load Balancer for IPv4 and IPv6 installation, you might receive the message unable to add server. This can be caused by the way the Solaris operating system handles the ping request for an IPv6 address.

When adding a server to the configuration on Solaris systems, Load Balancer tries to ping the server to obtain the MAC address of the server. The Solaris machine might choose a configured cluster address as the source address of the ping request, instead of using the NFA address of the machine. If the cluster address is configured on the server loopback, the ping response is not received at the Load Balancer machine; therefore, it does not add the server to the configuration.

The solution is to configure another IPv6 address on the Load Balancer machine either before or after configuring the IPv6 cluster address. This address must be an address that is not aliased on the loopback of the backend server on which you are trying to add to the Load Balancer configuration. Then add the server to the Load Balancer configuration.

Upgrading the Java file set provided with the Load Balancer installation

During the Load Balancer installation process, a Java file set also gets installed. Load Balancer will be the only application that uses the Java version which installs with the product. You should not upgrade this version of the Java file set on your own. If there are problem which requires an upgrade for the Java file set, you should report the problem to IBM Service so the Java file set which is shipped within Load Balancer will be upgraded with an official fix level.

Problem: Client requests fail when using IPv6 MAC forwarding with HP-UX back-end servers

After setting up Load Balancer for IPv6 on the HP-UX operating system, client requests to the cluster address fail. This error is a result of the interaction between the neighbor discovery function for the operating system and the Load Balancer.

When a back-end server is added to the configuration, Load Balancer tries to ping the new server for the MAC address. The HP-UX server might choose a configured cluster address as the source address of the ping request, instead of using the nonforwarding address (NFA) of the machine. In this case, a new entry is created for the cluster address in the routing table of the back-end HP-UX server in addition to the local loopback entry. The new entry has a higher routing priority than the local loopback. Thus, the client requests that reach the back-end server are then redirected back to the Load Balancer server, which does not respond.

To work around this problem, after Load Balancer is set up completely, configure the loopback alias on the back-end server as a final step. If the cluster address is aliased on the loopback when the Load Balancer configuration is loaded, remove the cluster loopback alias on the back-end server completely and then re-alias it. To bring down the loopback alias, use the `ifconfig lo0:1 inet6` command from the terminal window. Re-alias the loopback alias.

On AIX systems, Load Balancer conflicts with IP security (IPsec)

If you are using Load Balancer with IP security (IPsec) enabled, output packets might be incorrect and dispatcher configuration information might display incorrectly in the command line interface and administrative console for WebSphere Application Server. Load Balancer reports that it is forwarding connections, but clients do not receive responses.

If you are using Load Balancer function and IP security on the same host, there might be communication problems between Load Balancer and the application server. The Load Balancer component is not fully compatible with IPsec features and it transmits data from both sides of the security layer. Load Balancer receives packets below IPsec and, as a result, receives encrypted packets that it does not decrypt. When sending data, Load Balancer transmits them above IPsec, so it sends unencrypted packets to the application server that are encrypted on the other end by IPsec. The application server, therefore, receives encrypted data that cannot be used.

The serverUp script might run when you issue commands for Load Balancer that affect the status of servers

Weights are set by the manager during a manager cycle. At the start of the manager cycle, the manager retrieves the current weights from the executor function. The manager uses these values as the last known weight to determine if the status of a server has changed:

- If you issue a quiesce command for a server, the executor function saves the current weight of the server and associates a new weight to the server with a value of -1. This is an example of the quiesce command:

```
dscontrol manager quiesce server
```
- If you issue the unquiesce command, a call is made to the executor function to revert the weight of the server to the saved value. The system sets a flag to indicate that the server is no longer marked down by the user. This is an example of the quiesce command:

```
dscontrol manager unquiesce server
```

If the unquiesce command occurs after the manager has retrieved the weights, the executor function overwrites the weight that is used to determine if the server state has changed. This process does not cause any side effects unless the server is also quiesced.

The chances of experiencing this problem increase with larger configurations because the manager cycle takes longer to run. Also, there is a higher probability that the manager cycle will be in progress when the unquiesce command is issued.

Related concepts:

“Content Based Routing (CBR) component” on page 16
CBR works with Caching Proxy to proxy client requests to specified HTTP or HTTPS (SSL) servers. Use CBR to manipulate caching details for faster web document retrieval with low network bandwidth requirements. CBR and Caching Proxy examines HTTP requests by using specified rule types.

Related tasks:

“Configuring the Dispatcher machine” on page 69
Configure Dispatcher for IPv4 and IPv6 on the machine that you will use to load balance server traffic.

Checking Dispatcher port numbers

If you are experiencing problems running Dispatcher, it may be that one of your applications is using a port number that the Dispatcher normally uses.

Be aware that the Dispatcher server uses the following port numbers:

- 10099 to receive commands from **dscontrol**
- 10004 to send metric queries to Metric Server

If another application is using one of the Dispatcher port numbers, you can either change the Dispatcher port numbers or change the application port number.

Change the Dispatcher port numbers, as follows:

- Change the port that receives the commands by modifying the *LB_RMIPORT* variable at the top of the **dsserver** file to the port that you want Dispatcher to receive commands.
- Change the port that receives the metric reports from the Metric Server by modifying the *RMI_PORT* variable in the metric server file to the port that you

want Dispatcher to communicate with Metric Server. Provide the `metric_port` argument when the manager starts. Read the description of the manager start command syntax for more information.

Note: You can find `dsserver` in the `<install_root>Edge/ULB/servers/bin` directory.

Related concepts:

“Troubleshooting Dispatcher” on page 189

Use the information that is provided to help you solve problems that can occur in Content Based Routing.

Related reference:

“dscontrol manager” on page 239

You can control the manager function with the `dscontrol manager` command.

Troubleshooting Site Selector

Use the information that is provided to help you solve problems that can occur in Site Selector.

Click a link in the table to go to a full description and possible solution for the problem that you are experiencing.

Table 12. Site Selector Troubleshooting table

Symptom	Possible cause
Site Selector not running correctly	Conflicting port number
Site Selector does not round-robin incoming requests from Solaris client	Solaris systems run a "name service cache daemon"
The <code>sscontrol</code> or <code>lbadmin</code> command fails with 'Server not responding' or 'unable to access RMI server' message	Commands fail due to socksified stack. Or commands fail due to not starting <code>ssserver</code> .
<code>ssserver</code> fails to start on Windows platform	Windows systems do not require the host name to be in the DNS.
Machine with duplicate routes not load balancing correctly - name resolution appears to fail	Site Selector machine with multiple adapters attached to the same subnet
Unexpected GUI behavior when using Windows platform paired with Matrox AGP video card	Problem occurs when using Matrox AGP video cards while running the Load Balancer GUI
GUI hangs (or unexpected behavior) when trying to load a large configuration file.	Java does not have access to enough memory to handle such a large change to the GUI
On Windows platform, corrupted Latin-1 national characters appear in command prompt	Change font properties of command prompt window
On HP-UX platform, the following message occurs: <code>java.lang.OutOfMemoryError</code> unable to create new native thread	Some HP-UX installations by default allow 64 threads per process. This is insufficient.
On Windows platform, advisors and reach targets mark all servers down	Task offloading is not disabled or may need to enable <code>icmp</code> .
On Solaris systems, Load Balancer processes end when you exit the terminal session window from which they started	Use the <code>nohup</code> command to prevent the processes that you started from receiving a hangup signal when you exit the terminal session.

Site Selector will not run

This problem can occur when another application is using one of the ports used by Site Selector. For more information, go to Checking Site Selector port numbers.

Site Selector does not round-robin traffic from Solaris clients

Symptom: Site Selector component does not round-robin incoming requests from Solaris clients.

Possible cause: Solaris systems run a name service cache daemon. If this daemon is running, the subsequent resolver request is answered from this cache instead of querying Site Selector.

Solution: Turn off the name service cache daemon on the Solaris machine.

sscontrol or lbadmin command fails

1. The **sscontrol** command returns: Error: Server not responding. Or, the **lbadmin** command returns: Error: unable to access RMI server. These errors can result when your machine has a socksified stack. To correct this problem, edit the `socks.cnf` file to contain the following lines:

```
EXCLUDE-MODULE java
EXCLUDE-MODULE javaw
```

2. The administration consoles for Load Balancer interfaces (command line, graphical user interface, and wizards) communicate with **ssserver** using remote method invocation (RMI). The default communication uses three ports; each port is set in the **ssserver** start script:
 - 12099 to receive commands from **sscontrol**
 - 10004 to send metric queries to Metric Server
 - 12199 for the RMI server port
 - 53 for sending and receiving DNS traffic

This can cause problems when one of the administration consoles runs on the same machine as a firewall or through a firewall. For example, when Load Balancer runs on the same machine as a firewall, and you issue **sscontrol** commands, you might see errors such as Error: Server not responding.

To avoid this problem, edit the **ssserver** script file to set the port used by RMI for the firewall (or other application). Change the line: `LB_RMISERVERPORT=10199` to `LB_RMISERVERPORT=yourPort`. Where *yourPort* is a different port.

When complete, restart **ssserver** and open traffic for ports 12099, 10004, 12199, and 12100, or for the chosen port for the host address from which the administration console will run.

3. These errors can also occur if you have not already started **ssserver**.

The ssserver is failing to start on Windows platform

Site Selector must be able to participate in a DNS. All the machines involved in the configuration should also be participants of this system. Windows systems do not always require the configured host name to be in the DNS. Site Selector requires that its host name be defined in the DNS to start properly.

Verify this host is defined in the DNS. Edit the `ssserver.cmd` file and remove the `w` from `javaw`. This change should provide more information about errors.

Site Selector with duplicate routes not load balancing correctly

The Site Selector name server does not bind to any one address on the machine. It will respond to requests destined for any valid IP on the machine. Site Selector relies on the operating system to route the response back to the client. If the Site Selector machine has multiple adapters and any number of them are attached to the same subnet, it is possible the O/S will send the response to the client from a different address than it was received. Some client applications will not accept a response received from an address other than where it was sent. As a result, the name resolution will appear to fail.

On Windows platform, unexpected GUI behavior when using Matrox AGP video cards

On Windows platform when using a Matrox AGP card, unexpected behavior can occur in the Load Balancer GUI. When clicking the mouse, a block of space slightly larger than the mouse pointer can become corrupted causing possible highlighting reversal or images to shift out of place on the screen. Older Matrox cards have not shown this behavior. There is no known fix when using Matrox AGP cards.

On Windows platform, corrupted Latin-1 national characters appear in command prompt window

In a command prompt window on the Windows operating system, some national characters of the Latin-1 family might appear corrupted. For example, the letter "a" with a tilde may display as a pi symbol. To fix this, you must change the font properties of the command prompt window. To change the font, do the following:

1. Click the icon in the upper left corner of the command prompt window
2. Select properties, then click the font tab
3. The default font is raster fonts; change this to Lucida Console and click OK

On HP-UX, Java out of memory/ thread error occurs

Some HP-UX 11i installations are pre-configured to allow only 64 threads per process. However, some Load Balancer configurations require more than this amount. For HP-UX systems, set the threads per process to at least 256. To increase this value, use the sam utility to set the max_thread_proc kernel parameter. If heavy use is expected, you might need to increase max_thread_proc beyond 256.

To increase max_thread_proc, refer to the steps to increase the max_thread_proc parameter.

On Windows systems, advisors and reach targets mark all servers down

When configuring your adapter on a Load Balancer machine, you must ensure that the following two settings are correct for the advisor to work:

- Disable Task Offloading, which is most commonly used on 3Com adapter cards.
- Enable Protocol 1 (ICMP) for IP protocols if you are enabling TCP/IP filtering. If ICMP is not enabled, the ping test to the backend server will not succeed.

Refer to the section on disabling task offloading for instructions on configuring this setting.

Related concepts:

“Site Selector component” on page 19

Site Selector acts as a name server that works with other name servers in a domain name system to load balance among a group of servers by using measurements and weights that are gathered. You can create a site configuration to load balance traffic among a group of servers based on the domain name used for a client's request.

Checking Site Selector port numbers

If you are experiencing problems running the Site Selector component, it may be that one of your applications is using a port number that Site Selector normally uses.

Be aware that Site Selector uses the following port numbers:

- 12099 to receive commands from **sscontrol**
- 10004 to send metric queries to Metric Server

If another application is using one of the Site Selector port numbers, you can either change the Site Selector port numbers or change the application port number.

Change the Site Selector port numbers, as follows:

- Change the port that receives the commands by modifying the *LB_RMIPORT* variable at the top of the **ssserver** file to the port that you want Site Selector to receive commands.
- Change the port that receives the metric reports from the Metric Server by modifying the *RMI_PORT* variable in the metric server file to the port that you want Site Selector to communicate with Metric Server. Provide the *metric_port* argument when the manager starts. Read the description of the manager start command syntax for more information.

Note: You can find **ssserver** in the `<install_root>Edge/ULB/ss/servers/bin` directory.

Related concepts:

“Troubleshooting Site Selector” on page 203

Use the information that is provided to help you solve problems that can occur in Site Selector.

Related reference:

“sscontrol manager” on page 261


Use the **sscontrol manager** command to control the manager.

Chapter 9. Reference

Reference information is organized to help you locate particular facts quickly.

Advanced configuration

Settings are properties that you can configure using the configuration files, or by other means..

To open the information center table of contents to the location of this reference information, click the **Show in Table of Contents** button () on your information center border.

Directory conventions

References in product information to *install_root* and other directories infer specific default directory locations. This topic describes the conventions in use for IBM WebSphere Edge Components.

These file paths are default locations. You can install the product and other components in any directory where you have write access. You can create profiles in any valid directory where you have write access.

install_root - the root directory in which the product was installed

Load Balancer install paths include the following:

- Administration - /opt/IBM/WebSphere/Edge/ULB/admin
- Dispatcher - /opt/IBM/WebSphere/Edge/ULB/servers
- Content Based Routing- /opt/IBM/WebSphere/Edge/ULB/cbr/servers
- Site Selector- /opt/IBM/WebSphere/Edge/ULB/ss/servers
- Metric Server - /opt/IBM/WebSphere/Edge/ULB/ms
- Documentation - /opt/IBM/WebSphere/Edge/ULB/docs/

install_root - the root directory in which the product was installed

Load Balancer install paths include the following:

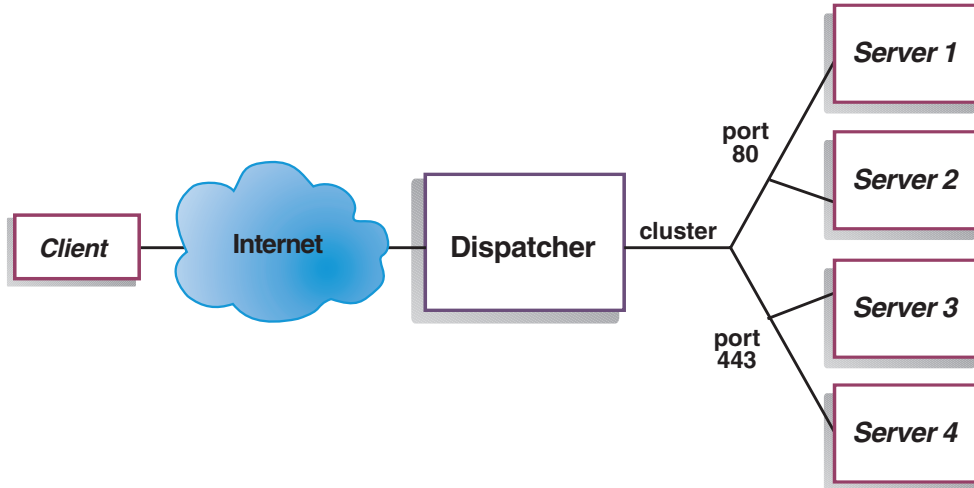
- Administration - C:\Program Files\IBM\WebSphere\Edge\ULB\admin
- Dispatcher - C:\Program Files\IBM\WebSphere\Edge\ULB\servers
- Content Based Routing- C:\Program Files\IBM\WebSphere\Edge\ULB\cbr\servers
- Site Selector- C:\Program Files\IBM\WebSphere\Edge\ULB\ss\servers
- Metric Server - C:\Program Files\IBM\WebSphere\Edge\ULB\ms
- Documentation - C:\Program Files\IBM\WebSphere\Edge\ULB\docs

Types of cluster, port, and server configurations for Dispatcher

There are many ways that you can configure Dispatcher to support your site.

1 cluster with 2 ports

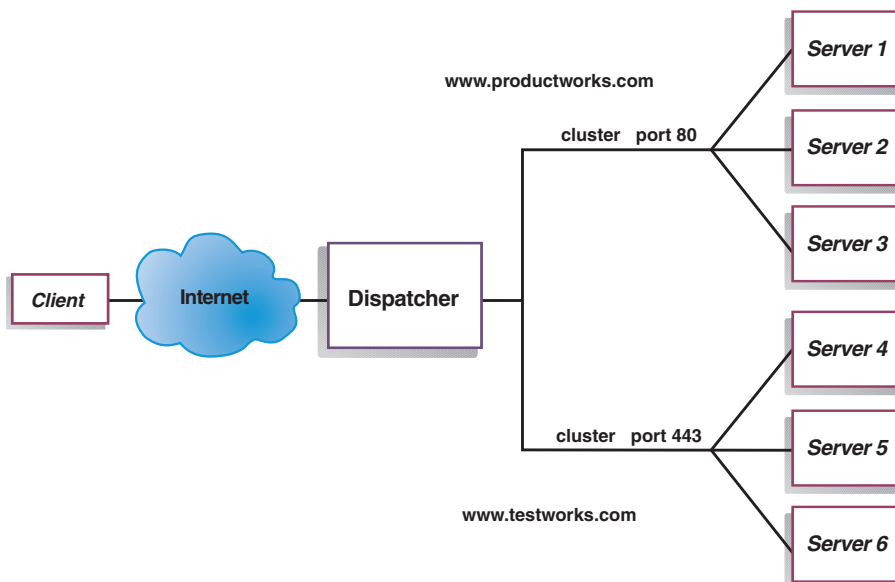
If you have only one host name for your site to which all of your customers will connect, you can define a single cluster of servers. For each of these servers, configure a port through which Load Balancer communicates.



In this example for the Dispatcher component, one cluster is defined at `www.productworks.com`. This cluster has two ports: port 80 for HTTP and port 443 for SSL. A client making a request to `http://www.productworks.com` (port 80) goes to a different server than a client requesting `https://www.productworks.com` (port 443).

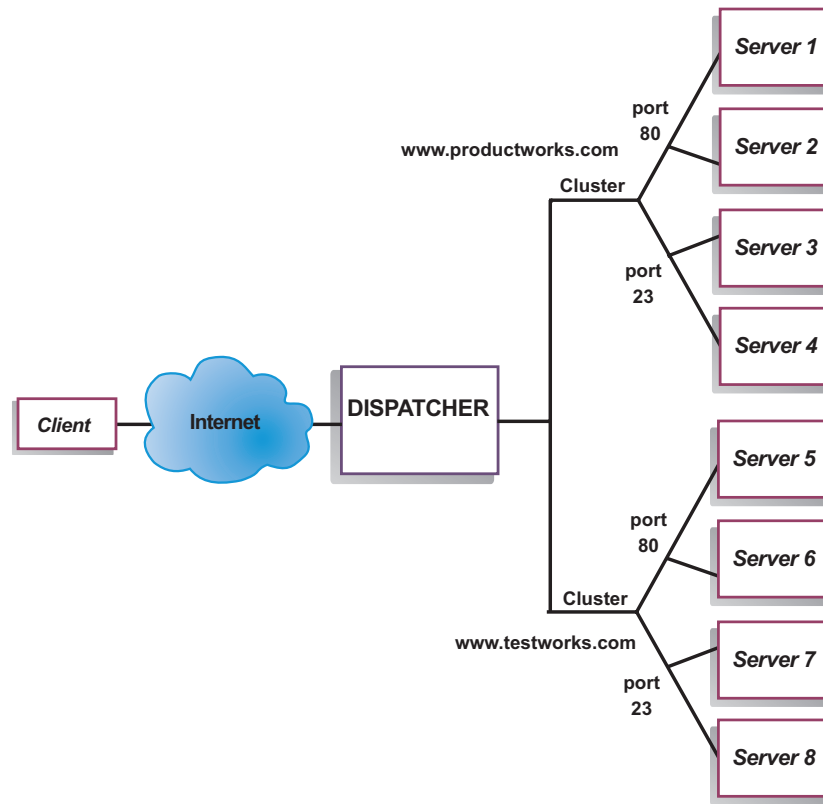
2 clusters, each with 1 port

Another way of configuring Load Balancer might be appropriate if you have a very large site with many servers dedicated to each protocol supported. In this case, you might want to define a cluster for each protocol with a single port but with many servers.



In this example for the Dispatcher component, two clusters are defined: `www.productworks.com` for port 80 (HTTP) and `www.testworks.com` for port 443 (SSL). A third way of configuring Load Balancer might be necessary if your site does content hosting for several companies or departments, each one coming into your site with a different URL. In this case, you might want to define a cluster for each company or department and then define any ports to which you want to receive connections at that URL, as shown in the configuration for 2 clusters, each with two ports.

2 clusters, each with 2 ports



In this example for the Dispatcher component, two clusters are defined with port 80 for HTTP and port 23 for Telnet for each of the sites at `www.productworks.com` and `www.testworks.com`.

Related tasks:

“Quick start configuration” on page 61

This quick start example shows how to configure three locally attached workstations to load-balance web traffic between two web servers.

GUI: General instructions

Load Balancer graphical user interface.

Start the GUI for the Load Balancer components with the following commands:

Dispatcher GUI

Start `dsserver` by issuing `<install_root>/bin/dsserver`. Issue `<install_root>/bin/lbadmin`.

Content Based Router GUI

Start **cbrserver** by issuing `<install_root>/cbr/servers/bin/cbrserver`.
Issue `<install_root>/cbr/servers/bin/lbadmin`.

Site Selector GUI

Start **sssserver** by issuing `<install_root>/ss/servers/bin/sssserver`. Issue
`<install_root>/ss/servers/bin/lbadmin`.

Start configuration wizards for the Load Balancer components with the following commands:

Dispatcher GUI

Start **dssserver** by issuing `<install_root>/bin/dssserver`. Issue
`<install_root>/bin/dswizard`.

Content Based Router GUI

Start **cbrserver** by issuing `<install_root>/cbr/servers/bin/cbrserver`.
Issue `<install_root>/cbr/servers/bin/cbrwizard`.

Site Selector GUI

Start **sssserver** by issuing `<install_root>/ss/servers/bin/sssserver`. Issue
`<install_root>/ss/servers/bin/sswizard`.

Related tasks:

“Configuring Dispatcher” on page 61

Configure Dispatcher for production environments and realistic test environments.

“Configuring Site Selector” on page 83

Configure Site Selector for production environments and realistic test environments.

“Configuring Content Based Routing” on page 92

Configure Content Based Routing (CBR) for production environments and realistic test environments.

Custom advisor methods and function calls

Use the following advisor methods and function calls in your custom advisors.

Be aware that custom advisors need to have all the required routines. Advisors must have the following base class methods:

- A constructor routine. The constructor calls the base class constructor.
- An `ADV_AdvisorInitialize` method. This method provides a way to perform additional steps after the base class completes its initialization.
- A `getLoad` routine. The base advisor class performs the socket opening; the `getLoad` function only needs to issue the appropriate send and receive requests to complete the advising cycle.

Constructor (provided by advisor base)

```
public <advisor_name> {  
    String sName;  
    String sVersion;  
    int iDefaultPort;  
    int iInterval;  
    String sDefaultLogFileName;  
    boolean replace  
}
```

sName

The name of the custom advisor

sVersion

The version of the custom advisor.

iDefaultPort

The port number on which to contact the server if no port number is specified in the call.

iInterval

The interval at which the advisor will query the servers.

sDefaultLogFileName

This parameter is required but not used. The only acceptable value is a null string, ""

replace

Whether or not this advisor functions in replace mode. Possible values are the following:

- *true* – Replace the load calculated by the advisor base code with the value reported by the custom advisor.
- *false* – Add the load value reported by the custom advisor to the load value calculated by the advisor base code.

ADV_AdvisorInitialize() method

```
void ADV_AdvisorInitialize()
```

This method is provided to perform any initialization that might be required for the custom advisor. This method is called after the advisor base module starts. In many cases, including the standard advisors, this method is not used and its code consists of a return statement only. This method can be used to call the “suppressBaseOpeningSocket()” on page 131 method, which is valid only from within this method.

In many cases, including the standard advisors, this method is not used and its code consists of a return statement only. You can use this method to call the suppressBaseOpeningSocket method, which is valid only from within the ADV_AdvisorInitialize method.

ADVLOG() method

The ADVLOG function allows a custom advisor to write a text message to the advisor base log file. The format follows:

```
void ADVLOG (int logLevel, String message)
```

This command has the following parameters:

logLevel

The status level at which the message is written to the log file. The advisor log file is organized in stages; the most urgent messages are given status level 0 and less urgent messages receive higher numbers. The most verbose type of message is given status level 5. These levels are used to control the types of messages that the user receives in real time (The dscontrol command is used to set verbosity). Catastrophic errors should always be logged at level 0.

message

The message to write to the log file. The value for this parameter is a standard Java string.

getAdvisorName function

The `getAdvisorName` function returns a Java string with the suffix portion of your custom advisor name. For example, for an advisor named `ADV_cdload.java`, this function returns the value `cdload`.

This function does not take parameters.

Note: It is not possible for this value to change during one instantiation of an advisor.

caller.getCurrentServerId()

The `getCurrentServerId` function returns a Java string which is a unique representation for the current server. Typically, this value changes each time you call your custom advisor, because the advisor base code queries all server machines in series.

This function takes no parameters.

caller.getCurrentClusterId()

The `getCurrentClusterId` function call returns a Java string which is a unique representation for the current cluster. Typically, this value changes each time you call your custom advisor, because the advisor base queries all clusters in series.

This function takes no parameters.

caller.getSocket()

The `getSocket` function call returns a Java socket which represents the socket opened to the current server for communication.

This function takes no parameters.

caller.getLatestLoad()

The `getLatestLoad` function allows a custom advisor to obtain the latest load value for a given server object. The load values are maintained in internal tables by the advisor base code and the manager daemon. This function call is useful if you want to make the behavior of one protocol or port dependent on the behavior of another. For example, you might use this function call in a custom advisor that disabled a particular application server if the Telnet server on that same machine was disabled.

The syntax is:

```
int caller.getLatestLoad (String clusterId, int port, String serverId)
```

The three arguments together define one server object.

This command has the following parameters:

clusterId

The cluster identifier of the server object for which to obtain the current load value. This argument must be a Java string.

port

The port number of the server object for which to obtain the current load value.

serverId

The server identifier of the server object for which to obtain the current load value. This argument must be a Java string. The return value is an integer.

- A positive return value represents the actual load value assigned for the object that was queried.
- The value -1 indicates that the server asked about is down.
- The value -2 indicates that the status of the server asked about is unknown.

caller.receive()

The receive function gets information from the socket connection. The syntax is:
`caller.receive(StringBuffer *response)`

This command has the following parameters:

response

This is a string buffer into which the retrieved data is placed. Additionally, the function returns an integer value with the following significance:

- 0 indicates data was sent successfully.
- A negative number indicates an error.

caller.send()

The send function uses the established socket connection to send a packet of data to the server, using the specified port. The syntax is as follows:

`caller.send(String command)`

This command has the following parameters:

command

This is a string containing the data to send to the server. The function returns an integer value with the following significance:

- 0 indicates data was sent successfully.
- A negative number indicates an error.

getLoad()

`int getLoad(int iConnectTime; ADV_Thread *caller)`

This function has the following parameters:

iConnectTime

The length of time, in milliseconds, that it took the connection to complete. This load measurement is performed by the advisor base code and passed to the custom advisor code, which can use or ignore the measurement when returning the load value. If the connection fails, this value is set to -1.

caller

The instance of the advisor base class where advisor base methods are provided. Function calls available to custom advisors: The methods, or functions, described in the following sections can be called from custom advisors. These methods are supported by the advisor base code. Some of these function calls can be made directly, for example, `function_name()`, but

others require the prefix caller. Caller represents the base advisor instance that supports the custom advisor that is being executed.

getAdviseOnPort()

The getAdviseOnPort function returns the port number on which the calling custom advisor is running.

The return value is a Java integer (int), and the function does not take parameters.

Note: It is not possible for this value to change during one instantiation of an advisor.

getAdvisorName()

The getAdvisorName function returns a Java string with the suffix portion of your custom advisor's name. For example, for an advisor named ADV_cdload.java, this function returns the value cdload. This function takes no parameters. Note that it is not possible for this value to change during one instantiation of an advisor.

getInterval()

The getInterval function returns the advisor interval, that is, the number of seconds between advisor cycles. This value is equal to the default value set in the custom advisor's constructor, unless the value has been modified at run time by using the dscontrol command. The return value is a Java integer (int).

The function takes no parameters.

suppressBaseOpeningSocket()

The suppressBaseOpeningSocket function call allows a custom advisor to specify whether the base advisor code opens a TCP socket to the server on the custom advisor's behalf. If your advisor does not use direct communication with the server to determine its status, it might not be necessary to open this socket. This function call can be issued only once, and it must be issued from the "ADV_AdvisorInitialize() method" on page 128 routine.

The function takes no parameters.

Related tasks:

“Creating a custom advisor” on page 125

A custom advisor is a small piece of Java code, provided as a class file, that is called by the Load Balancer base code to determine the load on a server. The base code provides all necessary administrative services, including starting and stopping an instance of the custom advisor, providing status and reports, recording history information in a log file, and reporting advisor results to the manager component.

Related reference:

“Example: Sample advisor” on page 132

This is a sample advisor file called ADV_sample.

List of advisors

Advisors are agents within Load Balancer. Their purpose is to assess the health and loading of server machines. This list of advisors are already provided with Load Balancer, but you can also write a custom advisor to suit specific needs.

Table 13. List of advisors

Advisor Name	Description
caching proxy	The caching proxy advisor opens a connection, sends a caching proxy specific HTTP GET request, and interprets the response as a caching proxy load. Important: When using the Caching Proxy advisor, Caching Proxy needs to be running on all servers being load balanced. The machine on which the Load Balancer resides does not need to have Caching Proxy installed unless it is collocated on the same machine it is load balancing.
connect	The connect advisor does not exchange any protocol-specific data with the server. It simply measures the time it takes to open and close a TCP connection with the server. This advisor is useful for server applications which use TCP, but with a higher-level protocol for which an IBM-supplied or custom advisor is not available.
Custom advisors	Dispatcher provides the ability for a customer to write a custom (customizable) advisor. This enables support for proprietary protocols (on top of TCP) for which IBM has not developed a specific advisor. For more information, see “Creating a custom advisor” on page 125.
db2	The DB2 advisor works in conjunction with the DB2 servers. Dispatcher has the built in capability of checking the health of DB2 servers without the need for customers to write their own custom advisors. The DB2 advisor communicates with the DB2 connection port only, not the Java connection port.

Table 13. List of advisors (continued)

Advisor Name	Description
dns	The DNS advisor opens a connection, sends a pointer query for DNS, waits for a response, closes the connection and returns the elapsed time as a load.
ftp	The FTP advisor opens a connection, sends a SYST request, waits for a response, closes the connection, and returns the elapsed time as a load.
http	The HTTP advisor opens a connection, sends a HEAD request by default, waits for a response connection, and returns the elapsed time as a load. See “Getting service-specific advice with the advisor request or response option” on page 118 for more information on how to change the type of request sent by the HTTP advisor.
https	<p>The HTTPS advisor is a heavyweight advisor for SSL connections. It performs a full SSL socket connection with the server. The HTTPS advisor opens an SSL connection, sends an HTTPS request, waits for a response, closes the connection, and returns the elapsed time as a load. (See also the SSL advisor, which is a lightweight advisor for SSL connections.)</p> <p>Note: The HTTPS advisor has no dependency upon server key or certificate content, but they must not be expired.</p>
imap	The IMAP advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.
ldap	The LDAP advisor opens a connection, sends an anonymous BIND request, waits for a response, closes the connection, and returns the elapsed time as a load.
ldapuri	<p>Note: The LDAP URI advisor allows you better gauge LDAP availability by processing a complete request to the LDAP server.</p> <p>The advisor:</p> <ol style="list-style-type: none"> 1. Opens a connection. 2. Sends a BIND request, which is based on the advisorrequest field that you define on the server object. 3. Waits for a response. 4. Closes the connection. 5. Returns the elapsed time as a load. <p>Read “Configuring the LDAP URI advisor” on page 119 for more information on configuring this advisor.</p>

Table 13. List of advisors (continued)

Advisor Name	Description
nntp	The NNTP advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.
ping	The ping advisor does not open a TCP connection with the servers, but instead reports whether the server responds to a ping. While the ping advisor may be used on any port, it is also designed for configurations using the wildcard port, over which multiple protocol traffic may be flowing. It is also useful for configurations using non-TCP protocols with their servers.
pop3	The POP3 advisor opens a connection, waits for an initial message from the server, sends a quit command, closes the connection, and returns the elapsed time as a load.
reach	The reach advisor pings its target machines. This advisor is also designed for the Dispatcher's high availability components to determine reachability of its reach targets. Its results flow to high availability component and do not appear in the manager report. Unlike the other advisors, the reach advisor starts automatically by the manager function of the Dispatcher component.
sip	The SIP advisor opens a connection, sends an OPTIONS request, waits for a response, closes the connection, and returns the elapsed time as a load. The SIP advisor that is supported runs on TCP only and requires an application to be installed on a server that responds to an OPTIONS request.
smtp	The SMTP advisor opens a connection, waits for an initial message from the server, sends a quit, closes the connection, and returns the elapsed time as a load.

Table 13. List of advisors (continued)

Advisor Name	Description
ssl	<p>The SSL advisor is a <i>lightweight</i> advisor for SSL connections. It does not establish a full SSL socket connection with the server. The SSL advisor opens a connection, sends an SSL CLIENT_HELLO request, waits for a response, closes the connection, and returns the elapsed time as a load. (See also the HTTPS advisor, which is a <i>heavyweight</i> advisor for SSL connections.)</p> <p>Note: Be aware of the following:</p> <ul style="list-style-type: none"> • The SSL advisor has no dependency upon key management or certificates. • You might want to use the TLS advisor if you find that the SSL advisor is marking servers down, you know that the server is still up, and you see messages stating "No ciphers found."
ssl2http	<p>The ssl2http advisor starts and advises on the servers listed under port 443, but the advisor will open a socket to the <i>mapport</i> for HTTP requests. Only use the ssl2http advisor for CBR if the client-to-proxy protocol list is SSL and the proxy-to-server protocol is HTTP.</p>
telnet	<p>The Telnet advisor opens a connection, waits for an initial message from the server, closes the connection, and returns the elapsed time as a load.</p>
tls	<p>Transport Layer Security (TLS) is the successor to the SSL protocol. TLS 1.0 was the follow up to SSLv3.</p> <p>You might want to use the TLS advisor if you find that the SSL advisor is marking servers down, you know that the server is still up, and you see messages stating "No ciphers found."</p>
was	<p>The WAS (WebSphere Application Server) advisor works in conjunction with the WebSphere Application servers. Customizable sample files for this advisor are provided in the installation directory. For more information, see "Example: Implementing the WAS advisor" on page 142.</p>
wlm	<p>The WLM (Workload Manager) advisor is designed to work in conjunction with servers on OS/390 mainframes running the MVS™ Workload Manager (WLM) component. For more information, see "The Workload Management Advisor" on page 124.</p>

Related tasks:

“Enabling advisors to manage load balancing” on page 111

Advisors are software agents that work within Load Balancer to provide information about the load on a given server. A different advisor exists for each standard protocol (HTTP, SSL, and others). Periodically, the Load Balancer base code performs an advisor cycle, during which it individually evaluates the status of all servers in its configuration.

“Configuring the LDAP URI advisor” on page 119

The LDAP URI advisor allows you better gauge Lightweight Directory Access Protocol (LDAP) availability by processing a complete request to the LDAP server. The LDAP URI advisor opens a connection to the LDAP server and sends a BIND request that is based on the `advisorrequest` field that you define on the server object. The advisor then waits for a response from the LDAP server and returns the elapsed time as a load.

Sample scripts to generate alerts and record server failure

Load Balancer provides user exits that trigger scripts that you can customize. You can create the scripts to perform automated actions, such as alerting an Administrator when servers are marked down by the manager or simply record the event of the failure.

Sample scripts, which you can customize, are in the `install_root/servers/samples` directory. In order to run the files, you must move them to the `install_root/servers/bin` directory and remove the `△sample△` file extension. The following sample scripts are provided:

- **serverDown** — a server is marked down by the manager.
- **serverUp** — a server is marked back up by the manager.
- **managerAlert** — all servers are marked down for a particular port.
- **managerClear** — at least one server is now up, after all were marked down for a particular port.

If all servers on a cluster are marked down (either by the user or by the advisors), the `managerAlert` (if configured) starts, and Load Balancer attempts to route traffic to the servers using a round-robin technique. The `serverDown` script does not start when the last server in the cluster is detected as offline. By design, Load Balancer attempts to continue to route the traffic in case a server comes back online and responds to the request. If Load Balancer instead dropped all traffic, the client would receive no response. When Load Balancer detects that the first server of a cluster is back online, the `managerClear` script (if configured) starts, but the `serverUp` script (if configured) is not run until an additional server is brought back online.

Here are some considerations for using the `serverUp` and `serverDown` scripts:

- If you define the manager cycle to be less than 25% of the advisor time, false reports of servers up or down can result. By default, the manager runs every 2 seconds, but the advisor runs every 7 seconds. Therefore, the manager expects new advisor information within 4 cycles. However, removing this restriction (that is, defining the manager cycle to be greater than 25% of the advisor time) significantly decreases performance because multiple advisors can advise on a single server.

- When a server goes down, the serverDown script starts. However, if you issue a serverUp command, it is assumed that the server is up until the manager obtains new information from the advisor cycle. If the server is still down, the serverDown script runs again.

High Availability recovery strategy for failed servers

The recovery strategy dictates how Load Balancer behaves when one Dispatcher machine fails and there is another configured as a backup.

Two Dispatcher machines are configured: the primary machine, and a second machine called the backup. At startup, the primary machine sends all the connection data to the backup machine until that machine is synchronized. The primary machine becomes active, that is, it begins load balancing. The backup machine, meanwhile, monitors the status of the primary machine, and is said to be in standby state.

If the backup Load Balancer machine detects that the primary machine has failed, it performs a takeover load balancing functions and becomes the active machine. After the primary machine has once again become operational, the machines respond according to how the recovery strategy has been configured by the user.

There are two kinds of strategy:

- **Automatic** - The primary machine resumes routing packets as soon as it becomes operational again.
- **Manual** - intervention is required to return the primary machine to active state and reset the backup machine to standby. The manual recovery strategy allows you to force the routing of packets to a particular machine, using the takeover command. Manual recovery is useful when maintenance is being performed on the other machine

Note: The strategy parameter must be the same for both machines.

Related tasks:

“Configuring high availability” on page 146

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the "health" of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

“Detecting server failures with heartbeats and reach targets” on page 149

Configure heartbeats and reach targets to detect server failures and control when failovers can occur.

“Scripts to run with high availability” on page 151

Scripts to run with high availability

The following sample script can be used:

highAvailChange

The HighAvailChange script runs whenever the high availability state changes within the Load Balancer. You can create this script to use state change information, for instance, to alert an Administrator or simply record the event.

On Linux for s/390: Load Balancer issues a gratuitous ARP to move IP addresses from one Load Balancer to another. This mechanism is therefore tied to the underlying network type. When running Linux for s/390, Load Balancer can natively do high availability takeovers, complete with IP address moves, only on those interfaces which can issue a gratuitous ARP and configure the address on the local interface. This mechanism will not work properly on point-to-point interfaces such as IUCV and CTC and will not work properly in certain configurations of qeth/QDIO.

Related tasks:

“Configuring high availability” on page 146

The high availability feature involves the use of a second Dispatcher machine. The first Dispatcher machine performs load balancing for all the client traffic as it does in a single Dispatcher configuration. The second Dispatcher machine monitors the “health” of the first, and takes over the task of load balancing if it detects that the first Dispatcher machine has failed.

“Detecting server failures with heartbeats and reach targets” on page 149

Configure heartbeats and reach targets to detect server failures and control when failovers can occur.

Related reference:

“High Availability recovery strategy for failed servers” on page 150

The recovery strategy dictates how Load Balancer behaves when one Dispatcher machine fails and there is another configured as a backup.

Sample 2-tier high availability configuration with Dispatcher, CBR, and Caching Proxy

Set up a 2-tier, high availability configuration that combines the capabilities of two Load Balancer components (the Dispatcher component and the CBR component) along with Caching Proxy.

Server setup

Example of a 2-tier, high availability configuration that uses Dispatcher, CBR, and Caching Proxy

This figure describes the following scenario:

- EdgeServer1: primary (high availability) Dispatcher machine that is collocated with CBR and Caching Proxy that load balances across web servers
- EdgeServer2: standby (high availability) Dispatcher machine that is collocated with CBR and Caching Proxy
- EdgeServer3: CBR and Caching Proxy server
- WebServerA, WebServerB, WebServerC: backend web servers

The figure shows a basic representation of multiple servers (EdgeServer1, EdgeServer2, EdgeServer3) load balancing across multiple backend web servers. The CBR component uses Caching Proxy to forward requests that are based on the content of the URL to the backend web servers. The Dispatcher component is used to load balance the CBR components across the EdgeServers. The high availability feature of the Dispatcher component is used to ensure that requests to the backend servers continue even if the primary high availability machine (EdgeServer1) fails at any time.

Basic configuration guidelines

- Configure Caching Proxy to be the same on all the EdgeServers. To improve the overall accessibility to the web pages on the backend servers, set up Caching Proxy to do memory caching. Setting up Caching Proxy enables the EdgeServers to cache web pages that are requested more frequently. For more information about setting up Caching Proxy, refer to the Caching Proxy Administration Guide.
 - Define the cluster address and ports to be the same in both the CBR and Dispatcher components of Load Balancer.
 - Configure the CBR component to be the same across all EdgeServers. Use web servers A, B, and C as your servers on the ports you want to define for the cluster. For more information to configure CBR, read the Configuring Content Based Routing topics.
 - Configure the Dispatcher component to be the same on EdgeServer1 and EdgeServer2. Define all the EdgeServers as your servers on the ports you want to be defined on the cluster to be load balanced by Dispatcher. For more information about how to configure Dispatcher, read the Configuring Dispatcher topics.
 - Configure EdgeServer1 as the primary high availability machine and EdgeServer2 as the standby (backup) high availability machine. For more information, read the High Availability topic.
1. To avoid backend server addresses displayed in the URL on a client, set the ReversePass directive for each backend server address in the Caching Proxy configuration file.
 2. To ensure that web memory caching is being used effectively, set the Caching directive to ON and increase the CacheMemory directive to the size required in the Caching Proxy configuration file.
 3. Sample lines referred to in steps 1-2:

```
Caching          ON
CacheMemory      128000 K
ReversePass /* http://websrvA.company.com/* http://www.company.com/*
```
 4. Alias the cluster address on the network interface card for EdgeServer1 and to alias the cluster address on the loopback device on the remaining EdgeServers.
 5. If using the Linux platform for the EdgeServers, you might install a patch to the Linux kernel or use an alternative to aliasing the loopback device. .
 6. For CBR, do not use port affinity (stickytime) when you use content rules because the content rules do not fire while processing requests to the backend web servers.

Sample configuration files

The following sample configuration files are similar to files that are created when you set up an Edge Components configuration as shown in the figure. The sample configuration files represent the files for the Dispatcher and CBR components of Load Balancer. In the sample configuration, a single Ethernet adapter is used for each of the EdgeServer machines and all addresses are represented within a private subnet. The sample configuration files use the following IP addresses for the specified machines:

- EdgeServer1 (Primary high availability EdgeServer): 192.168.1.10
- EdgeServer2 (Backup high availability EdgeServer): 192.168.1.20
- EdgeServer3 (web caching EdgeServer): 192.168.1.30
- Website cluster address: 192.168.1.11

- WebServers A-C (Backend web servers): 192.168.1.71, 192.168.1.72, and 192.168.1.73

Sample configuration file for Dispatcher component on primary high availability EdgeServer

```
dscontrol executor start

dscontrol cluster add 192.168.1.11 primaryhost 192.168.1.10

dscontrol port add 192.168.1.11@80

dscontrol server add 192.168.1.11@80@edgeserver1 address 192.168.1.10
dscontrol server add 192.168.1.11@80@edgeserver2 address 192.168.1.20
dscontrol server add 192.168.1.11@80@edgeserver3 address 192.168.1.30

dscontrol manager start manager.log 10004

dscontrol highavailability heartbeat add 192.168.1.10 192.168.1.20
dscontrol highavailability backup add primary auto 4567
```

Sample configuration file for CBR component on the EdgeServers

```
cbrcontrol set loglevel 1
cbrcontrol executor start

cbrcontrol cluster add 192.168.1.11

cbrcontrol port add 192.168.1.11@80

cbrcontrol server add 192.168.1.11@80@webserverA address 192.168.1.71
cbrcontrol server add 192.168.1.11@80@webserverB address 192.168.1.72
cbrcontrol server add 192.168.1.11@80@webserverC address 192.168.1.73

cbrcontrol rule add 192.168.1.11@80@webA_rule type content
  pattern (URI=*WSA*)|(URI=*wsA*) priority 21
cbrcontrol rule useserver 192.168.1.11@80@webA_rule webserverA

cbrcontrol rule add 192.168.1.11@80@webB_rule type content
  pattern (URI=/WS_B*) priority 22
cbrcontrol rule useserver 192.168.1.11@80@webB_rule webserverB

cbrcontrol rule add 192.168.1.11@80@webC_rule type content
  pattern URI=*webC* priority 23
cbrcontrol rule useserver 192.168.1.21@80@webC_rule webserverC
```

Related concepts:

“High availability with Load Balancer” on page 21

The Dispatcher component offers a built-in high availability feature, eliminating Dispatcher as a single point of failure from your network. This feature involves the use of a second Dispatcher machine that monitors the main, or primary, machine and stands by to take over the task of load balancing should the primary machine fail at any time.

Related information:

“Configuring Content Based Routing” on page 92

Configure Content Based Routing (CBR) for production environments and realistic test environments.

“Configuring Dispatcher” on page 61

Configure Dispatcher for production environments and realistic test environments.

Content rule (pattern) syntax

Learn how to use the content rule (pattern) syntax for the CBR component, along with scenarios and examples of its usage.

Content rule (pattern) syntax

Attention: Applicable only if you selected content for the rule type.

Enter the pattern syntax that you want to use, with the following restrictions:

- No spaces can be used within the pattern
- Special characters, unless you precede the character with a backward slash (\):
 - * Wildcard (matches 0 to x of any character)
 - (left parenthesis that is used for logic grouping
 -) right parenthesis that is used for logic grouping
 - & Logical AND
 - | Logical OR
 - ! Logical NOT

Reserved keywords

Reserved keywords are always followed by an equal sign “=”.

Method

HTTP method in the request, for example GET, POST

URI Path of the URL request (case sensitive)

Version

Specific version of request, either HTTP/1.0 or HTTP/1.1

Host Value from the host: header (not case-sensitive)

Important: Optional in HTTP/1.0 protocols

<key> Any valid HTTP header name that Dispatcher can search for. Examples of HTTP headers are User-Agent and Connection, Referrer

A browser that targets `http://www.company.com/path/webpage.htm` might result in values, as follows:

```
Method=GET
URI=/path/webpage.htm
Version=HTTP/1.1
Host=www.company.com
Connection=Keep-Alive
Referer=http://www.company.com/path/parentwebpage.htm
```

The operating system shell might interpret special characters, such as “&”, and convert them to alternative text before **cbrcontrol** evaluates them. If you are entering the command from the **dscontrol**, **cbrcontrol**, or a configuration file, use double quotation marks (“ ”) around the special characters. For example, the following command is valid only when you use the **cbrcontrol>>** prompt or from a configuration file:

```
rule add 10.1.203.4@80@cbr_prod_rule_ek type content
  pattern "uri=/nipoeck/*"
```

For this command to work at the operating system prompt when you use special characters, place double quotation marks around the entire command:

```
cbrcontrol "rule add 10.1.203.4@80@cbr_prod_rule_ek type content
  pattern uri=/nipoeck/*"
```

If the quotation marks are not used, some of the pattern might be truncated when the rule is saved in CBR.

Scenarios and examples for using pattern syntaxes

Scenario 1

The setup for one cluster name involves one set of web servers for standard HTML content, another set of web servers with WebSphere Application Server for servlet requests, and another set of Lotus® Notes® servers for NSF files. Access to the client data is required to distinguish between those requested pages. It is also required to send them to the appropriate servers. The content pattern matching rules provide the separation that is needed to accomplish these tasks. A series of rules are configured so that the necessary separation of requests occurs automatically. For example, the following commands accomplish the three splits mentioned:

```
>>rule add cluster1@80@servlets type content pattern "uri=*/servlet/*" priority 1
>>rule uses cluster1@80@servlets server1+server2
>>rule add cluster1@80@notes type content pattern "uri=*.nsf*" priority 2
>>rule uses cluster1@80@notes server3+server4
```

If a request for an NSF file arrives at Load Balancer, the servlets rule is checked first, but does not match. The request is then checked by the notes rule and returns a match. The client is load-balanced between server3 and server4.

Scenario 2

Another common scenario is when the main website controls several distinct internal groups. For example, `www.company.com/software` involves a different set of servers and content from `www.company.com/hardware` division. Because the requests are all based off the root `www.company.com` cluster, content rules are required to find the URI differences and complete load balancing. The scenario's rule looks similar to the following examples:

```
>>rule add cluster1@80@div1 type content pattern "uri=/software/*" priority 1
>>rule uses cluster1@80@div1 server1+server2
```

```
>>rule add cluster1@80@div2 type content pattern "uri=/hardware/*" priority 2
>>rule uses cluster1@80@div2 server3+server4
```

Scenario 3

Certain combinations are sensitive to the order in which rules are searched. For example, in Scenario 2, clients were split based on a directory in their request path; however, the target directory might display at multiple levels of the path and mean different things on placement. For example, `www.company.com/pcs/fixed/software` is a different target from `www.company.com/mainframe/fixed/software`. Define the rules to account for this possibility and not catch too many scenarios at the same time. For example, the `uri=*/software/*` test is too broad of a wildcard search in this case. Alternative rules might be structured as follows:

- A combination search can narrow this search down:

```
>>rule add cluster1@80@pcs type content pattern "(uri=/pcs*)&(uri=*/software/*)"
>>rule uses cluster 1@80@pcs server1
```
- In cases where there are no combinations to use, the order becomes important:

```
>>rule add cluster1@80@pc1 type content pattern "uri=/pcs/*"
>>rule uses cluster1@80@pc1 server2
```
- The second rule catches when “pcs” displays in later directory spots instead of the first.

```
>>rule add cluster1@80@pc2 type content pattern "uri=*/pcs/*"
>>rule uses cluster1@80@pc2 server3
```
- In almost every case, you want to complete the rules with a default always true rule to catch anything that falls through the other rules. This rule can also be a “Sorry, the site is currently down, please try again later” server for scenarios where all other servers fail for this client.

```
>>rule add cluster1@80@sorry type true priority 100
>>rule uses cluster1@80@sorry server5
```

Commands

Look up a command by its name to find detailed syntax and usage of the command.

Content Based Routing (CBR) uses a subset of the Dispatcher commands listed in this command reference. When using these syntax diagrams for CBR, substitute **cbrcontrol** for **dscontrol**. The dispatcher commands that are omitted in CBR are as follows:

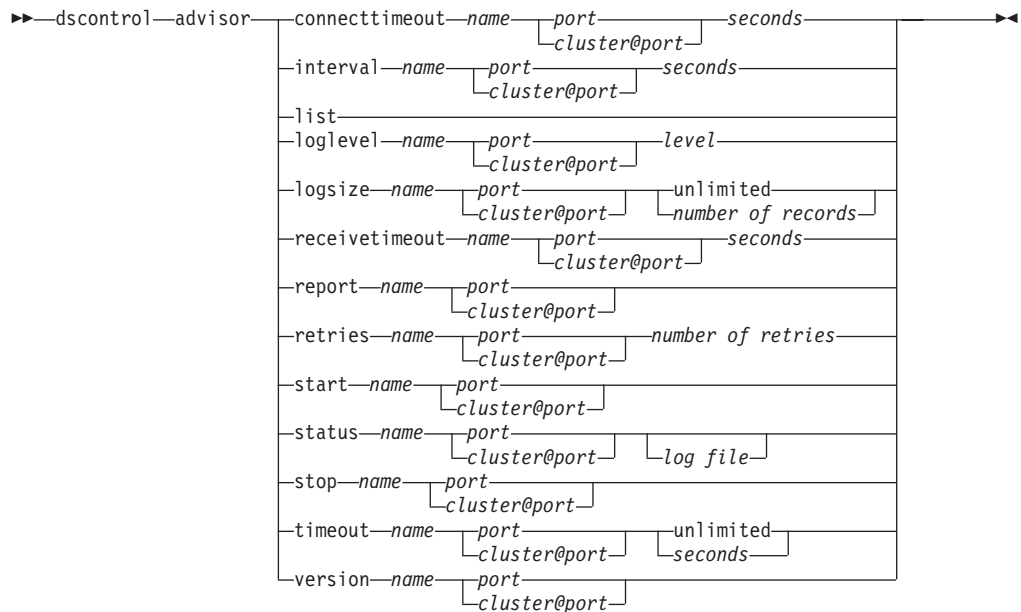
- **highavailability**
- **subagent**
- **executor**
 - report
 - set nfa <value>
 - set fintimeout <value>
 - set hatimeout <value>
 - set hasynctimeout <value>
 - set porttype <value>
- **cluster**
 - report {c}
 - set {c} porttype
- **port**

- add {c@p} porttype
- add {c@p} protocol
- set {c@p} porttype
- **rule add {c@p@r} type port**
- **server**
 - add {c@p@s} router
 - set {c@p@s} router

dscontrol advisor

Use this command to control various features of the advisor function.

Syntax



Parameters

connecttimeout

Set how long an advisor waits before reporting that a connect to a server for a particular port on a server (a service) fails. For more information, see “Enabling advisors to manage load balancing” on page 111.

- *name*

The name of the advisor. Possible values include **connect**, **db2**, **dns**, **ftp**, **http**, **https**, **cachingproxy**, **imap**, **ldap**, **ldapuri**, **nntp**, **ping**, **pop3**, **self**, **sip**, **smtp**, **ssl**, **ssl2http**, **telnet**, and **wlm**.

See the topic “List of advisors” on page 114 for more information on the advisors that Load Balancer provides.

Names of customized advisors are of the format **ADV_xxxx**, where **xxxx** is the name of the class that implements the custom advisor. See “Creating a custom advisor” on page 125 for more information.

- *port*

The number of the port that the advisor is monitoring.

- *cluster@port*

The cluster value is optional on the advisor commands, but the port value is required. If the cluster value is not specified, then the advisor will start running on the port for all clusters. If you specify a cluster, then the advisor will start running on the port, but only for the cluster you have specified. See the topic “Enabling advisors to manage load balancing” on page 111 for more information on starting and stopping advisors.

The cluster is the address in IP address format or symbolic name. The port is the number of the port that the advisor is monitoring.

- *seconds*

A positive integer representing the timeout in seconds at which the advisor waits before reporting that a connect to a server fails. The default is 3 times the value specified for the advisor interval.

interval

Set how often the advisor will query the servers for information.

- *seconds*

A positive integer representing the number of seconds between requests to the servers about their current status. The default is 7.

list

Show list of advisors that are currently providing information to the manager.

loglevel

Set the logging level for an advisor log.

- *level*

The number of the level (0 to 5). The default is 1. The higher the number, the more information that is written to the advisor log. The following are the possible values: 0 is None, 1 is Minimal, 2 is Basic, 3 is Moderate, 4 is Advanced, 5 is Verbose.

logsize

Set the maximum size of an advisor log. When you set a maximum size for the log file, the file will wrap; when the file reaches the specified size, the subsequent entries are written from the top of the file, overwriting the previous log entries. Log size cannot be set smaller than the current size of the log. Log entries are time-stamped so you can tell the order in which they were written. The higher you set the log level, the more carefully you should choose the log size, because you can quickly run out of space when logging at the higher levels.

- *number of records*

The maximum size in bytes for the advisor log file. You can specify either a positive number greater than zero, or the word unlimited. The log file may not reach the exact maximum size before overwriting because the log entries themselves vary in size. The default value is 1 MB.

receivetimeout

Set how long an advisor waits before reporting that a receive from a particular port on a server (a service) fails. For more information, see “Enabling advisors to manage load balancing” on page 111.

- *seconds*

A positive integer representing the timeout in seconds at which the advisor waits before reporting that a receive from a server fails. The default is 3 times the value specified for the advisor interval.

report

Display a report on the state of the advisor.

retry

Retry sets the number of retries that an advisor can make before marking a server down.

- *number of retries*

An integer greater than or equal to zero. This value should be no larger than 3. If retries keyword is not configured, the number of retries defaults to zero.

start

Start the advisor. There are advisors for each protocol. The default ports are as follows:

Table 14. Default ports for advisors

Advisor Name	Protocol	Port
connect	ICMP	12345
db2	private	50000
dns	DNS	53
ftp	FTP	21 Note: The FTP advisor should advise only on the FTP control port (21). Do not start an FTP advisor on the FTP data port (20).
http	HTTP	80
https	SSL	443
imap	IMAP	143
ldap	LDAP	389
ldapuri	LDAP	389
nntp	NNTP	119
ping	PING	0
pop3	POP3	110
self	private	12345
sip	SIP	5060
smtp	SMTP	25
ssl	SSL	443
ssl2http	SSL	443
telnet	Telnet	23
WLM	private	10007

- *log file*

File name to which the management data is logged. Each record in the log is time-stamped.

The default file is `advisorname_port.log`, for example, `http_80.log`. To change the directory where the log files are kept, see “Logging with Load Balancer” on page 158. The default log files for cluster (or site) specific advisors are created with the cluster address, for example, `http_127.40.50.1_80.log`.

status

Display the current status of all the values in an advisor that can be set globally and their defaults.

stop

Stop the advisor.

timeout

Set the number of seconds for which the manager will consider information from the advisor as valid. If the manager finds that the advisor information is older than this timeout period, the manager will not use that information in determining weights for the servers on the port the advisor is monitoring. An exception to this timeout is when the advisor has informed the manager that a specific server is down. The manager will use that information about the server even after the advisor information has timed out.

- *seconds*

A positive number representing the number of seconds, or the word unlimited. The default value is unlimited.

version

Display the current version of the advisor.

Samples

- To start the http advisor on port 80 for cluster 127.40.50.1:
`dscontrol advisor start http 127.40.50.1@80`
- To start the http advisor on port 88 for all clusters:
`dscontrol advisor start http 88`
- To stop the http advisor at port 80 for cluster 127.40.50.1:
`dscontrol advisor stop http 127.40.50.1@80`
- To set the time (30 seconds) an HTTP advisor for port 80 waits before reporting that a connect to a server fails:
`dscontrol advisor connecttimeout http 80 30`
- To set the time (20 seconds) an HTTP advisor for port 80 on cluster 127.40.50.1 waits before reporting that a connect to a server fails:
`dscontrol advisor connecttimeout http 127.40.50.1@80 20`
- To set the interval for the FTP advisor (for port 21) to 6 seconds:
`dscontrol advisor interval ftp 21 6`
- To display the list of advisors currently providing information to the manager:
`dscontrol advisor list`
- To change the log level of the advisor log to 0 for better performance:
`dscontrol advisor loglevel http 80 0`
- To change the ftp advisor log size for port 21 to 5000 bytes:
`dscontrol advisor logsize ftp 21 5000`
- To set the time (60 seconds) an HTTP advisor (for port 80) waits before reporting that a receive from a server fails:
`dscontrol advisor receivetimeout http 80 60`
- To display a report on the state of the ftp advisor (for port 21):
`dscontrol advisor report ftp 21`
- To display the current status of values associated with the http advisor for port 80:
`dscontrol advisor status http 80`
- To set the timeout value for the ftp advisor information on port 21 to 5 seconds:
`dscontrol advisor timeout ftp 21 5`
- To display the current version number of the ssl advisor for port 443:
`dscontrol advisor version ssl 443`

This command produces output similar to the following:
Version: 04.00.00.00 - 07/12/2001-10:09:56-EDT

Related reference:

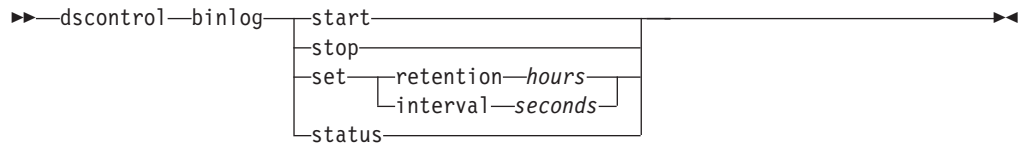
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol binlog

You can control the settings and operation of the binary log file with the dscontrol binlog command.

Syntax



Parameters

start

Starts the binary log.

stop

Stops the binary log.

set

- **retention** *hours*

The number of hours that binary log files are kept. The default value for retention is 24.

- **interval** *seconds*

The number of seconds between log entries. The default value for interval is 60.

status

Shows the retention and intervals of the binary log.

Related reference:

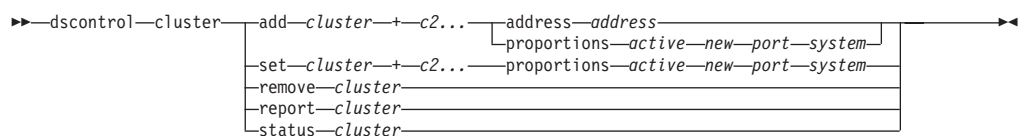
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol cluster

Configure clusters and cluster properties with the dscontrol cluster command.

Syntax



Parameters

add

- *cluster*

The cluster name or address to which clients connect. The cluster value is either a symbolic name or in IP address format.

With the exception of the `dscontrol cluster add` command, you can use an `@` symbol to act as a wild card. For example, the following command will result in setting a weightbound of 80 to all clusters:

```
dscontrol cluster set @ weightbound 80
```

Note: Additional clusters are separated by a plus sign (+).

- **address** *address*

The unique IP address of the TCP machine as either a host name or in IP address format. If the cluster value is unresolvable, you must provide the IP address of the physical machine.

- **proportions**

At the cluster level, set the proportion of importance for active connections (*active*), new connections (*new*), information from any advisors (*port*), and information from a system monitoring program such as Metric Server (*system*) that are used by the manager to set server weights. Each of these values, described below, is expressed as a percentage of the total and they therefore always total 100. For more information see “Tuning the proportion of importance given to status information” on page 178.

- *active*

A number from 0–100 representing the proportion of weight to be given to the active connections. The default is 50.

- *new*

A number from 0–100 representing the proportion of weight to be given to new connections. The default is 50.

- *port*

A number from 0–100 representing the proportion of weight to be given to the information from the advisors. The default is 0.

Note: When an advisor is started and if the port proportion is 0, Load Balancer automatically sets this value to 1 in order for the manager to use the advisor information as input for calculating server weight.

- *system*

A number from 0–100 representing the proportion of weight to be given to the information from the system metrics, such as from Metric Server. The default is 0.

set

Set the properties of the cluster.

remove

Remove this cluster.

report

Show the internal fields of the cluster.

status

Show current status of a specific cluster.

Samples

- To add cluster address 130.40.52.153:
dscontrol cluster add 130.40.52.153
- To remove cluster address 130.40.52.153:
dscontrol cluster remove 130.40.52.153
- To set the relative importance placed on input (active, new, port, system) received by the manager for servers residing on cluster 9.6.54.12:
dscontrol cluster set 9.6.54.12 proportions 60 35 5 0
- To show the status for cluster address 9.67.131.167:
dscontrol cluster status 9.67.131.167

Related reference:

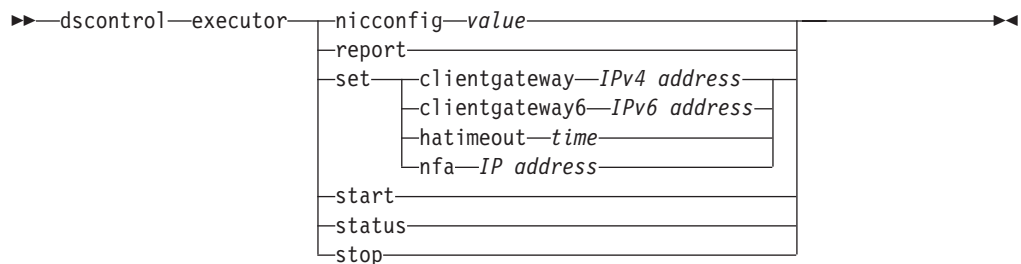
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol executor

Control the executor function with the dscontrol executor command.

Syntax



Parameters

nicconfig

Use this command to set network configuration update mode. For more information see page <page number of section "Automatic network configuration update".

- *value* can be:

auto Specifies to automatically determine the network configuration changes.

manual Specifies to manually determine the network configuration changes.

update Specifies to update the network configuration. This command should only be used when the mode is set to manual.

report

Display a statistics snapshot report.

set

Set the fields of the executor.

- **clientgateway** *IPv4 address*

Clientgateway is an IPv4 address that is used for routing. It is the router address through which traffic in the return direction is forwarded from Load Balancer to clients for IPv4 traffic. Clientgateway must be set to a nonzero value.

- **clientgateway6** *IPv6 address*

Clientgateway6 is an IPv6 address that is used for routing. It is the router address through which traffic in the return direction is forwarded from Load Balancer to clients for IPv6 traffic. Clientgateway6 must be set to a nonzero value.

- **hatimeout** *seconds*

The number of seconds that the executor uses to timeout high availability heartbeats. The default value is 2.

- **nfa** *IP address*

Set the non-forwarding address. Any packet sent to this address will not be forwarded by the Dispatcher machine.

The Internet Protocol address as either a symbolic name or in dotted decimal format.

start

Start the executor.

status

Display the current status of the values in the executor that can be set and their defaults.

stop

Stop the executor.

Samples

- To display the internal counters for Dispatcher:
dscontrol executor status
- To set the non-forwarding address to 130.40.52.167:
dscontrol executor set nfa 130.40.52.167
- To start the executor:
dscontrol executor start
- To stop the executor:
dscontrol executor stop

Related reference:

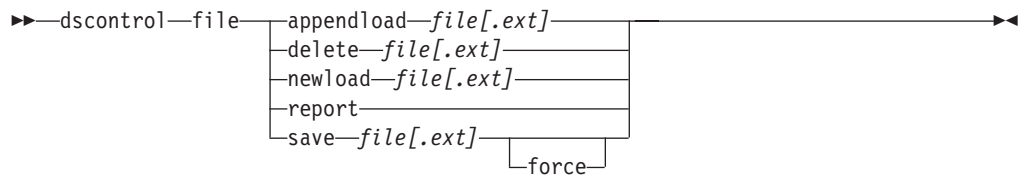
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol file

Manage your configuration files with the dscontrol file command.

Syntax



Parameters

appendload

To update the current configuration, the appendload command runs the executable commands from your script file.

- *file[.ext]*

A configuration file consisting of dscontrol commands. The file extension (.ext) can be anything you like and can be omitted.

delete

Delete the file.

newload

Loads and runs a new configuration file into the Load Balancer. The new configuration file replaces the current configuration.

report

Report on the available file or files.

save

Save the current configuration for Load Balancer to the file.

Note: Files are saved into and loaded from the *install_root/servers/configurations/dispatcher* directory.

- **force**

To save your file to an existing file of the same name, use force to delete the existing file before saving the new file. If you do not use the force option, the existing file is not overwritten.

Samples

- To delete a file:
dscontrol file delete file3
- To load a new configuration file to replace the current configuration:
dscontrol file newload file1.sv
- To append a configuration file to the current configuration and load:
dscontrol file appendload file2.sv
- To view a report of your files (files that you saved earlier):
dscontrol file report
- To save your configuration into a file named file3:
dscontrol file save file3

Related reference:

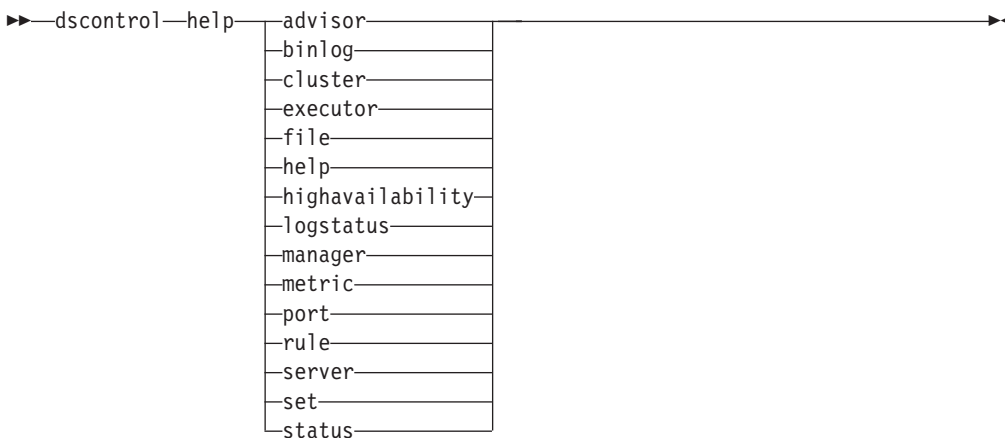
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol help

Display or print help for any dscontrol command with the dscontrol help command.

Syntax



Sample

To get help on the dscontrol command:

```
dscontrol help
```

Related reference:

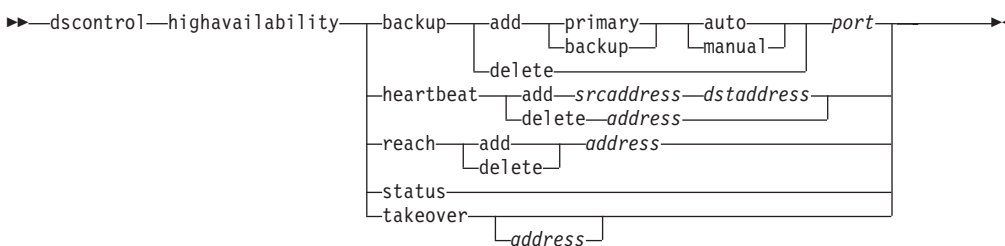
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol highavailability

You can control high availability functions with the dscontrol highavailability command.

Syntax



Parameters

backup

Specify information for either the primary or backup machine.

- **add**

Defines and runs the high availability functions for this machine.

- **primary**

Identifies the Dispatcher machine that has a primary role.

- **backup**

Identifies the Dispatcher machine that has a backup role.

- **auto**

Specifies an automatic recovery strategy, in which the primary machine will resume routing packets as soon as it comes back into service.

- **manual**

Specifies a manual recovery strategy, in which the primary machine does not resume routing packets until the administrator issues a takeover command.

- **delete**

Removes this machine from high availability, so that it will no longer be used as a backup or primary machine.

- *port*

An unused TCP port on both machines, to be used by Dispatcher for its heartbeat messages. The port must be the same for both the primary and backup machines.

heartbeat

Defines a communication session between the primary and backup Dispatcher machines.

- **add**

Tell the source Dispatcher the address of its partner (destination address).

- *source_address*

Source address. The address (IP or symbolic) of this Dispatcher machine.

- *destination_address*

Destination address. The address (IP or symbolic) of the other Dispatcher machine.

The *source_address* and *destination_address* must be the NFAs of the machines for at least one heartbeat pair.

- **delete address**

Removes the address pair from the heartbeat information. You can specify either the destination or source address of the heartbeat pair. The address (IP or symbolic) of either the destination or the source.

reach

Add or delete target address for the primary and backup Dispatchers, the reach advisor sends out pings from both the backup and the primary Dispatchers to determine how reachable their targets are.

- **add address**

Adds a target address for the reach advisor. *address* is the IP address, format or symbolic, of the target node.

- **delete address**

Removes a target address from the reach advisor. *address* is the IP address, format or symbolic, of the target node.

Note: When configuring the reach target, you must also start the reach advisor. The reach advisor starts automatically when you use the dscontrol manager reach command.

status

Return a report on high availability. Machines are identified as having one of three status conditions or states:

- **Active:** A given machine (either a primary, backup, or both) is routing packets.
- **Standby:** A given machine (either a primary, backup, or both) is not routing packets; it is monitoring the state of an active Dispatcher.
- **Idle:** A given machine is routing packets, and is not trying to establish contact with its partner Dispatcher.

takeover

Simple high availability configuration (role of the Dispatcher machines are either primary or backup).

Takeover instructs a standby Dispatcher to become active and to begin routing packets. This will force the currently active Dispatcher to become standby. The takeover command must be issued on the standby machine and works only when the strategy is manual. The substate must be synchronized.

- *address*

The takeover address value is optional. It should only be used when the role of the machine is both primary and backup (mutual high availability configuration). The address specified is the NFA of the Dispatcher machine which normally routes this cluster's traffic. When there is a takeover of both clusters, specify the Dispatcher's own NFA address.

Note:

- The roles of the machines (primary and backup) do not change. Only their relative status (active or standby) changes.
- There are three possible takeover scripts, which are goActive, goStandby, and goInOp. See "Scripts to run with high availability" on page 151 for more information on these scripts.

Samples

- To check the high availability status of a machine:
`dscontrol highavailability status`
- To add the backup information to the primary machine using the automatic recovery strategy and port 80:
`dscontrol highavailability backup add primary auto 80`
- To add an address that the Dispatcher must be able to reach:
`dscontrol highavailability reach add 9.67.125.18`
- To add heartbeat information for the primary and backup machines
Primary - `highavailability heartbeat add 9.67.111.3 9.67.186.8`
Backup - `highavailability heartbeat add 9.67.186.8 9.67.111.3`
- To tell the standby Dispatcher to become active, forcing the active machine to become standby:
`dscontrol highavailability takeover`

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol logstatus

Use this command to display the log settings for a server.

Syntax

▶▶ dscontrol logstatus ◀◀

Parameters

There are no parameters for this command.

Sample

To display the log status:

```
dscontrol logstatus
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol manager

You can control the manager function with the dscontrol manager command.

Syntax



Parameters

interval

Set how often the manager will update the weights of the servers to the executor, updating the criteria that the executor uses to route client requests.

- *seconds*

A positive number representing in seconds how often the manager will update weights to the executor. The default is 2.

loglevel

Set the logging level for the manager log.

- *level*

The number of the level (0 to 5). The higher the number, the more information that is written to the manager log. The default is 1. The following are the possible values: 0 is None, 1 is Minimal, 2 is Basic, 3 is Moderate, 4 is Advanced, 5 is Verbose.

logsize

Set the maximum size of the manager log. When you set a maximum size for the log file, the file will wrap; when the file reaches the specified size, the subsequent entries are written from the top of the file, overwriting the previous log entries. Log size cannot be set smaller than the current size of the log. Log entries are time stamped so you can tell the order in which they were written. The higher you set the log level, the more carefully you should choose the log size, because you can quickly run out of space when logging at the higher levels.

- *bytes*

The maximum size in bytes for the manager log file. You can specify either a positive number greater than zero, or the word unlimited. The log file may not reach the exact maximum size before overwriting because the log entries themselves vary in size. The default value is 1 MB.

metric set

Sets the **loglevel** and **logsize** for the metric monitor log. The loglevel is the metric monitor logging level (0 - None, 1 - Minimal, 2 - Basic, 3 - Moderate, 4 - Advanced, or 5 - Verbose). The default log level is 1. The log size is the maximum number of bytes to be logged in the metric monitor log file. You can specify either a positive number greater than zero, or unlimited. The default logsize is 1 MB.

quiesce

Specify no more connections to be sent to a server except subsequent new connections from the client to the quiesced server if the connection is designated as sticky and stickytime has not expired. The manager sets the weight for that server to 0 in every port to which it is defined. Use this command if you want to do some quick maintenance on a server and then unquiesce it. If you delete a quiesced server from the configuration and then add it back, it will not retain its status prior to being quiesced. For more information, see "Quiesce servers or services for maintenance windows" on page 156.

- *server*

The IP address of the server as either a symbolic name or in dotted decimal format.

- *cluster@port@server*

When you want to bring down a particular service on a server, like HTTP, you can now quiesce individual daemons. If you have added server1 to port

80 and port 22 of a cluster, for example, the usual quiesce command will stop forwarding traffic to server1 on port 80 and port 22. This new option allows you to stop traffic only on port 80 if you need to stop that particular service.

- **daily** *start_hour end_hour*

Note: This setting specifies to quiesce the server at a time of day, *start_hour*, and unquiesce the server at a later point, *end_hour*. The values for *start_hour* and *end_hour* can range from 0 to 23. For example, (0 0) indicates to quiesce the server from 12:00 AM to 12:59 AM. (12 13) indicates to quiesce the server from 12:00 PM to 1:59 PM, which is a 2 hour period. Specify (-1 -1) to disable the daily quiesce for a particular server.

reach set

Sets the interval, loglevel, and logsize for the reach advisor.

refresh

Set the number of intervals before querying the executor for a refresh of information about new and active connections.

- *refresh cycle*

A positive number representing the number of intervals. The default is 2.

report

Display a statistics snapshot report.

- *cluster*

The address of the cluster you want displayed in the report. The address can be either a symbolic name or in IP address format. The default is a manager report display for all the clusters.

Note: Additional clusters are separated by a plus sign (+).

restart

Restart all servers (that are not down) to normalized weights (1/2 of maximum weight).

- *message*

A message that you want written to the manager log file.

sensitivity

Set minimum sensitivity to which weights update. This setting defines when the manager should change its weighting for the server based on external information.

- *weight*

A number from 1 to 100 to be used as the weight percentage. The default of 5 creates a minimum sensitivity of 5%.

smoothing

Set an index that smooths the variations in weight when load balancing. A higher smoothing index will cause server weights to change less drastically as network conditions change. A lower index will cause server weights to change more drastically.

- *index*

A positive floating point number. The default is 1.5.

start

Start the manager.

- *log file*

File name to which the manager data is logged. Each record in the log is time stamped. The default file is installed in the logs directory. See “Examples” on page 273. To change the directory where the log files are kept, see “Logging with Load Balancer” on page 158.

- *metric_port*

Port that Metric Server will use to report system loads. If you specify a metric port, you must specify a log file name. The default metric port is 10004.

status

Display the current status of all the values in the manager that can be set globally and their defaults.

stop

Stop the manager.

unquiesce

Specify that the manager can begin to give a weight higher than 0 to a server that was previously quiesced, in every port to which it is defined.

- *server*

The IP address of the server as either a symbolic name or in dotted decimal format.

- *cluster@port@server*

This command will unquiesce that particular port for service, but all other ports will remain quiesced.

version

Display the current version of the manager.

Samples

- To set the updating interval for the manager to every 5 seconds:
`dscontrol manager interval 5`
- To set the level of logging to 0 for better performance:
`dscontrol manager loglevel 0`
- To set the manager log size to 1,000,000 bytes:
`dscontrol manager logsize 1000000`
- To specify that no more connections be sent to the server at 130.40.52.153:
`dscontrol manager quiesce 130.40.52.153`
- To quiesce the HTTP service on server 9.40.25.67, which is on cluster1:
`dscontrol manager quiesce cluster1@9.40.25.67@80`
- To set the number of updating intervals before the weights are refreshed to 3:
`dscontrol manager refresh 3`
- To get a statistics snapshot of the manager:
`dscontrol manager report`
- To restart all the servers to normalized weights and write a message to the manager log file:
`dscontrol manager restart Restarting the manager to update code`
- To set the sensitivity to weight changes to 10:
`dscontrol manager sensitivity 10`
- To set the smoothing index to 2.0:
`dscontrol manager smoothing 2.0`

- To start the manager and specify the log file named ndmgr.log (paths cannot be set):
dscontrol manager start ndmgr.log
- To display the current status of the values associated with the manager:
dscontrol manager status
- To stop the manager:
dscontrol manager stop
- To specify that no more new connections be sent to a server at 130.40.52.153 between 2:00 AM and 4:59 PM:
dscontrol manager quiesce 130.40.52.153 daily 2 16
- To specify that the manager can begin to give a weight higher than 0 to a server at 130.40.52.153 that was previously quiesced:
dscontrol manager unquiesce 130.40.52.153
- To display the current version number of the manager:
dscontrol manager version

Related reference:

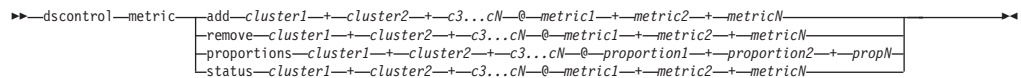
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol metric

You can configure system metrics with the dscontrol metric command.

Syntax



Parameters

add

Add the specified metric.

- *cluster*

The address to which clients connect. The address can be either the host name of the machine, or the IP address notation format. Additional clusters are separated by a plus sign (+).

- *metric*

The system metric name. This must be the name of an executable or script file in the metric server’s script directory.

remove

Remove the specified metric.

proportions

Set the proportions for all the metrics associated with this object.

status

Display the current values of this metric.

Samples

- To add a system metric:
dscontrol metric add site1@metric1

- To set proportions for a sitename with two system metrics:
dscontrol metric proportions site1 0 100
- To display the current status of values associated with the specified metric:
dscontrol metric status site1@metric1

Related reference:

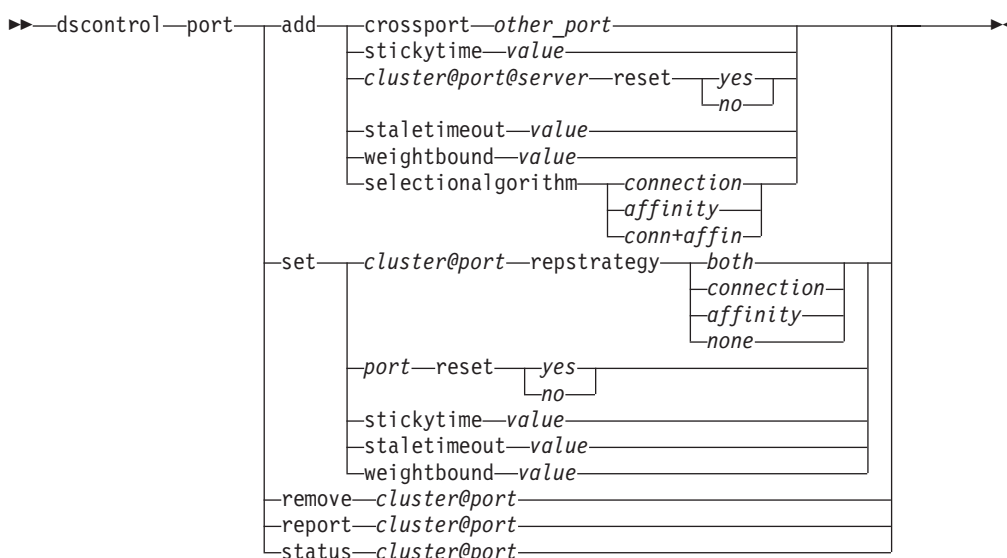
Commands

Browse the subtopics for a quick reference to all the dscontrol commands, syntax diagrams for each, and examples for their use.

dscontrol port

Configure ports and port settings with the dscontrol port command.

Syntax



Parameters

add

Add a port to a cluster. You must add a port to a cluster before you can add any servers to that port. If there are no ports for a cluster, all client requests are processed locally. You can add more than one port at one time using this command.

- **crossport**

Crossport affinity allows you to expand the affinity feature across multiple ports so that client requests received on different ports can still be sent to the same server for subsequent requests. For the crossport value, specify the *other_port* number for which you want to share the cross port affinity feature.

In order to use this feature, the ports must:

- Share the same cluster address
- Share the same servers
- Use the affinity or conn+aff selection algorithm
- Have the same stickytime value, which is not zero

To remove the crossport feature, set the crossport value back to its own port number.

- *other_port*: specifies the value of crossport. The default value is the same as its own port number.

- **stickytime** *value*

The interval between the opening of one connection and the opening of a new connection, during which a client will be sent back to the same server that was used for the first connection. After the stickytime value has elapsed, the client might be sent to a server different from the server that was used for the first connection. If you use affinity as the selection algorithm, the stickytime value is the interval between any two successive packets from the same client.

Note: stickytime is only valid for the conn+aff and affinity selection algorithms.

- *value* is the value of stickytime in seconds.

- **staletimeout**

The number of seconds during which there can be no activity on a connection before that connection is removed. The default value is 900 for port 21 (FTP) and 259,200 for port 23 (Telnet).

Note: staletimeout is only valid for the connection and conn+affin selection algorithms.

- *value*

The value of staletimeout in number of seconds.

weightbound

Set the maximum weight for servers on this port. This affects how much difference there can be between the number of requests the executor will give each server. The default value is 20.

- *weight*

A number from 1–100 representing the maximum weight bound.

selectionalgorithm

Defines the method for selecting the next server.

- **affinity**

Specifies that the server selection is based on client affinity.

- **connection**

Specifies that the server selection is based on simple round-robin selection (default).

- **conn+affin**

Note: Specifies that server selection is based on an existing connection. For new connections, the server selection is based on affinity.

set

Set the fields of a port.

- *cluster@port* **repstrategy** *value*:

Specifies that connection and affinity records will be replicated between high availability partners. When the records are replicated, connection and affinity states are preserved so that the connections can continue even after takeover has taken place.

Value can be:

- *both*

Replicate both connection and affinity records.

- *affinity*
Replicate affinity records only.
- *connection*
Replicate connection records only.
- *none*
No replication.

reset

Activate or deactivate the TCP reset feature. Load Balancer will send a TCP reset to the client when the client has a connection to a server with a weight of 0. A server's weight can be 0 if it is configured to 0 or if an advisor marks the server down. A TCP reset will cause the connection to be immediately closed.

Note: Be aware of the following:

- Load Balancer will not send a TCP reset command when all backend servers are down.
- To use TCP reset, you must set the `clientgateway` for IPv4 addresses, or `clientgateway6` for IPV6 addresses, to a router address. Use the `dscontrol executor` command. `Clientgateway` has to be set for a port that is on IPv4 cluster, and `clientgateway6` has to be set on a port that is on IPv6 cluster.

remove

Remove this port.

report

Report on this port.

status

Show status of servers on this port. If you want to see the status on all ports, do not specify a port with this command, but remember to still include the `@` symbol.

Sample

- To set the selection algorithm for a port:
`dscontrol port add cluster@port selectionalgorithm affinity`
- To add port 80 and 23 to a cluster address 130.40.52.153:
`dscontrol port add 130.40.52.153@80+23`
- To set the maximum weight of 10 to port 80 at a cluster address of 130.40.52.153:
`dscontrol port set 130.40.52.153@80 weightbound 10`
- To set the stickytime value to 60 seconds for port 80 and port 23 at a cluster address of 130.40.52.153:
`dscontrol port set 130.40.52.153@80+23 stickytime 60`
- To set the cross port affinity of port 80 to port 23 at a cluster address of 130.40.52.153:
`dscontrol port add 130.40.52.153@80 crossport 23`
- To activate TCP reset feature for port 80:
`dscontrol port set 80 reset yes`
- To remove port 23 from a cluster address of 130.40.52.153:
`dscontrol port remove 130.40.52.153@23`
- To get the status of port 80 at a cluster address of 9.67.131.153:
`dscontrol port status 9.67.131.153@80`
- To get the report of port 80 at a cluster address of 9.62.130.157:

```
dscontrol port report 9.62.130.157@80
```

- To enable connection record replication for high availability partners:

```
dscontrol port set 130.40.52.153@23 repstrategy connection
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol rule

Control the executor function with the `dscontrol rule` command.

Syntax



The following options are available for this command:



Parameters

add

Add this rule to a port.

- *cluster*: specifies the address of the cluster as either a symbolic name or in IP address format. You can use an at symbol (@) to act as a wild card. For instance, the following command will result in adding RuleA to port 80 for all clusters:

```
dscontrol rule add @80@RuleA type type
```

Separate additional clusters with a plus sign (+).

- *port*: specifies the number of the port. You can use an at symbol (@) to act as a wild card. For instance, the following command will result in adding RuleA to all ports for ClusterA:

```
dscontrol rule add clusterA@@RuleA type type
```

Separate additional ports with a plus sign (+).

- *rule*: specifies the name that you choose for the rule. This name can contain any alphanumeric character, underscore, hyphen, or period. It can be from 1 to 20 characters and cannot contain any blanks. Separate additional rules with a plus sign (+).
- *type value*

- **active**: based on the number of active connections total for the port. This rule will work only if the manager is running.
- **true**: specifies that this rule will always evaluate as true.
- *beginrange*: specifies the lower value in the range used to determine whether or not the rule is true. This is an integer with a default value of 0.
- *endrange*: specifies the higher value in the range used to determine whether or not the rule is true. This is an integer with a default value of 2 to the 31st power minus 1.
- **priority value**: The order in which the rules are reviewed, where *value* is an integer.

If you do not specify the priority of the first rule you add, Load Balancer will set it to 1 by default. When a subsequent rule is added, by default its priority is calculated to be 10 + the current lowest priority of any existing rule. For example, assume you have an existing rule whose priority is 30. You add a new rule and set its priority at 25 (which, remember, is a higher priority than 30). Then you add a third rule without setting a priority. The priority of the third rule is calculated to be 40 (30 + 10).

- **evaluate value**: specifies whether to evaluate the rule's condition across all servers within the port or across servers within the rule. *Value* can be:
 - **port**: specifies to evaluate rule's condition across all the servers on the port. This is the default value.
 - **rule**: specifies to evaluate the rule's condition across the servers within the rule.

Evaluate servers within the rule

The option to measure the rule's condition across the servers within the rule allows you to configure two rules with the following characteristics:

1. The first rule that gets evaluated contains all the servers maintaining the Web site content, and the evaluate option is set to rule (evaluate the rule's condition across the servers within the rule).
2. The second rule is an always true rule that contains a single server that responds with a "site busy" type response.

The result is that when traffic exceeds the threshold of the servers within the first rule, traffic is sent to the "site busy" server within the second rule. When traffic falls below the threshold of the servers within the first rule, new traffic continues once again to the servers in the first rule.

Evaluate servers on the port

Using the two rules described above, if you set the evaluate option to port for the first rule (evaluate rule's condition across all the servers on the port), when traffic exceeds the threshold of that rule, traffic is sent to the "site busy" server associated to the second rule. The first rule measures all server traffic (including the "site busy" server) on the port to determine whether the traffic exceeds the threshold. As congestion decreases for the servers associated to the first rule, an unintentional result may occur where traffic continues to the "site busy" server because traffic on the port still exceeds the threshold of the first rule.

dropserver

Remove a server from a rule set.

- *server*: specifies the name of the server to remove. This is the IP address of the TCP server machine as either a symbolic name or in IP address format. Or, if you used server partitioning, use the logical server's unique name. Separate additional servers with a plus sign (+).

remove

Remove one or more rules, separated from one another by plus signs.

report

Display the internal values of one or more rules.

set

Set values for this rule.

useserver

Insert servers into a rule set.

status

Display the values that are configured of one or more rules.

Samples

- For example, to route a range of connections to a certain cluster and port:
`dscontrol rule add 130.40.52.153@80@pool2 type active beginrange 250 endrange 500`
- Create a rule that always evaluates as true with a priority of 100:
`dscontrol rule add 130.40.52.153@80@jamais type true priority 100`

Related tasks:

“Configure rules to manage traffic to busy or unavailable servers” on page 162
Use rules-based load balancing to fine tune when and why packets are sent to which servers. Load Balancer reviews any rules you add from first priority to last priority, stopping on the first rule that it finds to be true, then load balancing the traffic between any servers associated with the rule. It already balances the load based on the destination and port, but using rules expands your ability to distribute connections.

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol server

Configure servers and modify existing server configurations with the dscontrol server command.

Syntax



Parameters

add

Add this server.

- *cluster*

The address of the cluster as either a symbolic name or in IP address format. You can use an at symbol (@) to act as a wild card. For instance, the following command will result in adding ServerA to port 80 on all clusters:

```
dscontrol server add @80@ServerA
```

Note: Additional clusters are separated by a plus sign (+)

.

- *port*

The number of the port. You can use an at symbol (@) to act as a wild card. For instance, the following command will result in adding ServerA to all clusters on all ports:

```
dscontrol server add @@ServerA
```

Note: Additional ports are separated by a plus sign (+).

- *server*

The server is the unique IP address of the TCP server machine as either a symbolic name or in IP address format. Or, if you use a unique name that does not resolve to an IP address, you must provide the server address parameter on the dscontrol server add command.

Note: Additional servers are separated by a plus sign (+).

- **address**

The unique IP address of the TCP server machine as either a host name or in IP address format. If the server is not able to be resolved, you must provide the address of the physical server machine.

– *address*

Value of the address of the server.

- **advisorrequest**
String
- **advisorresponse**
String
- **collocation** *value*
 - yes
 - no

The default value is no.
- **encapforward** *value*
Specifies to enable encapsulation forwarding. *Value* can be yes or no.

Note: Use encapsulation forwarding when the backend server is not located on the same network segment or if you are using virtualization technology and need to forward packets that are otherwise unable to be forwarded.

- **encaptype** *value*
Specifies the type of encapsulation forwarding. *Value* can be:
 - ipip
 - gre
- **encapcond** *value*
Specifies the conditions in which to enable encapsulation forwarding. *Value* can be:
 - always
 - auto
- **fixedweight**
The fixedweight option allows you to specify whether you want the manager to modify the server weight or not. If you set the fixedweight value to yes, when the manager runs it will not be allowed to modify the server weight. For more information, see “Managing traffic with server weights” on page 179.
 - *value*
Specifies the value of fixedweight. The value can be yes or no. Default is no.
 - *portvalue*
Value of the map port number. The default is the client request's destination port number.
- **NAT** *value*
Specifies whether or not nat is enabled.
 - yes
 - no
- **return address** *value*
The IP address to use as the source address when forwarding client requests to the server by the dispatcher when using NAT forwarding.
 - *address*
Value of the return IP address.
- **router address** *value*
The IP address of the router to the remote server that is used by the dispatcher when using the NAT forwarding method

- *address*
Value of the router IP address.

- **sticky value**

The sticky option allows a server to override the stickytime setting on a specific port. This feature can be useful in certain situations in which you are using rules to control traffic. The sticky option behaves like this:

- With a default value of yes, the server retains the normal affinity that is defined at the port.
- With a value of no, the client will not return to that server the next time it issues a request on that port, regardless of the stickytime setting of the port.

Note: Sticky is valid only for conn+affin port selection algorithm. For more information, see the step on enabling port affinity override in the “Optimize connections with client-to-server affinity for the Dispatcher component” on page 166 topic. *Value* can be:

- yes. This is the default value.
- no

- **weight**

A number from 0–100 (but not to exceed the specified port's weightbound value) representing the weight for this server. Setting the weight to zero will prevent any new requests from being sent to the server, but will not end any currently active connections to that server. The default is one-half the specified port's maximum weightbound value. If the manager is running, this setting will be quickly overwritten.

- *value*
Value of the server weight.

remove

Remove this server.

report

Report on this server. The report contains the following information per server: current number of connections per second (CPS), kilobytes transferred in a one second interval (KBPS), total number of connections (Total), number of connections that are in the active state (Active), number of connections that are in the FIN state (FINed), and number of completed connections (Comp).

set

Set values for this server.

status

Show status of the servers.

Samples

- To add the server at 27.65.89.42 to port 80 on a cluster address 130.40.52.153:
`dscontrol server add 130.40.52.153@80@27.65.89.42`
- To remove the server at 27.65.89.42 on all ports on all clusters:
`dscontrol server remove @@27.65.89.42`
- To allow the HTTP advisor to query an HTTP URL request HEAD / HTTP/1.0 for server 27.65.89.42 on HTTP port 80:
`dscontrol server set 130.40.52.153@80@27.65.89.42 advisorrequest "\"HEAD / HTTP/1.0\""`
- To show the status for server 9.67.143.154 on port 80:
`dscontrol server status 9.67.131.167@80@9.67.143.154`

- To set the server 27.65.89.42 as nonsticky, which is the port affinity override feature:
dscontrol server set 130.40.52.153@80@27.65.89.42 sticky no
- To set the weight to 10 for server 27.65.89.42 at port 80 on cluster address 130.40.52.153:
dscontrol server set 130.40.52.153@80@27.65.89.42 weight 10
- To set the server 27.65.89.42 as nonsticky, which enables the port affinity override feature:
dscontrol server set 130.40.52.153@80@27.65.89.42 sticky no

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

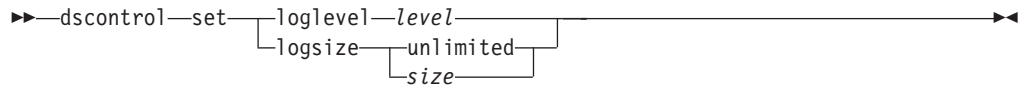
“dscontrol executor” on page 233

Control the executor function with the dscontrol executor command.

dscontrol set

Configure the settings for the server log file with the dscontrol set command.

Syntax



Parameters

loglevel

The level at which the dsserver logs its activities.

- *level*

The default value of **loglevel** is 0. The range is 0–5. The following are the possible values: 0 is None, 1 is Minimal, 2 is Basic, 3 is Moderate, 4 is Advanced, 5 is Verbose.

logsize

The maximum number of bytes to be logged in the log file.

- *size*

The default value of logsize is 1 MB.

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

dscontrol status

You can display status for the managers or advisors with the dscontrol status command.

Syntax



Parameters

There are no parameters for this command.

Sample

To see which managers and advisors are running:

```
dscontrol status
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol advisor

Use the **sscontrol advisor** command to control the advisor.

Syntax

```
▶▶ sscontrol file delete filename.ext
▶▶ sscontrol file appendload filename.ext
▶▶ sscontrol file report
▶▶ sscontrol file save filename.ext
▶▶ sscontrol file logsize name port [cluster@port] [unlimited | number of records]
▶▶ sscontrol file receivetimeout name port [cluster@port] seconds
▶▶ sscontrol file report save filename.ext
▶▶ sscontrol file force newload filename.ext

▶▶ sscontrol connecttimeout name port sitename@port seconds interval name port sitename@port
▶▶ sscontrol seconds list loglevel name port sitename@port level logsize name port sitename@port
▶▶ sscontrol size | unlimited bytes receivetimeout name port sitename@port seconds report name
▶▶ sscontrol port sitename@port retries name port sitename@port numretries start name port
▶▶ sscontrol tename@port log_file status name port sitename@port stop name port sitename@port
▶▶ sscontrol meout name port sitename@port version name port sitename@port seconds
```

Parameters

connecttimeout

Set how long an advisor waits before reporting that a connect to a server for a particular port on a server (a service) fails. For more information, see the Advisor connect timeout and receive timeout for servers topic.

- *name*

The name of the advisor. Possible values include http, https, ftp, sip, ssl, smtp, imap, pop3, ldap, ldaps, nntp, telnet, connect, ping, WLM, and WTE. Names of customized advisors are of the format xxxx, where ADV_xxxx is the name of the class that implements the custom advisor.

- *port*

The number of the port that the advisor is monitoring.

- *seconds*

A positive integer representing the timeout in seconds at which the advisor waits before reporting that a connect to a server fails. The default is 3 times the value specified for the advisor interval.

interval

Set how often the advisor will query the servers for information.

- *seconds*

A positive integer representing the number of seconds between requests to the servers about their current status. The default is 7.

list

Show list of advisors that are currently providing information to the manager.

loglevel

Set the logging level for an advisor log.

- *level*

The number of the level (0 to 5). The default is 1. The higher the number, the more information that is written to the advisor log. The possible values are:

- 0 is None
- 1 is Minimal
- 2 is Basic
- 3 is Moderate
- 4 is Advanced
- 5 is Verbose

logsize

Set the maximum size of an advisor log. When you set a maximum size for the log file, the file will wrap; when the file reaches the specified size, the subsequent entries are written from the top of the file, overwriting the previous log entries. Log size cannot be set smaller than the current size of the log. Log entries are time-stamped so you can tell the order in which they were written. The higher you set the log level, the more carefully you should choose the log size, because you can quickly run out of space when logging at the higher levels.

- *size | unlimited*

The maximum size in bytes for the advisor log file. You can specify either a positive number greater than zero, or unlimited. The log file may not reach the exact maximum size before being overwritten because the log entries vary in size. The default value is 1 MB.

receivetimeout

Set how long an advisor waits before reporting that a receive from a particular port on a server (a service) fails.

- *seconds*

A positive integer representing the timeout in seconds at which the advisor waits before reporting that a receive from a server fails. The default is 3 times the value specified for the advisor interval.

report

Display a report on the state of the advisor.

retries

The number of retries that an advisor can make before marking a server down.

- *numretries*

An integer greater than or equal to zero. This value should be no larger than 3. If retries keyword is not configured, the number of retries defaults to zero.

start

Start the advisor. There are advisors for each protocol. The default ports are as follows:

Table 15. Default ports for advisors

Advisor name	Protocol	Port
connect	n/a	user-defined
db2	private	50000
ftp	FTP	21
http	HTTP	80
https	SSL	443
imap	IMAP	143
ldap	LDAP	389
nntp	NNTP	119
ping	PING	n/a
pop3	POP3	110
sip	SIP	5060
smtp	SMTP	25
ssl	SSL	443
telnet	Telnet	23

name

The advisor name.

sitename@port

The *sitename* value is optional on the advisor commands; however, the port value is required. If the *sitename* value is not specified, the advisor starts running on all available site names configured. If you specify a site name, the advisor starts running for only the site name you specify. Additional site names are separated by a plus sign (+).

- *log file*

File name to which the management data is logged. Each record in the log is time-stamped.

The default file is *advisorname_port.log*, for example, *http_80.log*. To change the directory where the log files are kept, see the Changing the log file paths topic.

You can start only one advisor for each site name.

status

Display the current status of all the values in an advisor that can be set globally and their defaults.

stop

Stop the advisor.

timeout

Set the number of seconds for which the manager will consider information from the advisor as valid. If the manager finds that the advisor information is older than this timeout period, the manager will not use that information in

determining weights for the servers on the port the advisor is monitoring. An exception to this timeout is when the advisor has informed the manager that a specific server is down. The manager will use that information about the server even after the advisor information has timed out.

- *seconds*

A positive number representing the number of seconds, or the word unlimited. The default value is unlimited.

version

Display the current version of the advisor.

Samples

- To set the time (30 seconds) an HTTP advisor (for port 80) waits before reporting that a connect to a server fails:

```
sscontrol advisor connecttimeout http 80 30
```

- To set the interval for the FTP advisor (for port 21) to 6 seconds:

```
sscontrol advisor interval ftp 21 6
```

- To display the list of advisors currently providing information to the manager:

```
sscontrol advisor list
```

This command produces output similar to:

```
-----
| ADVISOR | SITENAME@PORT | TIMEOUT |
-----
| http    |                | 80      | unlimited
| ftp     |                | 21      | unlimited
-----
```

- To change the log level of the http advisor log for the site name of *mysite* to 0 for better performance:

```
sscontrol advisor loglevel http mysite@80 0
```

- To change the ftp advisor log size for the site name of *mysite* to 5000 bytes:

```
sscontrol advisor logsize ftp mysite@21 5000
```

- To set the time (60 seconds) an HTTP advisor (for port 80) waits before reporting that a receive from a server fails:

```
sscontrol advisor receivetimeout http 80 60
```

- To display a report on the state of the ftp advisor (for port 21):

```
sscontrol advisor report ftp 21
```

This command produces output similar to:

```
Advisor Report:
-----
Advisor name ..... http
Port number ..... 80

sitename ..... mySite
Server address ..... 9.67.129.230
Load ..... 8
```

- To start the advisor with the `ftpadv.log` file:

```
sscontrol advisor start ftp 21 ftpadv.log
```

- To display the current status of values associated with the http advisor:

```
sscontrol advisor status http 80
```

This command produces output similar to the following:

```
Advisor Status:
-----
Interval (seconds) ..... 7
```

```

Timeout (seconds) ..... Unlimited
Connect timeout (seconds).....21
Receive timeout (seconds).....21
Advisor log filename ..... Http_80.log
Log level ..... 1
Maximum log size (bytes) ..... Unlimited
Number of retries ..... 0

```

- To stop the http advisor at port 80:
`sscontrol advisor stop http 80`
- To set the timeout value for advisor information to 5 seconds:
`sscontrol advisor timeout ftp 21 5`
- To find out the current version number of the ssl advisor:
`sscontrol advisor version ssl 443`

Related reference:

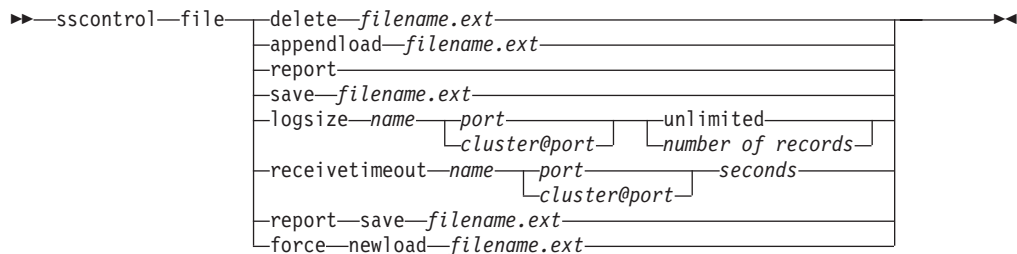
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol file

Use the **sscontrol file** command to manage configuration files.

Syntax



Parameters

delete

Delete the file.

- *file.ext*

File.ext is a configuration file that can be any value and is optional.

appendload

Append a configuration file to the current configuration and load into the Site Selector.

report

Report on the available file or files.

save

Save the current configuration for Site Selector to the file.

Important: Files are saved into and loaded from the following directories:

- /opt/IBM/WebSphere/Edge/ULB/servers/configurations/ss
- <install_root>ss\servers\configurations\component

force

To save your file to an existing file of the same name, use the force option to

delete the existing file before saving the new file. If you do not use the force option, the existing file is not overwritten.

newload

Load a new configuration file into Site Selector. The new configuration file will replace the current configuration.

status

Display the current status of all the values in an advisor that can be set globally and their defaults.

Samples

- To delete a file:

```
sscontrol file delete file3
```

```
File (file3) was deleted.
```

- To load a new configuration file to replace the current configuration:

```
sscontrol file newload file1.sv
```

```
File (file1.sv) was loaded into the Dispatcher.
```

- To append a configuration file to the current configuration and load:

```
sscontrol file appendload file2.sv
```

```
File (file2.sv) was appended to the current configuration and loaded.
```

- To view a report of your files (the files that you saved previously):

```
sscontrol file report
```

```
FILE REPORT:  
file1.save  
file2.sv  
file3
```

- To save your configuration into a file named *file3*:

```
sscontrol file save file3
```

```
The configuration was saved into file (file3).
```

Related reference:

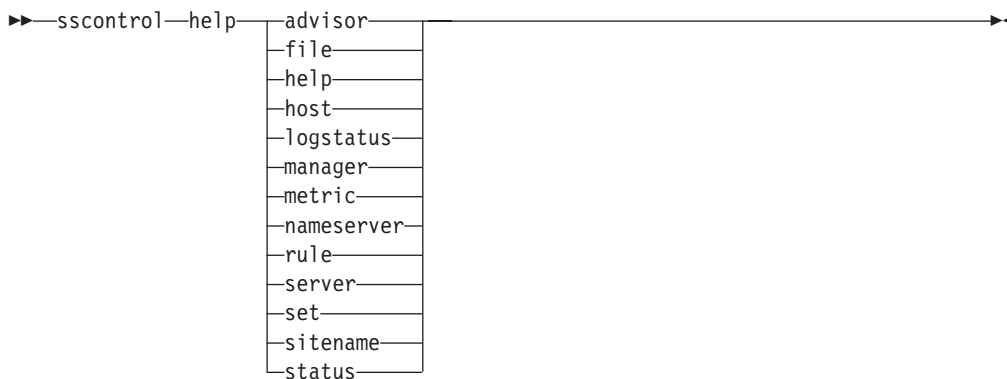
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol help

Use the **sscontrol help** command to display or print help.

Syntax



Samples

- To get help on the **sscontrol** command:

```
sscontrol help
```

This command produces output similar to:

```
HELP COMMAND ARGUMENTS:
```

```
-----
```

```
Usage: help <help option>
```

```
Example: help name
```

```
help           - print complete help text  
advisor        - help on advisor command  
file           - help on file command  
host           - help on host command  
manager        - help on manager command  
metric         - help on metric command  
sitename       - help on sitename command  
nameserver     - help on nameserver command  
rule           - help on rule command  
server         - help on server command  
set            - help on set command  
status         - help on status command  
logstatus      - help on logstatus command
```

Parameters with <> tags are variables.

- The help might display choices for the variables using | to separate the options:

```
logsize <number of bytes | unlimited>
```

```
-Set the maximum number of bytes to be logged in the log file
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol logstatus

Use the **sscontrol logstatus** command to display server log settings.

Syntax



Parameters

logstatus

Displays the server log settings (log file name, logging level, and log size).

Related reference:

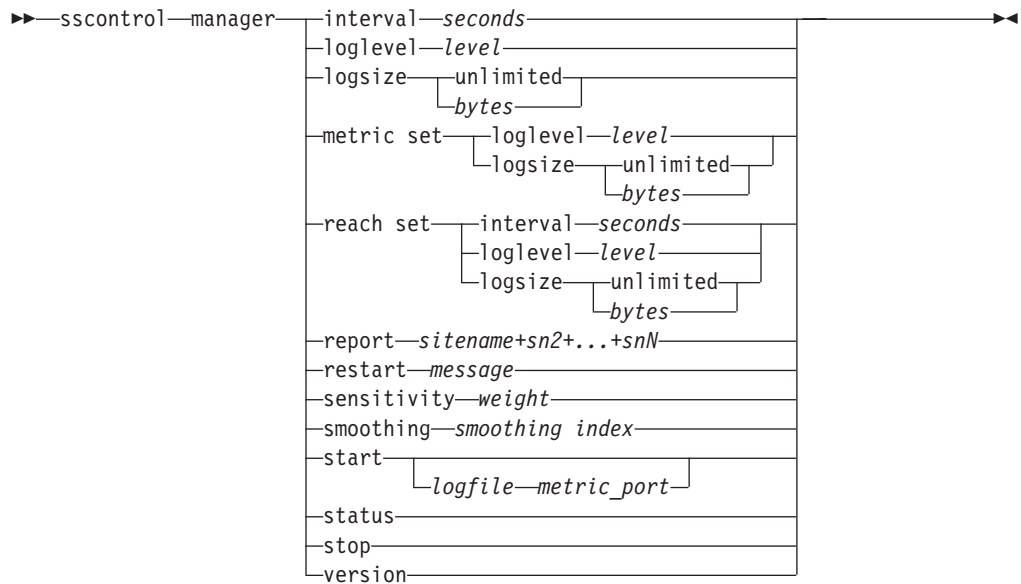
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol manager

Use the **sscontrol manager** command to control the manager.

Syntax



Parameters

interval

Set how often the manager updates the weights of the servers.

interval

Set how often the advisor will query the servers for information.

- *seconds*

A positive number in seconds that represents how often the manager updates weights. The default is 2.

loglevel

Set the logging level for an advisor log.

- *level*

The number of the level (0 to 5). The default is 1. The higher the number, the more information that is written to the advisor log. The possible values are:

- 0 is None
- 1 is Minimal
- 2 is Basic

- 3 is Moderate
- 4 is Advanced
- 5 is Verbose

logsize

Set the maximum size of an advisor log. When you set a maximum size for the log file, the file will wrap; when the file reaches the specified size, the subsequent entries are written from the top of the file, overwriting the previous log entries. Log size cannot be set smaller than the current size of the log. Log entries are time-stamped so you can tell the order in which they were written. The higher you set the log level, the more carefully you should choose the log size, because you can quickly run out of space when logging at the higher levels.

- *bytes*

The maximum size in bytes for the manager log file. You can specify either a positive number greater than zero, or unlimited. The log file may not reach the exact maximum size before being overwritten because the log entries vary in size. The default value is 1 MB.

metric set

Sets the loglevel and logsize for the metric monitor log. The loglevel is the metric monitor logging level (0 - None,1 - Minimal,2 - Basic,3 - Moderate, 4 - Advanced, or 5 - Verbose). The default loglevel is 1. The logsize is the maximum number of bytes to be logged in the metric monitor log file. You can specify either positive number greater than zero, or unlimited. The default logsize is 1.

reach set

Sets the interval, loglevel, and logsize for the reach advisor.

report

Display a statistics snapshot report.

- *sitename*

An integer greater than or equal to zero. This value should be no larger than 3. If retries keyword is not configured, the number of retries defaults to zero.

start

Start the advisor. There are advisors for each protocol. The default ports are as follows:

log file

File name to which the manager data is logged. Each record in the log is time-stamped.

The default file is installed in the logs directory.

metric_port

Port that Metric Server uses to report system loads. If you specify a metric port, you must specify a log file name. The default metric port is 10004.

status

Display the current status and defaults of all the global values in the manager.

stop

Stop the manager.

version

Display the current version of the manager.

Samples

- To set the updating interval for the manager to every 5 seconds:
`sscontrol manager interval 5`
- To set the level of logging to 0 for better performance:
`sscontrol manager loglevel 0`
- To set the manager log size to 1,000,000 bytes:
`sscontrol manager logsize 1000000`
- To get a statistics snapshot of the manager:
`sscontrol manager report`

This command produces output similar to:

```

-----
|          SERVER          | STATUS |
-----
|      9.67.129.221      | ACTIVE |
|      9.67.129.213      | ACTIVE |
|      9.67.134.223      | ACTIVE |
-----

| MANAGER REPORT LEGEND |
-----
| CPU | CPU Load |
| MEM | Memory Load |
| SYS | System Metric |
| NOW | Current Weight |
| NEW | New Weight |
| WT  | Weight |
-----

| mySite | WEIGHT | CPU 49% | MEM 50% | PORT 1% | SYS 0% |
-----
|          | NOW  NEW | WT  LOAD | WT  LOAD | WT  LOAD | WT  LOAD |
-----
| 9.37.56.180 | 10  10 | -99  -1 | -99  -1 | -99  -1 | 0    0 |
-----
|          TOTALS: | 10  10 |          |          |          |          | 0 |
-----

| ADVISOR | SITENAME:PORT | TIMEOUT |
-----
| http    |                80 | unlimited |
-----

```

- To restart all the servers to normalized weights and write a message to the manager log file:
`sscontrol manager restart` Restarting the manager to update code
This command produces output similar to:
320-14:04:54 Restarting the manager to update code
- To set the sensitivity to weight changes to 10:
`sscontrol manager sensitivity 10`
- To set the smoothing index to 2.0:
`sscontrol manager smoothing 2.0`
- To start the manager and specify the log file named `ndmgr.log` (paths cannot be set)
`sscontrol manager start ndmgr.log`
- To display the current status of the values associated with the manager:
`sscontrol manager status`

This command produces output similar to the following:

```
Manager status:
=====
Metric port..... 10004
Manager log filename..... manager.log
Manager log level..... 1
Maximum manager log size (bytes)..... unlimited
Sensitivity level..... 5
Smoothing index..... 1.5
Update interval (seconds)..... 2
Weights refresh cycle..... 2
Reach log level..... 1
Maximum reach log size (bytes)..... unlimited
Reach update interval (seconds)..... 7
```

- To stop the manager:
sscontrol manager stop
- To display the current version number of the manager:
sscontrol manager version

Related reference:

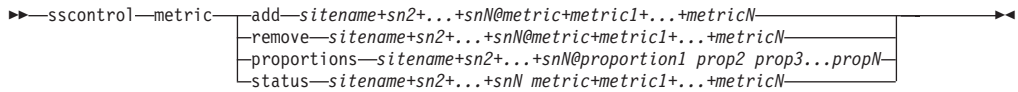
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol metric

Use the sscontrol metric command to configure system metrics.

Syntax



Parameters

add

Add the specified metric.

- *sitename*

The configured site name. Additional site names are separated by a plus sign (+).

- *metric*

The system metric name. This must be the name of an executable or script file in the metric server's script directory.

remove

Remove the specified metric.

proportions

Proportions determines the significance of each metric as compared to the others when they are combined into a single system load for a server.

status

Display the current server values for this metric.

Samples

- To add a system metric:
sscontrol metric add site1@metric1

- To set proportions for a site name with two system metrics:
`sscontrol metric proportions site1 0 100`
- To display the current status of values associated with the specified metric:

```
sscontrol metric status site1@metric1
```

This command produces output similar to:

```
Metric Status:
-----
sitename ..... site1
Metric name ..... metric1
Metric proportion ..... 50
  Server ..... 9.37.56.100
  Metric data .... -1
```

Related reference:

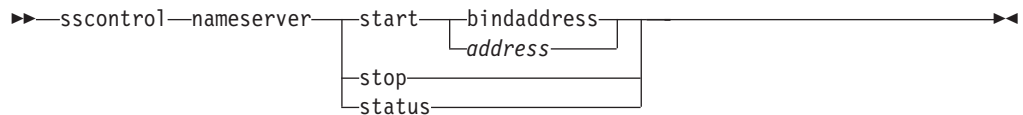
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol nameserver

Use the **sscontrol nameserver** command to control the name server.

Syntax



Parameters

start

Starts the name server.

bindaddress

Starts the name server bound to the specified address. The name server responds only to a request destined for this address.

- *address*

An address (IP or symbolic) configured on the Site Selector machine.

stop

Stops the name server.

status

Displays the status of the name server.

Related reference:

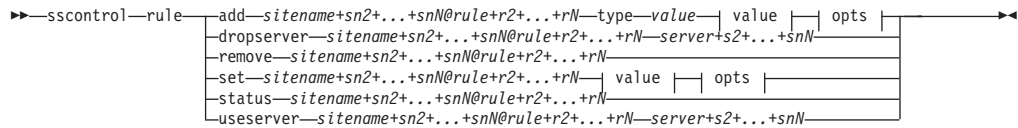
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

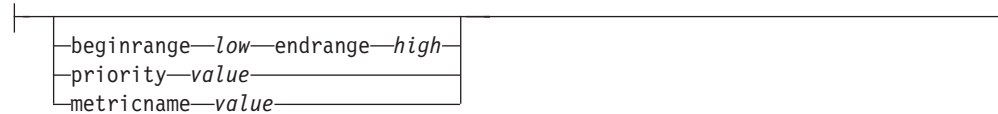
sscontrol rule

Use the **sscontrol rule** command to configure rules.

Syntax



opts:



Parameters

add

Add this rule to a site name.

- *sitename*

An unresolvable hostname that the client will request. The site name must be a fully qualified domain name. Additional site names are separated by a plus sign (+).

Note: An unresolvable hostname that the client will request. The site name must be a fully qualified domain name. Additional site names are separated by a plus sign (+).

- *rule*

The name you choose for the rule. This name can contain any alphanumeric character, underscore, hyphen, or period. It can be from 1 to 20 characters and cannot contain any blanks.

Note: An unresolvable hostname that the client will request. The site name must be a fully qualified domain name. Additional site names are separated by a plus sign (+).

type

The type of rule.

- *type*

Your choices for type are:

ip The rule is based on the client IP address.

metricall

The rule is based on the current metric value for all the servers in the server set.

metricavg

The rule is based on the average of the current metric values for all the servers in the server set.

time The rule is based on the time of day.

true This rule is always true. Think of it as an else statement in programming logic.

beginrange

The lower value in the range used to determine whether or not the rule is true.

- *low*

Depends on the type of rule. The kind of value and its default are listed here by the type of rule:

ip The address of the client as either a symbolic name or in IP address format. The default is 0.0.0.0.

time An integer. The default is 0, representing midnight.

metricall
An integer. The default is 100.

metricavg
An integer. The default is 100.

endrange

The higher value in the range used to determine whether or not the rule is true.

- *high*

Depends on the type of rule. The kind of value and its default are listed here by the type of rule:

ip The address of the client as either a symbolic name or in IP address format. The default is 255.255.255.254.

time An integer. The default is 24, representing midnight.

Note: When defining the beginrange and endrange of time intervals, note that each value must be an integer representing only the hour portion of the time; portions of an hour are not specified. For this reason, to specify a single hour—say, the hour between 3:00 and 4:00 am—you would specify a beginrange of 3 and an endrange also of 3. This will signify all the minutes beginning with 3:00 and ending with 3:59. Specifying a beginrange of 3 and an endrange of 4 would cover the two-hour period from 3:00 through 4:59.

metricall
An integer. The default is 2 to the 32nd power minus 1.

metricavg
An integer. The default is 2 to the 32nd power minus 1.

priority

The order in which the rules are reviewed.

status

Show status of the servers.

- *level*

An integer. If you do not specify the priority of the first rule you add, Site Selector sets it by default to 1. When a subsequent rule is added, by default its priority is calculated to be 10 + the current lowest priority of any existing rule. For example, assume you have an existing rule whose priority is 30. You add a new rule and set its priority at 25 (which is a higher priority than 30). Then you add a third rule without setting a priority. The priority of the third rule is calculated to be 40 (30 + 10).

metricname

Name of the metric measured for a rule.

dropserver

Remove a server from a rule set.

- *server*

The IP address of the TCP server machine as either a symbolic name or in IP address format.

Note: Additional site names are separated by a plus sign (+).

remove

Remove one or more rules, separated from one another by plus signs.

set

Set values for this rule.

status

Display all the values of one or more rules.

useserver

Insert server into a rule set.

Samples

- To add a rule that will always be true, do not specify the beginning range or end range:
`sscontrol rule add sitename@rulename type true priority 100`
- To create a rule forbidding access to a range of IP addresses, in this case those beginning with "9":
`sscontrol rule add sitename@rulename type ip b 9.0.0.0 e 9.255.255.255`
- To create a rule that will specify the use of a given server from the hour of 11:00 a.m. through the hour of 3:00 p.m.:
`sscontrol rule add sitename@rulename type time beginrange 11 endrange 14`
`sscontrol rule useserver sitename@rulename server05`

Related reference:

"Commands" on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol server

Use the sscontrol server command to configure servers.

Syntax



Parameters

add

Add this server

- *sitename*
 An unresolvable hostname that the client requests. The site name must be a fully qualified domain name. Additional site names are separated by a plus sign (+).
- *server*

The IP address of the TCP server machine as either a symbolic name or in IP address format.

Note: Additional servers are separated by a plus sign (+).

weight

A number from 0–100 (not to exceed the specified site name's maximum weightbound value) representing the weight for this server. Setting the weight to zero will prevent any new requests from being sent to the server. The default is one-half the specified site name's maximum weightbound value. If the manager is running, this setting will be quickly overwritten.

- *value*

The server weight value.

down

Mark this server down. This command prevents any other request from being resolved to that server.

remove

Remove this server.

set

Set values for this server.

status

Show status of the servers.

up Mark this server up. Site Selector will now resolve new requests to that server.

Samples

- To add the server at 27.65.89.42 to a site name of *site1*:
sscontrol server add site1@27.65.89.42
- To mark the server at 27.65.89.42 as down:
sscontrol server down site1@27.65.89.42
- To remove the server at 27.65.89.42 for all site names:
sscontrol server remove @27.65.89.42
- To mark the server at 27.65.89.42 as up:
sscontrol server up site1@27.65.89.42

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol set

Use the **sscontrol set** command to configure the server log.

Syntax

►►—sscontrol—set—loglevel—*level*—loglevel—*level*—logsize—*unlimited*—*size*—◄◄

Parameters

loglevel

The level at which the **ssserver** logs its activities.

- *level*
The default value of loglevel is 0. The possible values are:
 - 0 is None
 - 1 is Minimal
 - 2 is Basic
 - 3 is Moderate
 - 4 is Advanced
 - 5 is Verbose

logsize

The maximum number of bytes to be logged in the log file.

- *size*
The default value of log size is 1 MB.

Related reference:

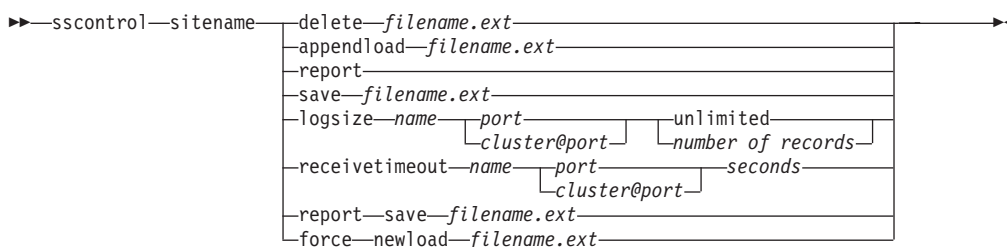
“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol sitename

Use the **sscontrol sitename** command to configure a site name.

Syntax



Parameters

add

Add a new site name.

- *sitename*
An irresolvable host name, requested by the client. Additional site names are separated by a plus sign (+).

cachelife

The amount of time a proximity response is valid and saved in the cache. The default is 1800.

- *value*
A positive number representing the number of seconds a proximity response is valid and saved in the cache.

networkproximity

Determines each server's network proximity to the requesting client. Use this proximity response in the load balancing decision. Set the proximity on or off.

- *value*
The choices are yes or no. The default is no, which means network proximity is turned off.

proportions

Set the proportion of importance for cpu, memory, port (information from any advisors), and system metrics for the Metric Server that are used by the manager to set server weights. Each of these values is expressed as a percentage of the total and the total is always 100.

cpu The percentage of CPU in use on each load balanced server machine (input from Metric Server agent).

memory

The percentage of memory in use (input from Metric Server agent) on each load balanced server

port The input from advisors listening on the port.

system

The input from the Metric Server.

proximitypercentage

Sets the importance of the proximity response versus the health of the server (manager weight).

- *value*

The default is 50.

stickytime

The interval during which a client will receive the same server ID previously returned for the first request. The default value of stickytime is 0, which signifies that the site name is not sticky.

- *time*

A positive, non-zero number representing the number of seconds during which the client receives the same server ID previously returned for the first request.

ttl

Sets the time to live. This indicates how long another name server will cache the resolved response. The default value is 5.

- *value*

A positive number representing the number of seconds the name server will cache the resolved response.

waitforallresponses

Sets whether to wait for all proximity responses from the servers before responding to the client request.

- *value*

The choices are yes or no. The default is yes.

weightbound

A number representing the maximum weight that can be set for servers on this site name. The weightbound value set for the site name may be overridden for individual servers using server weight. The default value of site name weightbound is 20.

- *weight*

The value of weightbound.

set

Set the properties of the site name.

remove

Remove this site name.

status

Show current status of a specific site name.

Samples

- To add a site name:
`sscontrol sitename add 130.40.52.153`
- To turn on network proximity:
`sscontrol sitename set mySite networkproximity yes`
- To set a cache life of 1900000 seconds:
`sscontrol sitename set mySite cachelife 1900000`
- To set a proximity percent of 45:
`sscontrol sitename set mySite proximitypercentage 45`
- To set a site name to not wait for all responses before responding:
`sscontrol sitename set mySite waitforallresponses no`
- To set the time to live to 7 seconds:
`sscontrol sitename set mySite ttl 7`
- To set the proportions of importance for CpuLoad, MemLoad, Port, and System Metric, respectively:
`sscontrol sitename set mySite proportions 50 48 1 1`
- To remove a site name:
`sscontrol sitename remove 130.40.52.153`
- To show the status for site name *mySite*:
`sscontrol sitename status mySite`

This command produces output similar to:

```
SiteName Status:
-----
SiteName ..... mySite
WeightBound ..... 20
TTL ..... 5
StickyTime ..... 0
Number of Servers ..... 1
Proportion given to CpuLoad ..... 49
Proportion given to MemLoad ..... 50
Proportion given to Port ..... 1
Proportion given to System metric .. 0
Advisor running on port ..... 80
Using Proximity ..... N
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

sscontrol status

Use the **sscontrol status** command to display the manager and advisors that are running.

Syntax

►—sscontrol—status—◄

Samples

- To see what is running, type:

```
sscontrol status
```

This command produces output similar to:

```
NameServer has been started.  
Manager has been started.
```

```
-----  
| ADVISOR | SITENAME:PORT | TIMEOUT |  
-----  
| http | 80 | unlimited |  
-----
```

Related reference:

“Commands” on page 226

Look up a command by its name to find detailed syntax and usage of the command.

Examples

This section provides examples of code snippets, command syntax, and configuration values that are relevant to performing tasks with Load Balancer.

To open the information center table of contents to the location of this reference

information, click the **Show in Table of Contents** button () on your information center border.

Example: Sample advisor

This is a sample advisor file called ADV_sample.

```
/ * *  
* ADV_sample: The Load Balancer HTTP advisor  
*  
*  
* This class defines a sample custom advisor for Load Balancer. Like all  
* advisors, this custom advisor extends the function of the advisor base,  
* called ADV_Base. It is the advisor base that actually performs most of  
* the advisor's functions, such as reporting loads back to the Load Balancer  
* for use in the Load Balancer's weight algorithm. The advisor base also  
* performs socket connect and close operations and provides send and receive  
* methods for use by the advisor. The advisor itself is used only for  
* sending and receiving data to and from the port on the server being  
* advised. The TCP methods within the advisor base are timed to calculate  
* the load. A flag within the constructor in the ADV_base overwrites the  
* existing load with the new load returned from the advisor if desired.  
*  
* Note: Based on a value set in the constructor, the advisor base supplies  
* the load to the weight algorithm at specified intervals. If the actual  
* advisor has not completed so that it can return a valid load, the advisor  
* base uses the previous load.  
*  
* NAMING  
*  
* The naming convention is as follows:  
*  
* - The file must be located in the following Load Balancer directory:  
*  
* ulb/servers/lib/CustomAdvisors/ (ulb\servers\lib\CustomAdvisors on Windows)  
*  
* - The Advisor name must be preceded with "ADV_". The advisor can be  
* started with only the name, however; for instance, the "ADV_sample"  
* advisor can be started with "sample".  
*
```

```

* - The advisor name must be in lowercase.
*
* With these rules in mind, therefore, this sample is referred to as:
*
*   <base directory="">/lib/CustomAdvisors/ADV_sample.class
*
*
* Advisors, as with the rest of Load Balancer, must be compiled with the
* prerequisite version of Java. To ensure access to Load Balancer classes, make
* sure that the ibmlb.jar file (located in the lib subdirectory of the base
* directory) is included in the system's CLASSPATH.
*
* Methods provided by ADV_Base:
*
* - ADV_Base (Constructor):
*
*   - Params
*     - String sName = Name of the advisor
*     - String sVersion = Version of the advisor
*     - int iDefaultPort = Default port number to advise on
*     - int iInterval = Interval on which to advise on the servers
*     - String sDefaultName = Unused. Must be passed in as "".
*     - boolean replace = True - replace the load value being calculated
*                           by the advisor base
*                           False - add to the load value being calculated
*                                 by the advisor base
*   - Return
*     - Constructors do not have return values.
*
* Because the advisor base is thread based, it has several other methods
* available for use by an advisor. These methods can be referenced using
* the CALLER parameter passed in getLoad().
*
* These methods are as follows:
*
* - send - Send a packet of information on the established socket connection
*         to the server on the specified port.
*   - Params
*     - String sDataString - The data to be sent in the form of a string
*   - Return
*     - int RC - Whether the data was successfully sent or not: zero indicates
*               data was sent; a negative integer indicates an error.
*
* - receive - Receive information from the socket connection.
*   - Params
*     - StringBuffer sbDataBuffer - The data received during the receive call
*   - Return
*     - int RC - Whether the data was successfully received or not; zero
*               indicates data was received; a negative integer indicates
*               an error.
*
* If the function provided by the advisor base is not sufficient,
* you can create the appropriate function within the advisor and
* the methods provided by the advisor base will then be ignored.
*
* An important question regarding the load returned is whether to apply
* it to the load being generated within the advisor base,
* or to replace it; there are valid instances of both situations.
*
* This sample is essentially the Load Balancer HTTP advisor. It functions
* very simply: a send request--an http head request--is issued. Once a
* response is received, the getLoad method terminates, flagging the advisor
* base to stop timing the request. The method is then complete. The
* information returned is not parsed; the load is based on the time
* required to perform the send and receive operations.
*/

```

```

package CustomAdvisors;
import com.ibm.internet.nd.advisors.*;
public class ADV_sample extends ADV_Base implements ADV_MethodInterface
{
    String COPYRIGHT =
        "(C) Copyright IBM Corporation 1997, All Rights Reserved.\n";
    static final String ADV_NAME = "Sample";
    static final int ADV_DEF_ADV_ON_PORT = 80;
    static final int ADV_DEF_INTERVAL = 7;

    // Note: Most server protocols require a carriage return ("\r") and line
    // feed ("\n") at the end of messages. If so, include them in
    // your string here.

    static final String ADV_SEND_REQUEST =
        "HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
        "IBM_Load_Balancer_HTTP_Advisor\r\n\r\n";

    /**
     * Constructor.
     *
     * Parms: None; but the constructor for ADV_Base has several parameters
     *        that must be passed to it.
     */
    public ADV_sample()
    {
        super( ADV_NAME,
              "2.0.0.0-03.27.98",
              ADV_DEF_ADV_ON_PORT,
              ADV_DEF_INTERVAL,
              "", // not used false);
        super.setAdvisor( this );
    }

    /**
     * ADV_AdvisorInitialize
     *
     * Any Advisor-specific initialization that must take place after the
     * advisor base is started. This method is called only once and is
     * typically not used.
     */
    public void ADV_AdvisorInitialize()
    {
        return;
    }

    /**
     * getLoad()
     *
     * This method is called by the advisor base to complete the advisor's
     * operation, based on details specific to the protocol. In this sample
     * advisor, only a single send and receive are necessary; if more complex
     * logic is necessary, multiple sends and receives can be issued. For
     * example, a response might be received and parsed. Based on the
     * information learned thereby, another send and receive could be issued.
     *
     * Parameters:
     *
     * - iConnectTime - The current load as it refers to the length of time it
     *                   took to complete the connection to the server through
     *                   the specified port.
     *
     * - caller - A reference to the advisor base class where the Load
     *             Balancer-supplied methods are to perform simple TCP requests,
     *             mainly send and receive.
     */

```

```

*
* Results:
*
* - The load - A value, expressed in milliseconds, that can either be added
*   to the existing load, or that can replace the existing load, as
*   determined by the constructor's "replace" flag.
*
* The larger the load, the longer it took the server to respond;
* therefore, the lower the weight will become within the Load Balancer.
*
* If the value is negative, an error is assumed. An error from an
* advisor indicates that the server the advisor is trying to reach is not
* accessible and has been identified as being down. Load Balancer will
* not attempt to load balance to a server that is down. Load Balancer will
* resume load balancing to the server when a positive value is received.
*/
public int getLoad(int iConnectTime, ADV_Thread caller)
{
    int iRc;
    int iLoad = ADV_HOST_INACCESSIBLE; // -1

    // Send tcp request iRc = caller.send(ADV_SEND_REQUEST);
    if (iRc >= 0)
    {
        // Perform a receive
        StringBuffer sbReceiveData = new StringBuffer("");
        iRc = caller.receive(sbReceiveData);

        /**
         * In the normal advisor mode ("replace" flag is false), the load
         * returned is either 0 or 1 indicating the server is up or down.
         * If the receive is successful, a load of zero is returned
         * indicating that the load built within the base advisor is to be used.
         *
         * Otherwise ("replace" flag is true), return the desired load value.
         */

        if (iRc >= 0)
        {
            iLoad = 0;
        }
    }
    return iLoad;
} // End - ADV_sample

```

Example: Implementing custom advisors

The following examples show how custom advisors can be implemented.

Examples for the following types of custom advisors are provided:

- Standard advisor
- Side stream advisor
- Two-port advisor
- WebSphere Application Server (WAS) advisor

Related tasks:

“Creating a custom advisor” on page 125

A custom advisor is a small piece of Java code, provided as a class file, that is called by the Load Balancer base code to determine the load on a server. The base code provides all necessary administrative services, including starting and stopping an instance of the custom advisor, providing status and reports, recording history information in a log file, and reporting advisor results to the manager component.

Example: Using data returned from advisors

Whether you use a standard call to an existing part of the application server or add a new piece of code to be the server-side counterpart of your custom advisor, you possibly want to examine the load values returned and change server behavior.

The Java StringTokenizer class, and its associated methods, make this investigation easy to do. The content of a typical HTTP command might be

```
GET /index.html HTTP/1.0 90
```

A typical response to this command might be the following:

```
HTTP/1.1 200 OK
Date: Mon, 20 November 2000 14:09:57 GMT
Server: Apache/1.3.12 (Linux and UNIX)
Content-Location: index.html.en
Vary: negotiate
TCN: choice
Last-Modified: Fri, 20 Oct 2000 15:58:35 GMT
ETag: "14f3e5-1a8-39f06bab;39f06a02"
Accept-Ranges: bytes
Content-Length: 424
Connection: close
Content-Type: text/html
Content-Language: en

<!DOCTYPE HTML PUBLIC "-//w3c//DTD HTML 3.2 Final//EN">
<HTML><HEAD><TITLE>Test Page</TITLE></HEAD>
<BODY><H1>Apache server</H1>
<HR>
<P><P>This Web server is running Apache 1.3.12.
</P>
<P><IMG SRC="apache_pb.gif" ALT="">
</P></P>
</HR>
</BODY></HTML>
```

The items of interest are contained in the first line, specifically the HTTP return code. The HTTP specification classifies return codes that can be summarized as follows:

- 2xx return codes are successes
- 3xx return codes are redirections
- 4xx return codes are client errors
- 5xx return codes are server errors

If you know precisely what codes the server can possibly return, your code might not need to be as detailed as this example. However, keep in mind that limiting the return codes you detect might limit the future flexibility of your program.

The following example is a stand-alone Java program that contains a minimal HTTP client. The example invokes a simple, general-purpose parser for examining HTTP responses.

```
import java.io.*;
import java.util.*;
import java.net.*;

public class ParseTest {
    static final int iPort = 80;
    static final String sServer = "www.ibm.com";
    static final String sQuery = "GET /index.html HTTP/1.0\r\n\r\n";
    static final String sHTTP10 = "HTTP/1.0";
    static final String sHTTP11 = "HTTP/1.1";

    public static void main(String[] Arg) {
        String sHTTPVersion = null;
        String sHTTPReturnCode = null;
        String sResponse = null; int iRc = 0;
        BufferedReader brIn = null;
        PrintWriter psOut = null;
        Socket soServer= null;
        StringBuffer sbText = new
        StringBuffer(40);

        try {
            soServer = new Socket(sServer, iPort);
            brIn = new BufferedReader(new InputStreamReader(
                soServer.getInputStream()));
            psOut = new PrintWriter(soServer.getOutputStream());
            psOut.println(sQuery);
            psOut.flush();
            sResponse = brIn.readLine();
            try {
                soServer.close();
            } catch (Exception sc) {}
        } catch (Exception swr) {}

        StringTokenizer st = new StringTokenizer(sResponse, " ");
        if (true == st.hasMoreTokens()) {
            sHTTPVersion = st.nextToken();
            if (sHTTPVersion.equals(sHTTP10) || sHTTPVersion.equals(sHTTP11)) {
                System.out.println("HTTP Version: " + sHTTPVersion);
            } else {
                System.out.println("Invalid HTTP Version: " + sHTTPVersion);
            }
        } else {
            System.out.println("Nothing was returned");
            return;
        }

        if (true == st.hasMoreTokens()) {
            sHTTPReturnCode = st.nextToken();
            try {
                iRc = Integer.parseInt(sHTTPReturnCode);
            } catch (NumberFormatException ne) {}

            switch (iRc) {
            case(200):
                System.out.println("HTTP Response code: OK, " + iRc);
                break;
            case(400): case(401): case(402): case(403): case(404):
                System.out.println("HTTP Response code: Client Error, " + iRc);
                break;
            case(500): case(501): case(502): case(503):
                System.out.println("HTTP Response code: Server Error, " + iRc);
                break;
            }
```

```

        default:
            System.out.println("HTTP Response code: Unknown, " + iRc);
            break;
        }
    }

    if (true == st.hasMoreTokens()) {
        while (true == st.hasMoreTokens()) {
            sbText.append(st.nextToken());
            sbText.append(" ");
        }
        System.out.println("HTTP Response phrase: " + sbText.toString());
    }
}
}
}

```

Related reference:

“Example: Sample advisor” on page 132

This is a sample advisor file called ADV_sample.

Example: Implementing a side stream advisor

The following example demonstrates how a side stream advisor can be implemented. This sample illustrates suppressing the standard socket opened by the advisor base. Instead, this advisor opens a side stream Java socket to query a server. This procedure can be useful for servers that use a different port from normal client traffic to listen for an advisor query.

In this example, a server is listening on port 11999 and when queried returns a load value with a hexadecimal int "4". This sample runs in replace mode, that is, the last parameter of the advisor constructor is set to true and the advisor base code uses the returned load value rather than the elapsed time.

Note the call to `supressBaseOpeningSocket()` in the initialization routine. Suppressing the base socket when no data will be sent is not required. For example, you might want to open the socket to ensure that the advisor can contact the server. Examine the needs of your application carefully before making this choice.

```

package CustomAdvisors;
import java.io.*;
import java.net.*;
import java.util.*;
import java.util.Date;
import com.ibm.internet.lb.advisors.*;
import com.ibm.internet.lb.common.*;
import com.ibm.internet.lb.server.SRV_ConfigServer;

public class ADV_sidea extends ADV_Base implements ADV_MethodInterface {
    static final String ADV_NAME = "sidea";
    static final int ADV_DEF_ADV_ON_PORT = 12345;
    static final int ADV_DEF_INTERVAL = 7;

    // create an array of bytes with the load request message
    static final byte[] abHealth = {(byte)0x00, (byte)0x00, (byte)0x00,
                                    (byte)0x04};

    public ADV_sidea() {
        super(ADV_NAME, "3.0.0.0-03.31.00", ADV_DEF_ADV_ON_PORT,
            ADV_DEF_INTERVAL, "",
            true); // replace mode parameter is true
        super.setAdvisor( this );
    }
}

//-----

```

```

// ADV_AdvisorInitialize
public void ADV_AdvisorInitialize()
{
    suppressBaseOpeningSocket();    // tell base code not to open the
                                    // standard socket
    return;
}

//-----
// getLoad
public int getLoad(int iConnectTime, ADV_Thread caller) {
    int iRc;
    int iLoad = ADV_HOST_INACCESSIBLE;    // -1
    int iControlPort = 11999;            // port on which to communicate
                                        // with the server
    String sServer = caller.getCurrentServerId(); // address of server to query
    try {
        socket soServer = new Socket(sServer, iControlPort); // open socket to
                                                            // server
        DataInputStream disServer = new DataInputStream(
                                                    soServer.getInputStream());
        DataOutputStream dosServer = new DataOutputStream(
                                                    soServer.getOutputStream());
        int iRecvTimeout = 10000;        // set timeout (in milliseconds)
                                        // for receiving data
        soServer.setSoTimeout(iRecvTimeout);
        dosServer.writeInt(4);           // send a message to the server
        dosServer.flush();
        iLoad = disServer.readByte();    // receive the response from the server
    } catch (exception e) {
        system.out.println("Caught exception " + e);
    }
    return iLoad;                        // return the load reported from the server
}
}

```

Related reference:

"Example: Sample advisor" on page 132

This is a sample advisor file called ADV_sample.

Example: Implementing standard advisors

The following example demonstrates how to use a standard custom advisor.

This sample source code is similar to the standard Load Balancer HTTP advisor. It functions as follows:

1. A send request, a "HEAD/HTTP" command, is issued.
2. A response is received. The information is not parsed, but the response causes the getLoad method to terminate.
3. The getLoad method returns 0 to indicate success or -1 to indicate a failure.

This advisor operates in normal mode, so the load measurement is based on the elapsed time in milliseconds required to perform the socket open, send, receive, and close operations.

```

package CustomAdvisors;
import com.ibm.internet.lb.advisors.*;
public class ADV_sample extends ADV_Base implements ADV_MethodInterface {
    static final String ADV_NAME = "Sample";
    static final int ADV_DEF_ADV_ON_PORT = 80;
    static final int ADV_DEF_INTERVAL = 7;
    static final String ADV_SEND_REQUEST =
        "HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
        "IBM_Load_Balancer_HTTP_Advisor\r\n\r\n";
}

```

```

//-----
// Constructor
public ADV_sample() {
    super(ADV_NAME, "3.0.0.0-03.31.00",
        ADV_DEF_ADV_ON_PORT, ADV_DEF_INTERVAL, "",
        false);
    super.setAdvisor( this );
}

//-----
// ADV_AdvisorInitialize
public void ADV_AdvisorInitialize() {
    return; // usually an empty routine
}

//-----
// getLoad

public int getLoad(int iConnectTime, ADV_Thread caller) {
    int iRc;
    int iLoad = ADV_HOST_INACCESSIBLE; // initialize to inaccessible

    iRc = caller.send(ADV_SEND_REQUEST); // send the HTTP request to
                                        // the server
    if (0 <= iRc) { // if the send is successful
        StringBuffer sbReceiveData = new StringBuffer(""); // allocate a buffer
                                                            // for the response
    }
    iRc = caller.receive(sbReceiveData); // receive the result

    // parse the result here if you need to

    if (0 <= iRc) { // if the receive is successful
        iLoad = 0; // return 0 for success
    } // (advisor's load value is ignored by
    } // base in normal mode)
    return iLoad;
}
}

```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called `ADV_sample`.

Example: Implementing the WAS advisor

The following examples show how custom advisors can be implemented.

A sample custom advisor for WebSphere Application Server is included in the `install_root/servers/samples/CustomAdvisors/` directory. The full code is not duplicated in this document. Ensure that the following will be implemented:

- `ADV_was.java` is the advisor source code file that is compiled and run on the Load Balancer machine.
- `LBAdvisor.java.servlet` is the servlet source code that must be renamed to `LBAdvisor.java`, compiled, and run on the WebSphere Application Server machine.

The complete advisor is only slightly more complex than the sample. It adds a specialized parsing routine that is more compact than the `StringTokenizer` example shown in the topic “Example: Using data returned from advisors” on page 144.

The more complex part of the sample code is in the Java servlet. Among other methods, the servlet contains two methods required by the servlet specification: `init()` and `service()`, and one method, `run()`, that is required by the `Java.lang.thread` class.

- `init()` is called once by the servlet engine at initialization time. This method creates a thread named `_checker` that runs independently of calls from the advisor and sleeps for a period of time before resuming its processing loop.
- `service()` is called by the servlet engine each time the servlet is invoked. In this case, the method is called by the advisor. The `service()` method sends a stream of ASCII characters to an output stream.
- `run()` contains the core of the code execution. It is called by the `start()` method that is called from within the `init()` method.

The relevant fragments of the servlet code appear below:

```
...
public void init(ServletConfig config) throws ServletException {
    super.init(config);
    ...
    _checker = new Thread(this);
    _checker.start();
}

public void run() {
    setStatus(GOOD);

    while (true) {
        if (!getKeepRunning())
            return;
        setStatus(figureLoad());
        setLastUpdate(new java.util.Date());

        try {
            _checker.sleep(_interval * 1000);
        } catch (Exception ignore) { ; }
    }
}

public void service(HttpServletRequest req, HttpServletResponse res)
    throws ServletException, IOException {
    ServletOutputStream out = null;
    try {
        out = res.getOutputStream();
    } catch (Exception e) { ... }
    ...
    res.setContentType("text/x-application-LBAdvisor");
    out.println(getStatusString());
    out.println(getLastUpdate().toString());
    out.flush(); return;
}
...
```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called `ADV_sample`.

Example: Implementing a two-port advisor

The following example shows how to implement a two-port advisor. This custom advisor sample demonstrates the capability to detect failure for one port of a server based upon both its own status and on the status of a different server daemon that is running on another port on the same server machine.

For example, if the HTTP daemon on port 80 stops responding, you might also want to stop routing traffic to the SSL daemon on port 443.

This advisor is more aggressive than standard advisors, because it considers any server that does not send a response to have stopped functioning, and marks it as down. Standard advisors consider unresponsive servers to be very slow. This advisor marks a server as down for both the HTTP port and the SSL port based on a lack of response from either port.

To use this custom advisor, the administrator starts two instances of the advisor: one on the HTTP port, and one on the SSL port. The advisor instantiates two static global hash tables, one for HTTP and one for SSL. Each advisor tries to communicate with its server daemon and stores the results of this event in its hash table. The value that each advisor returns to the base advisor class depends on both the ability to communicate with its own server daemon and the ability of the partner advisor to communicate with its daemon.

The following custom methods are used.

- `ADV_nte()` is a simple container object to hold information about a server. These objects are stored in the hash table as table elements. Each object has a time stamp that is used to determine whether the element is current.
- `putNte()` and `getNte()` are synchronized methods that ensure that the two advisor instances access the hash table in a controlled fashion.
- `getLoadHTTP` is a method that queries the responsiveness of an HTTP server. It is a low-level routine and does not gather or use information about SSL.
- `getLoadSSL()` is a method that queries the responsiveness of an SSL server. It is a low-level routine and does not gather or use information about HTTP.
- `getLoad()` is the entry point routine for this custom advisor. It can handle both protocols and can store and fetch information from the hash table. This is the routine that links the two ports.

The following error conditions are detected:

- Unresponsive server machine - The base advisor classes periodically send a ping signal to the server address. If the address is not reachable, the base advisor classes marks the server down. Neither of the two instances of the custom advisor is called, and both servers on that machine are marked down.
- One daemon on a server machine becomes unresponsive, but the other is working - When the base code attempts to open a socket with the server, the connection is refused, and the base advisor for this protocol marks the server as down. The custom advisor code for that protocol is not called. Although the custom advisor for the other protocol continues communicating with its server, it learns from the hash table that the other custom advisor cannot communicate with its server daemon. Therefore, the second protocol's advisor also marks its server as down.
- One daemon does not send a response, but the other daemon does - The custom advisor for the unresponsive protocol detects the failure to communicate, marks the server as down, and stores the data in the hash table. The custom advisor for the other port learns that information from the hash table and marks its server as down.

This sample is written to link ports 80 for HTTP and 443 for SSL, but it can be tailored to any combination of ports:

```

package CustomAdvisors;
import java.io.*;
import java.net.*;
import java.util.*;
import java.util.Date;
import com.ibm.internet.lb.advisors.*;
import com.ibm.internet.lb.common.*;
import com.ibm.internet.lb.manager.*;
import com.ibm.internet.lb.server.SRV_ConfigServer;

//-----
// Define the table element for the hash tables used in this custom advisor

class ADV_nte implements Cloneable {
    private String sCluster;
    private int iPort;
    private String sServer;
    private int iLoad;
    private Date dTimestamp;

//-----
// constructor

    public ADV_nte(String sClusterIn, int iPortIn, String sServerIn,
        int iLoadIn) {
        sCluster = sClusterIn;
        iPort = iPortIn;
        sServer = sServerIn;
        iLoad = iLoadIn;
        dTimestamp = new Date();
    }

//-----
// check whether this element is current or expired
    public boolean isCurrent(ADV_twop oThis) {
        boolean bCurrent;
        int iLifetimeMs = 3 * 1000 * oThis.getInterval();    // set lifetime as
                                                            // 3 advisor cycles

        Date dNow = new Date();
        Date dExpires = new Date(dTimestamp.getTime() + iLifetimeMs);

        if (dNow.after(dExpires)) {
            bCurrent = false;
        } else {
            bCurrent = true;
        } return bCurrent;
    }

//-----
// value accessor(s)

    public int getLoadValue() { return iLoad; }

//-----
// clone (avoids corruption between threads)

    public synchronized Object Clone() {
        try {
            return super.clone();
        } catch (cloneNotSupportedException e) {
            return null;
        }
    }
}

//-----

```



```

// define the custom advisor

public class ADV_twop extends ADV_Base
    implements ADV_MethodInterface, ADV_AdvisorVersionInterface {
    static final int ADV_TWOP_PORT_HTTP = 80;
    static final int ADV_TWOP_PORT_SSL = 443;

    //-----
    // define tables to hold port-specific history information

    static Hashtable htTwopHTTP = new Hashtable();
    static Hashtable htTwopSSL = new Hashtable();
    static final String ADV_TWOP_NAME = "twop";
    static final int ADV_TWOP_DEF_ADV_ON_PORT = 80;
    static final int ADV_TWOP_DEF_INTERVAL = 7;
    static final String ADV_HTTP_REQUEST_STRING =
        "HEAD / HTTP/1.0\r\nAccept: */*\r\nUser-Agent: " +
        "IBM_LB_Custom_Advisor\r\n\r\n";

    //-----
    // create byte array with SSL client hello message

    public static final byte[] abClientHello = {
        (byte)0x80, (byte)0x1c,
        (byte)0x01, // client hello
        (byte)0x03, (byte)0x00, // SSL version
        (byte)0x00, (byte)0x03, // cipher spec len (bytes)
        (byte)0x00, (byte)0x00, // session ID len (bytes)
        (byte)0x00, (byte)0x10, // challenge data len (bytes)
        (byte)0x00, (byte)0x00, (byte)0x03, // cipher spec
        (byte)0x1A, (byte)0xFC, (byte)0xE5, (byte)0x20, // challenge data
        (byte)0xFD, (byte)0x3A, (byte)0x3C, (byte)0x18,
        (byte)0xAB, (byte)0x67, (byte)0xB0, (byte)0x52,
        (byte)0xB1, (byte)0x1D, (byte)0x55, (byte)0x44, (byte)0x0D, (byte)0x0A };

    //-----
    // constructor

    public ADV_twop() {
        super(ADV_TWOP_NAME, VERSION, ADV_TWOP_DEF_ADV_ON_PORT,
            ADV_TWOP_DEF_INTERVAL, "",
            false); // false = load balancer times the response
        setAdvisor ( this );
    }

    //-----
    // ADV_AdvisorInitialize

    public void ADV_AdvisorInitialize() {
        return; }

    //-----
    // synchronized PUT and GET access routines for the hash tables

    synchronized ADV_nte getNte(Hashtable ht, String sName, String sHashKey) {
        ADV_nte nte = (ADV_nte)(ht.get(sHashKey));
        if (null != nte) {
            nte = (ADV_nte)nte.clone();
        }
        return nte;
    }
    synchronized void putNte(Hashtable ht, String sName, String sHashKey,
        ADV_nte nte) { ht.put(sHashKey,nte); return;
    }

    //-----

```

```

// getLoadHTTP - determine HTTP load based on server response

int getLoadHTTP(int iConnectTime, ADV_Thread caller) {
    int iLoad = ADV_HOST_INACCESSIBLE;
    int iRc = caller.send(ADV_HTTP_REQUEST_STRING); // send request message
                                                    // to server
    if (0 <= iRc) { // did the request return a failure?
        StringBuffer sbReceiveData = new StringBuffer("") // allocate a buffer
                                                    // for the response
        iRc = caller.receive(sbReceiveData); // get response from server

        if (0 <= iRc) { // did the receive return a failure?
            if (0 < sbReceiveData.length()) { // is data there?
                iLoad = SUCCESS; // ignore retrieved data and
                                // return success code
            }
        }
    }
    return iLoad;
}

//-----
// getLoadSSL() - determine SSL load based on server response

int getLoadSSL(int iConnectTime, ASV_Thread caller) {
    int iLoad = ADV_HOST_INACCESSIBLE;
    int iRc;

    CMNByteArrayWrapper cbawClientHello = new CMNByteArrayWrapper(
                                                abClientHello);
    Socket socket = caller.getSocket();

    try {
        socket.getOutputStream().write(abClientHello); // Perform a receive.
        socket.getInputStream().read(); // If receive is successful,
                                        // return load of 0. We are not
                                        // concerned with data's contents,
                                        // and the load is calculated by
                                        // the ADV_Thread thread.

        iLoad = 0;
    } catch (IOException e) { // Upon error, iLoad will default to it.
    }
    return iLoad;
}

//-----
// getLoad - merge results from the HTTP and SSL methods

public int getLoad(int iConnectTime, ADV_Thread caller) {
    int iLoadHTTP;
    int iLoadSSL;
    int iLoad;
    int iRc;

    String sCluster = caller.getCurrentClusterId(); // current cluster address
    int iPort = getAdviseOnPort();
    String sServer = caller.getCurrentServerId();
    String sHashKey = sCluster + ":" + sServer; // hash table key

    if (ADV_TWOP_PORT_HTTP == iPort) { // handle an HTTP server
        iLoadHTTP = getLoadHTTP(iConnectTime, caller); // get the load for HTTP

        ADV_nte nteHTTP = newADV_nte(sCluster, iPort, sServer, iLoadHTTP);
        putNte(htTwopHTTP, "HTTP", sHashKey, nteHTTP); // save HTTP load
                                                        // information
    }
}

```

```

ADV_nte nteSSL = getNte(htTwopSSL, "SSL", sHashKey); // get SSL
                                                    // information
if (null != nteSSL) {
    if (true == nteSSL.isCurrent(this)) {          // check the time stamp
        if (ADV_HOST_INACCESSIBLE != nteSSL.getLoadValue()) { // is SSL
                                                                // working?
            iLoad = iLoadHTTP;
        } else { // SSL is not working, so mark the HTTP server down
            iLoad= ADV_HOST_INACCESSIBLE;
        }
    } else { // SSL information is expired, so mark the
              // HTTP server down
        iLoad = ADV_HOST_INACCESSIBLE;
    }
} else { // no load information about SSL, report
        // getLoadHTTP() results
    iLoad = iLoadHTTP;
}
}
else if (ADV_TWOP_PORT_SSL == iPort) { // handle an SSL server
    iLoadSSL = getLoadSSL(iConnectTime, caller); // get load for SSL

    ADV_nte nteSSL = new ADV_nte(sCluster, iPort, sServer, iLoadSSL);
    putNte(htTwopSSL, "SSL", sHashKey, nteSSL); // save SSL load info.

    ADV_nte nteHTTP = getNte(htTwopHTTP, "SSL", sHashKey); // get HTTP
                                                            // information
    if (null != nteHTTP) {
        if (true == nteHTTP.isCurrent(this)) { // check the timestamp
            if (ADV_HOST_INACCESSIBLE != nteHTTP.getLoadValue()) { // is HTTP
                                                                    // working?
                iLoad = iLoadSSL;
            } else { // HTTP server is not working, so mark SSL down
                iLoad = ADV_HOST_INACCESSIBLE;
            }
        } else { // expired information from HTTP, so mark SSL down
            iLoad = ADV_HOST_INACCESSIBLE;
        }
    } else { // no load information about HTTP, report
              // getLoadSSL() results
        iLoad = iLoadSSL;
    }
}
}

//-----
// error handler

else {
    iLoad = ADV_HOST_INACCESSIBLE;
}
return iLoad;
}
}

```

Related reference:

“Example: Sample advisor” on page 132
This is a sample advisor file called ADV_sample.

Glossary

ACK A control bit (acknowledge) occupying no sequence space, which indicates that the acknowledgment field of this segment specifies the next sequence number the sender of this segment is expecting to receive, hence acknowledging receipt of all previous sequence numbers.

address

The unique code assigned to each device or workstation connected to a network. A standard IPv4 address is a 32-bit address field containing two parts. The first part is the network address, and the second part is the host number. An IPv6 address is a 128-bit address field that supports a much higher number of addresses than IPv4, and IPv6 addresses also support additional features like multicast and anycast addressing.

advisor

Advisors collect and analyze feedback from individual servers, and inform the manager function.

agent In systems management, a user that, for a particular interaction, has assumed an agent role.

An entity that represents one or more managed objects by (a) emitting notifications regarding the objects and (b) handling requests from managers for management operations to modify or query the objects.

alias An additional name assigned to a server. The alias makes the server independent of the name of its host machine. The alias must be defined in the domain name server.

API Application programming interface. The interface (calling conventions) by which an application program accesses operating system and other services. An API is defined at source code level and provides a level of abstraction between the application and the kernel (or other privileged utilities) to ensure the portability of the code.

backup

In a high availability configuration, this is the partner of the primary machine. It monitors the status of the primary machine and takes over if necessary. See also high availability, primary.

bandwidth

The difference between the highest and lowest frequencies of a transmission channel. This is the amount of data that can be sent through a given communication circuit per second.

begin range

In rules-based load balancing, a lower value specified on a rule. The default for this value depends on the type of rule.

binary logging

Allows server information to be stored in binary files, and then be processed to analyze the server information that is gathered over time.

Caching proxy

A caching proxy server that can help speed up end-user response time through highly-efficient caching schemes. Flexible PICS filtering helps network administrators control access to Web-based information at one central location.

CGI Common Gateway Interface. A standard for the exchange of information between a Web server and an external program. The external program can be written in any language supported by the operating system, and performs tasks not usually done by the server, such as forms processing.

CGI script

A CGI program written in a scripting language such as Perl or REXX that uses the Common Gateway Interface to perform tasks not usually done by the server, such as forms processing.

client A computer system or process that requests a service of another computer system or process. For example, a workstation or personal computer requesting HTML documents from a Lotus Domino® Go Webserver is a client of that server.

cluster

A group of TCP or UDP servers that are used for the same purpose and are identified by a single hostname. See also cell.

cluster address

The address to which clients connect.

clustered server

A server that Load Balancer groups with other servers into a single, virtual server. Load Balancer balances TCP or UDP traffic among these clustered servers.

consultant

Collects server metrics from the servers that are being load balanced and sends server weight information to the switch that performs the load balancing.

controller

A collection of one or more consultants.

cross port affinity

Cross port affinity is the affinity (sticky) feature expanded to cover across multiple ports. See also sticky time.

daemon

Disk And Execution Monitor. A program that is not involved explicitly, but lies dormant waiting for some condition(s) to occur. The idea is that the perpetrator of the condition need not be aware that a daemon is lurking (though often a program will commit an action only because it knows that it will implicitly invoke a daemon).

default

A value, attribute, or option that is assumed when none is explicitly specified.

destination address

The address of the high availability partner machine to which heartbeats and responses are sent.

Dispatcher

A component of Load Balancer that efficiently balances TCP or UDP traffic among groups of individual linked servers. The Dispatcher machine is the server running the Dispatcher code.

domain name server

DNS. A general-purpose distributed, replicated, data query service chiefly used on Internet for translating hostnames into Internet addresses. Also, the style of hostname used on the Internet, though such a name is properly

called a fully qualified domain name. DNS can be configured to use a sequence of name servers, based on the domains in the name being looked for, until a match is found.

dotted-decimal notation

The syntactical representation for a 32-bit integer that consists of four 8-bit numbers, written in base 10 and separated by periods (dots). It is used to represent IPv4 addresses.

dscontrol

Provides the interface to the Dispatcher component of Load Balancer.

dsserver

Handles the requests from the command line to the executor, manager, and advisors.

end range

In rules-based load balancing, a higher value specified on a rule. The default for this value depends on the type of rule.

Ethernet

A standard type of local area network (LAN). It allows multiple stations to access the transmission medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by using collision detection and transmission. Software protocols used by Ethernet systems vary, but include TCP/IP.

executor

One of several functions. The executor routes requests to the TCP or UDP servers, and also monitors the number of new, active, and finished connections and does garbage collection of completed or reset connections. The executor supplies the new and active connections to the manager function.

FIN A control bit (finis) occupying one sequence number, which indicates that the sender will send no more data or control occupying sequence space.

FIN state

The status of a transaction that has finished. When a transaction is in FIN state, the garbage collector can clear the memory reserved for the connection.

Firewall

A computer that connects a private network, such as a business, to a public network, such as the Internet. It contains programs that limit the access between two networks. See also proxy gateway.

FQDN

Fully Qualified Domain Name. The full name of a system, consisting of its local hostname and its domain name, including a top-level domain (tld). For example, "venera" is a hostname and "venera.isi.edu" is an FQDN. An FQDN should be sufficient to determine a unique Internet address for any host on the Internet. This process, called "name resolution", uses the Domain Name System (DNS).

FTP (File Transfer Protocol)

An application protocol used for transferring files to and from network

computers. FTP requires a user ID and sometimes a password to allow access to files on a remote host system.

gateway

A functional unit that interconnects two computer networks with different architectures.

GRE Generic Routing Encapsulation. A protocol which allows an arbitrary network protocol A to be transmitted over any other arbitrary protocol B, by encapsulating the packets of A within GRE packets, which in turn are contained within packets of B.

heartbeat

A simple packet sent between two machines in high availability mode used by the standby machine to monitor the health of the active machine.

high availability

A feature in which one Load Balancer machine can take over the function of another, if the primary machine is no longer available.

host A computer, connected to a network, that provides an access point to that network. A host can be a client, a server, or both simultaneously.

host name

The symbolic name assigned to a host. Host names are resolved to IP addresses through a domain name server.

HTML (Hypertext Markup Language)

The language that is used to create hypertext documents. Hypertext documents include links to other documents that contain additional information about the highlighted term or subject. HTML controls the format of text and position of form input areas, for example, as well as the navigable links.

HTTP (Hypertext Transfer Protocol)

The protocol used to transfer and display hypertext documents.

HTTPS (Hypertext Transfer Protocol, Secure)

The protocol used to transfer and display hypertext documents using SSL.

ICMP Internet Control Message Protocol. A message control and error-reporting protocol between a host server and a gateway to the Internet.

IMAP Internet Message Access Protocol. A protocol allowing a client to access and manipulate electronic mail messages on a server. It permits manipulation of remote message folders (mailboxes), in a way that is functionally equivalent to local mailboxes.

Internet

The worldwide collection of interconnected networks that use the Internet suite of protocols and permit public access.

intranet

A secure, private network that integrates Internet standards and applications (such as Web browsers) with an organization's existing computer networking infrastructure.

IP Internet Protocol. A connectionless protocol that routes data through a

network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical layer.

IP address

Internet Protocol address. The unique address that specifies the actual location of each device or workstation in a network. It is also known as an Internet address.

IPSEC Internet Protocol Security. A developing standard for security at the network or packet processing layer of network communication.

LAN Local Area Network. A computer network of devices connected within a limited geographical area for communication and which can be connected to a larger network.

loopback alias

An alternative IP address associated with the loopback interface. The alternative address has the useful side affect of not advertising on a real interface.

loopback interface

An interface that bypasses unnecessary communications functions when the information is addressed to an entity within the same system.

MAC address

Media Access Control address. The hardware address of a device connected to a shared network medium.

managed node

In Internet communications, a workstation, server, or router that contains a network management agent. In the Internet Protocol (IP), the managed node usually contains a Simple Network Management Protocol (SNMP) agent.

manager

Sets weights of servers based on internal counters in the executor and feedback that is provided by the advisors. The executor then uses the weights to perform load balancing.

mark down

To break all active connections to a server and stop any new connections or packets from being sent to that server.

mark up

To allow a server to receive new connections.

metric A process or command that returns a numeric value that can be used in load balancing on the network, for example, the number of users currently logged on.

metric address

The address where the metric server connects.

metric collector

Resides in the consultant and is responsible for collecting a metric or metrics.

Metric Server

Formerly known as Server Monitor Agent (SMA). Metric server provides system specific metrics to the manager.

MIB Management Information Base. A collection of objects that can be accessed by means of a network management protocol.

A definition for management information that specifies the information available from a host or gateway and the operations allowed.

netmask

For IPv4, a 32-bit mask used to identify the subnetwork address bits in the host portion of an IP address.

network

Hardware and software data communication system. Networks are often classified according to their geographical extent, local area network (LAN), metropolitan area network (MAN), wide area network (WAN) and also according to the protocols used.

Network Address Translation

NAT, or Network Address Translator, Virtual LAN. A hardware device currently being developed and used to extend the Internet addresses already in use. It allows duplicate IP addresses to be used within a corporation and unique addresses outside.

Network Address Port Translation

NAPT, also known as port mapping. This allows you to configure multiple server daemons within one physical server to listen on different port numbers.

network management station

In the Simple Network Management Protocol (SNMP), a station that runs management application programs that monitor and control network elements.

network proximity

The proximity of two networked entities, such as a client and server, which determines by measuring round-trip time.

NFA (nonforwarding address)

The primary IP address of the machine, used for administration and configuration.

NIC Network Interface Card. An adapter circuit board installed in a computer to provide a physical connection to a network.

NNTP Network News Transfer Protocol. A TCP/IP protocol for transferring news items.

packet The unit of data that is routed between an origin and a destination on the Internet or any other packet-switched network.

PICS Platform for Internet Content Selection. PICS-enabled clients allow the users to determine which rating services they want to use and, for each rating service, which ratings are acceptable and which are unacceptable.

ping A command that sends Internet Control Message Protocol (ICMP) echo-request packets to a host, gateway, or router with the expectation of receiving a reply.

POP3 Post Office Protocol 3. A protocol used for exchanging network mail and accessing mailboxes.

port A number that identifies an abstracted communication device. Web servers use port 80 by default.

primary In a high availability configuration, the machine that starts out as the machine actively routing packets. Its partner, the backup machine, monitors the status of the primary machine and takes over if necessary. See also backup, high availability.

priority In rules-based load balancing, the level of importance placed upon any given rule. The evaluates rules from the first priority level to the last priority level.

private network A separate network on which Load Balancer communicates with clustered servers for performance reasons.

protocol The set of rules governing the operation of functional units of a communication system if communication is to take place. Protocols can determine low-level details of machine-to-machine interfaces, such as the order in which bits from a byte are sent; they can also determine high-level exchanges between application programs, such as file transfer.

Quality of Service (QoS) The performance properties of a network service, including throughput, transit delay and priority. Some protocols allow packets or streams to include QoS requirements.

quiesce To end a process by allowing operations to complete normally.

reach An advisor that issues pings to a given target and reports whether that target is responding.

reach address In a high availability configuration, the address of the target to which the advisor should issue pings to see if the target is responding.

return address A unique IP address or hostname. It is configured on the Load Balancer machine and used as its source address when load balancing the client's request to the server.

RMI Remote Method Invocation. Part of the Java programming language library which enables a Java program running on one computer to access the objects and methods of another Java program running on a different computer.

root user The unrestricted authority to access and modify any part of the AIX, Red Hat Linux, or Solaris operating system, usually associated with the user who manages the system.

route The path of network traffic from origin to destination.

router A device which forwards packets between networks. The forwarding decision is based on network layer information and routing tables, often constructed by routing products.

RPM Red Hat Package Manager.

rule In rules-based load balancing, a mechanism for grouping servers such that a server can be chosen based on information other than the destination address and port.

rule type

In rules-based load balancing, an indicator of the information that should be evaluated to determine whether a rule is true.

scalable

Pertaining to the capability of a system to adapt readily to a greater or lesser intensity of use, volume, or demand. For example, a scalable system can efficiently adapt to work with larger or smaller networks performing tasks of varying complexity.

server A computer that provides shared services to other computers over a network; for example, a file server, a print server, or a mail server.

server address

The unique code assigned to each computer that provides shared services to other computers over a network; for example, a file server, a print server, or a mail server. The server address can be either the IP address or the host name.

server machine

A server that Load Balancer groups with other servers into a single, virtual server. Load Balancer balances traffic among the server machines. Synonymous with clustered server.

service

A function provided by one or more nodes; for example, HTTP, FTP, Telnet.

shell The software that accepts and processes command lines from a user's workstation. The bash shell is one of several UNIX shells available.

site name

A site name is an unresolvable host name that the client will request. For example, a web site has 3 servers (1.2.3.4, 1.2.3.5, and 1.2.3.6) configured for site name *www.dnsload.com*. When a client requests this site name, one of the three server IP addresses will be returned as the resolution. The site name must be a fully qualified domain name, for example: *dnsload.com*. An unqualified name, for example, *dnsload* is invalid for a site name.

Site Selector

A DNS-based load balancing component of . Site Selector balances the load on servers within a wide area network (WAN) using measurements and weights that are gathered from the Metric Server component running on those servers.

SMTP Simple Mail Transfer Protocol. In the Internet suite of protocols, an application protocol for transferring mail among users in the Internet environment. SMTP specifies the mail exchange sequences and message format. It assumes that the Transmission Control Protocol (TCP) is the underlying protocol.

SNMP

Simple Network Management Protocol. The Internet standard protocol, defined in STD 15, RFC 1157, developed to manage nodes on an IP network. SNMP is not limited to TCP/IP. It can be used to manage and monitor all sorts of equipment including computers, routers, wiring hubs, toasters and jukeboxes.

source address

In a high availability configuration, the address of the high availability partner machine that sends heartbeats.

SPARC

Scalable processor architecture.

sscontrol

Provides the interface to the Site Selector component of .

SSL Secure Sockets Layer. A popular security scheme developed by Netscape Communications Corp. along with RSA Data Security Inc. SSL allows the client to authenticate the server and all data and requests to be encrypted. The URL of a secure server protected by SSL begins with https (rather than HTTP).

sticky time

The interval between the closing of one connection and the opening of a new connection during which a client will be sent back to the same server used during the first connection. After the sticky time, the client may be sent to a server different from the first.

strategy

In a high availability configuration, a keyword for specifying how recovery takes place following the failure of the active machine.

subnet mask

For IPv4, a 32-bit mask used to identify the subnetwork address bits in the host portion of an IP address.

SYN A control bit in the incoming segment, occupying one sequence number, used at the initiation of a connection, to indicate where the sequence numbering will start.

TCP Transmission Control Protocol. A communications protocol used on the Internet. TCP provides reliable host-to-host exchange of information. It uses IP as the underlying protocol.

TCP/IP

Transmission Control Protocol/Internet Protocol. A suite of protocols designed to allow communication between networks regardless of the communication technologies used in each network.

TCP server machine

A server that links with other servers into a single, virtual server. balances TCP traffic among the TCP server machines. Synonymous with clustered server.

Telnet Terminal emulation protocol, a TCP/IP application protocol for remote connection service. Telnet allows a user at one site to gain access to a remote host as if the user's workstation were connected directly to that remote host.

timeout

The time interval allotted for an operation to occur.

TOS Type of service. A one byte field in the IP header of the SYN packet.

TTL A DNS TTL (time to live) is the number of seconds a client can cache the name resolution response.

UDP User Datagram Protocol. In the Internet suite of protocols, a protocol that provides unreliable, connectionless datagram service. It enables an application program on one machine or process to send a datagram to an application program on another machine or process. UDP uses the Internet Protocol (IP) to deliver datagrams.

URI Universal Resource Identifier. The encoded address for any resource on the Web, such as HTML document, image, video clip, program, and so forth.

URL Uniform Resource Locator. A standard way of specifying the location of an object, typically a web page, on the Internet. URLs are the form of address used on the World-Wide Web. They are used in HTML documents to specify the target of a hyperlink which is often another HTML document (possibly stored on another computer).

VPN Virtual Private Network (VPN). A network comprised of one or more secure IP tunnels connecting two or more networks.

WAN Wide Area Network. A network that provides communication services to a geographic area larger than that served by a local area network or a metropolitan area network, and that may use or provide public communication facilities.

WAP Wireless Application Protocol. This is an open international standard for applications that use wireless communication. For example, this standard includes internet access from a mobile phone.

WAS WebSphere Application Server.

Web The network of HTTP servers that contain programs and files, many of them hypertext documents that contain links to other documents on HTTP servers. Also World Wide Web.

wizard

A dialog within an application that uses step-by-step instructions to guide a user through a specific task.

WLM Workload Manager. An advisor that is provided with Load Balancer. It is designed to work only in conjunction with servers on OS/390® mainframes running the MVS™ Workload Manager (WLM) component.

